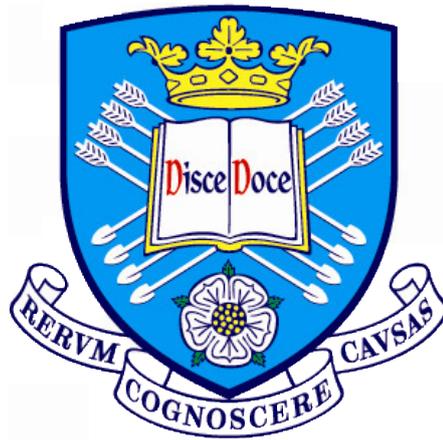


Physics-informed machine learning for structural dynamics: combining physics and data for offshore structures



A Thesis submitted to the University of Sheffield
for the degree of Doctor of Philosophy in the Faculty of Engineering

by

D. J. Pitchforth

Department of Mechanical Engineering

University of Sheffield

March 2025

ACKNOWLEDGEMENTS

Firstly, I would like to thank Professor Lizzy Cross, who has been an excellent supervisor and a role model throughout the thesis. Without her support and guidance, this thesis would not have been possible. I'm grateful for all you have taught me. I would also like to thank Professor Keith Worden for his valuable insight in to wave theory, NARX models and filtering methods amongst many other topics. Additionally, I thank Dr Tim Rogers for his endless generation of great research ideas, and always being available to 'take a quick look at some code'.

For the experimental component of this thesis, the design and manufacturing efforts of Dr Robin Mills and Michael Dutchman at the Laboratory for Verification and Validation are greatly appreciated. A key driver of the field of wave loading prediction is that wave loads are challenging to measure; Robin and Michael's bespoke force collar took a great deal of skill, patience and hard work to enable this. Thank you both.

I wish to extend thanks to everyone in the DRG for creating an enjoyable, welcoming and inspiring place to work. It is a pleasure to be a part of the group, and I'm grateful to all its members. Particular thanks to RD14a: Matt, Sam, Max, Aidan, Tina and Dan for creating such a nice office environment, and many forms of expertise they continue to share.

Lastly, but by no means least, I would like to thank my family and friends for their support and encouragement throughout the thesis. To all those at underwater hockey, thank you for getting away from the computer, into the pool, and able to switch off. A final special thank you to Elliott, whose care, understanding, and support at home, has helped me immeasurably.

ABSTRACT

Physics-Informed Machine Learning (PIML) aims to exploit the benefits of both physics and data-based modelling approaches; insight, structure and an enhanced ability to extrapolate are provided through physical knowledge, whilst a data-based component increases flexibility and allows for the capture of complex relationships directly from data. For many applications within engineering, our available physical knowledge might only be sufficient to represent part of a complete systems' behaviour, while the accompanying use of measured data can help capture variation due to environment, manufacturing tolerances, or other effects. When working effectively, PIML models often outperform the individual physics and data-based models from which they are constructed.

This thesis develops PIML models within the field of structural dynamics, with a focus toward Structural Health Monitoring (SHM), offshore structures and wave loading prediction. In harsh environments, including those offshore, conditions are often difficult to fully characterise with purely physics-based approaches, whilst a high variability of conditions places large demands on the collection of measured data. Typically, the predictions of machine learners are only suitable within the realms of previously observed conditions, requiring extensive, and expensive, monitoring campaigns. PIML has the potential to address both of these issues.

Amongst the thesis contributions, is the first instance of a physics-informed model for the prediction of wave loads on a real offshore structure. The widely used Morison's Equation, an empirical wave loading solution, is incorporated within the mean function of an autoregressive form of Gaussian process Regression (GP-NARX). The model achieved a 29.13% and 5.48% relative reduction in error over Morison's Equation and a purely data-based GP-NARX respectively. Enhanced improvement was seen when extrapolating, where the model was able to rely upon physical knowledge to overcome a scarcity of measured data.

Other novel contributions include the design and undertaking of an experiment, within a laboratory wave tank, to investigate wave loading across a range of representative ocean state spectra. A monopile structure, representing an offshore wind turbine, was heavily instrumented with accelerometers, strain gauges, flow meters, wave gauges and, most importantly, a force collar. The availability of wave load measurements, even within an experimental setting, are rare, making the dataset a valuable resource with which to develop and validate models. Within the thesis, the dataset was used to construct PIML models that relied on only incoming wave height as a model input, a commonly available variable on many offshore structures. Aspects of linear wave theory were integrated within a GP-NARX framework to remove the requirement for access to measured flow conditions close to a structure. The installation of flow meters in offshore environments is typically both expensive and challenging.

In the latter stages of the thesis, after a range of PIML models have been developed, a wider view of PIML is taken. Being a relatively new field, the rules of best practise of how to develop models for a given scenario (available level of physical knowledge and measured data) are not yet in place. Relationships between how different types of physics may be incorporated within a model, the effects of changing model structure and the capabilities of constructed models are studied and discussed. The concluding contribution of the thesis is the development of a framework to aid with how best to integrate a given piece of physical knowledge within a PIML model.

PUBLICATIONS

Journal papers: Published

E. J. Cross, T. J. Rogers, D. J. Pitchforth, S. G. Gibson, S. Zhang and M. R. Jones, A spectrum of physics-informed Gaussian processes for regression in engineering. *Data-Centric Engineering*, e8 5, 2024.

D. J. Pitchforth, T. J. Rogers, U. T. Tygesen and E. J. Cross, Grey-box models for wave loading prediction. *Mechanical Systems and Signal Processing*, 107741 159, 2021.

Journal papers: In preparation

D. J. Pitchforth, M. R. Jones, S. J. Gibson and E. J. Cross, Physically-informed change-point kernels for structural dynamics.

D. J. Pitchforth and R. S. Mills and T. J. Rogers and U. T. Tygesen and E. J. Cross, Physics-informed Gaussian processes for wave loading prediction: An experimental study.

Book chapters

E. J. Cross, S. G. Gibson, M. R. Jones, D. J. Pitchforth, S. Zhang and T. J. Rogers, Physics-Informed Machine Learning for Structural Health Monitoring. In *Structural Health Monitoring Based on Data Science Techniques*, pages 347–367, Springer, 2022.

Conference papers

D. J. Pitchforth, M. R. Jones and E. J. Cross, Physics-informed Gaussian processes - a guide to what physics where?. *ISMA 2024 - International Conference on Noise and Vibration Engineering and USD 2024 - International Conference on Uncertainty in Structural Dynamics*, 2024.

D. J. Pitchforth, M. R. Jones and E. J. Cross, Physically-informed change-point kernels for variable levels of physical knowledge inclusion in Gaussian processes. *Proceedings of the 11th European Workshop on Structural Health Monitoring (EWSHM 2024)*, 2024.

M. R. Jones and D. J. Pitchforth and E. J. Cross, Gaussian process kernels for partial physical insight. *Proceedings of the 11th European Workshop on Structural Health Monitoring (EWSHM 2024)*, 2024.

D. J. Pitchforth and R. S. Mills and T. J. Rogers and U. T. Tygesen and E. J. Cross, Physics-informed Gaussian processes for wave loading prediction. *Proceedings of the 14th International Workshop on Structural Health Monitoring (IWSHM 2023)*, 2023.

M. R. Jones and D. J. Pitchforth and E. J. Cross, Partially Structured Gaussian Processes for Grey-Box Learning in SHM. *Proceedings of the 14th International Workshop on Structural Health Monitoring (IWSHM 2023)*, 2023.

D. J. Pitchforth and T. J. Rogers and U. T. Tygesen and E. J. Cross, Incorporation of partial physical knowledge within Gaussian processes. *ISMA 2022 - International Conference on Noise and Vibration Engineering and USD 2022 - International Conference on Uncertainty in Structural Dynamics*, 2022.

D. J. Pitchforth and T. J. Rogers and U. T. Tygesen and E. J. Cross, Physically-informed kernels for wave loading prediction. *SHMII 2022 - The 11th International Conference on Structural Health Monitoring of Intelligent Infrastructure*, 2022.

D. J. Pitchforth and T. J. Rogers and U. T. Tygesen and E. J. Cross, A grey-box model for wave loading prediction with uncertainty propagation. *ISMA 2020 - International Conference on Noise and Vibration Engineering and USD 2020 - International Conference on Uncertainty in Structural Dynamics*, 2020.

Workshop contributions

D. J. Pitchforth and R. S. Mills and T. J. Rogers and U. T. Tygesen and E. J. Cross, Physics-informed Gaussian processes for wave loading prediction. *The 5th Sheffield Workshop on Structural Dynamics*, 2023.

E. J. Cross and S. J. Gibson and M. R. Jones and W. Lin and R. Nayek and D. J. Pitchforth and T. J. Rogers, Grey-box benchmarks System identification with Gaussian processes. *Workshop on Nonlinear System Identification Benchmarks 2021*, 2021.

TABLE OF CONTENTS

1	Introduction	1
1.1	Motivation	2
1.1.1	The balance between physics and data	2
1.2	Structural Health Monitoring	5
1.2.1	SHM for Offshore structures	7
1.2.2	Wave loading prediction	8
1.3	Contribution of this thesis	10
1.3.1	Outline of the thesis and contribution of each chapter	10
2	Integrating Physics and Data	12
2.1	Terminology	12
2.2	A machine learning pipeline	14
2.2.1	Problem focussed knowledge integration	15
2.2.2	Data focussed knowledge integration	18
2.2.3	Training focussed knowledge integration	21
2.2.4	Model focussed knowledge integration	23
2.2.5	Output focussed knowledge integration	29
2.3	Grey-box models for structural dynamics	31

3	Background Theory and Interpretation	32
3.1	Gaussian processes	32
3.2	Wave theory, sea states and wave loading	44
3.2.1	Linear waves	44
3.2.2	Representation of sea states	47
3.2.3	Empirical wave loading prediction	49
3.3	Concluding remarks on theory	52
4	Data and Experimentation	54
4.1	The Christchurch Bay Tower	54
4.2	Monopile Experiment	57
4.2.1	The structure	57
4.2.2	The wave tank	62
4.2.3	Measurement of the wave load	63
4.2.4	Measurement of flow conditions	66
4.3	Test cases	67
4.3.1	Sinusoidal waves	67
4.3.2	Representative sea states	68
4.3.3	Steady flow	69
4.4	Data within the thesis	71
5	A First Grey-box Model for Offshore Structures	72
5.1	Development of a grey-box model	72
5.1.1	White-box	73
5.1.2	Black-box	75
5.1.3	Grey-box	76
5.2	GP-NARX implementation	78
5.2.1	One step ahead and model predicted output	78

5.2.2	Uncertainty propagation in a GP-NARX	80
5.2.3	Hyperparameter optimisation	81
5.2.4	GP-NARX lag selection	82
5.3	Results on the Christchurch Bay Tower	86
5.3.1	Uncertainty propagation	86
5.3.2	Model predictive performance	87
5.3.3	Model performance during extrapolation	90
5.3.4	Computation time	96
5.4	Conclusions	97
6	Enhancing Gaussian Process Wave Loading Predictions Using Linear Wave Theory	98
6.1	Measurement of waves	98
6.2	Integrating physics and data	100
6.2.1	Input augmentation: Linear wave theory	100
6.2.2	Mean function: Morison's Equation	107
6.3	Results on the Monopile Experiment Data	109
6.3.1	Morison's Equation	109
6.3.2	Autoregressive modelling	110
6.3.3	GP-NARX	117
6.3.4	Interpretability through inclusion of prior knowledge	118
6.4	Conclusions	121
7	Partial Knowledge Inclusion Through Kernel Design	122
7.1	Kernel design	122
7.2	Morison's Equation in kernel form	125
7.2.1	Model Predictive Performance	128
7.2.2	Posterior contribution breakdown	130

7.2.3	Conclusions	131
7.3	Physical kernels for the free vibration of a cantilever beam	133
7.3.1	Relaxing assumptions through kernel design	136
7.3.2	Recovery of mode shapes	138
7.4	The balance between physics and data	141
7.5	Conclusions	142
8	Exploring the Effects of Prior Knowledge Inclusion on Gaussian	
	Process Model Structure	144
8.1	Prior knowledge through linear models	145
8.2	Modification of GP model structures	147
8.2.1	Deterministic mean function	147
8.2.2	Stochastic mean function	149
8.2.3	Physics-informed kernel with zero mean prior	150
8.2.4	Physics-informed kernel with known mean	151
8.2.5	Summary of methods	154
8.3	Toy case study: Learning from limited data	155
8.3.1	Black-box SE kernel fit	156
8.3.2	Stochastic mean function	158
8.3.3	Physics-informed kernel with zero mean prior	159
8.3.4	Physics-informed kernel with known mean	161
8.3.5	Performance summary	162
8.4	Conclusions	164
9	A Framework for Prior Knowledge Inclusion Within Gaussian	
	Processes	166
9.1	Categorising prior knowledge	167
9.1.1	Types of knowledge source	167

9.1.2	Knowledge forms	169
9.2	A framework for prior knowledge inclusion	173
9.2.1	Example A: Morison's Equation	175
9.2.2	Example B: Linear wave theory	178
9.2.3	Example C: Free vibration of a cantilever beam	181
9.3	Conclusions	184
10	Conclusions and Future Work	186
10.1	Summary of contributions	187
10.2	Future work	191
10.2.1	Continued PIML model development	191
10.2.2	Decision-making for PIML	192
10.3	Challenges for PIML in SHM	194
10.3.1	Dealing with large quantities of data	194
10.3.2	Operator confidence and benefit to the industry	195
	Bibliography	197

Chapter 1

INTRODUCTION

The construction of engineering models has, until recently, occurred from two opposing philosophies: physics-based approaches, which use our insight into how a system behaves to mathematically represent underlying phenomena; and data-based approaches, opting instead, to infer relationships directly from measured data. Despite preferences, varying by both research community and application, there is no singular correct answer as to the ‘right way’ to construct a model.

The use of physical knowledge and machine learning in combination creates models that aim to exploit the benefits of both physics and data-based approaches. Structure, insight into a problem and a degree of understanding into model operation is provided through physics, whilst increased flexibility and an enhanced ability to model complex relationships, that may not be fully understood, is provided through a data-based learner.

The work of this thesis applies Physics-Informed Machine Learning (PIML) models to the field of structural dynamics, with a particular focus towards offshore structures. The use of PIML for structural dynamics, though growing rapidly, is in its infancy with a huge potential for benefit within a range of real world engineering applications. A key goal of this thesis, along with the development of novel model architectures, is to explore how one might decide how to combine physics and data most effectively for a given modelling scenario.

This chapter motivates the need for physics-informed machine learning within structural dynamics applications, and details the contributions and structure of the thesis.

1.1 Motivation

Many engineering structures, ranging from offshore wind turbines to aircraft landing gear, operate in harsh environments, in which accurate representation of dynamic behaviour and prediction of remaining fatigue life is difficult. In such extreme environments, physics-based (**white-box**) models are often unable to fully capture the complexity of dynamic behaviour. For example, phenomena including movement of mechanical joints, thermal effects and humidity are difficult to characterise and, therefore, model and validate in a dynamic context, often leading to the adoption of large factors of safety at the design stage.

A data-based (**black-box**) approach aims to provide a more flexible alternative, where machine learning techniques may be used to characterise relationships between measured variables directly. The nature of the variables being modelled is arbitrary and no prior understanding of the physics is required. Although purely data-based approaches have proven to be effective across a range of predictive tasks within structural dynamics [1–3], machine learning models are not without their drawbacks. Black-box models are generally poor at extrapolation, with performance suffering in conditions outside the scope of provided training data. Overfitting and the adoption of unnecessarily complex model structures during training can also be an issue [4, 5].

A physics-informed (**grey-box**) approach combines physics-based modelling and data-based learning with the aim of having a flexible model that is informed by physical insight. A key area of expected improvement concerns extrapolation; it is hoped that the structured white-box component of the model will assist inference in areas where training data coverage is low. Improvement of performance outside the observed training conditions would reduce the current demand for increasingly large training datasets and associated monitoring efforts.

1.1.1 The balance between physics and data

One may consider the construction of physics-informed machine learning models to lie on a scale between the extremes of a purely physics-based and purely data-based models, similar to that shown in Figure 1.1. In the case of complete physical understanding, a purely physics-based model would be most appropriate whilst zero knowledge of a process would suggest the implementation of a purely data-driven learner. The intuitive sliding colour scale from white, through various shades of grey, to black is a key reason for the author’s preference for the terminology

‘grey-box’ model. For instance, one could quickly grasp what would be meant by a ‘light grey’ (more reliant on physics) or ‘dark grey’ (more reliant on data) model.

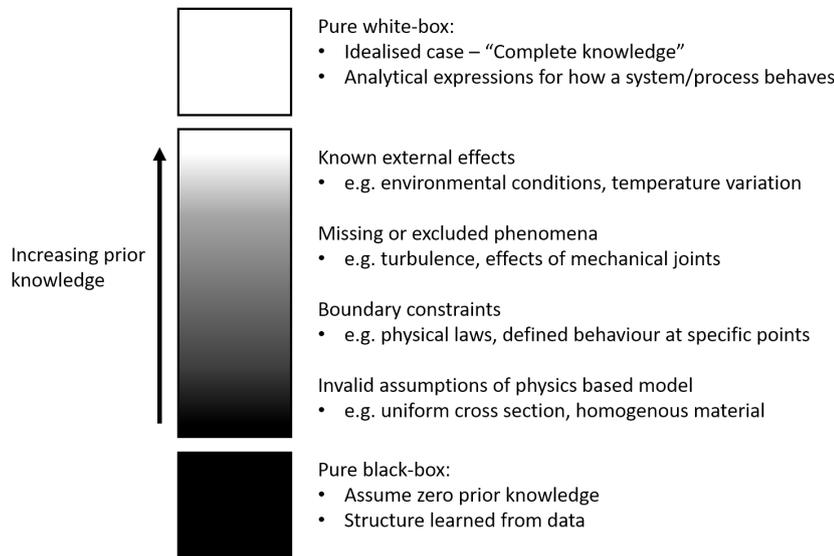


Figure 1.1: Sliding scale of prior knowledge inclusion within physics-informed machine learning. Examples of prior knowledge are placed in their approximate place on the scale for visualisation purposes, however, there will be significant overlap between categories.

External effects such as environmental conditions and weather can change the dynamic behaviour of a structure and failure to account for this can lead to increasing model error and misclassification of damage states [6, 7]. Knowledge of how a structure’s behaviour changes with its environment can be useful to improve prediction quality within varying environments. Zhang *et al.* [8] incorporated the relationship between temperature and longitudinal deflection of a bridge deck within the mean function of a Gaussian process to improve prediction quality across varying seasons.

An important, and recurring, theme throughout this thesis is the balance between physics and data. For any given modelling scenario faced by a user, one can assess the relative levels of both knowledge and data available. This will determine the location on the chart shown in Figure 1.2. The amount of knowledge here is defined as the extent to which a user understands and is able to define a process or phenomena, ranging from zero (no assumed knowledge), to complete understanding. An important point to note about the data axis, is that it is not strictly a measure of the quantity of data available (in terms of gigabytes etc.), it represents the proportion of possible conditions observed by a model. To highlight this, consider

the following two examples:

- Case 1: 100 GB of measured data for a structure, all recorded during a heatwave in the month of July.
- Case 2: 52 GB of measured data for a structure, where each GB was recorded during a different week of the year.

Although the raw quantity of data for Case 1 is higher, a much wider range of conditions will be observed for Case 2. Access to seasonal variations will allow structural behaviour to be captured within a more varied range of environmental conditions. The author argues that this would place Case 2 at a higher point on the data axis despite a lower quantity of data.

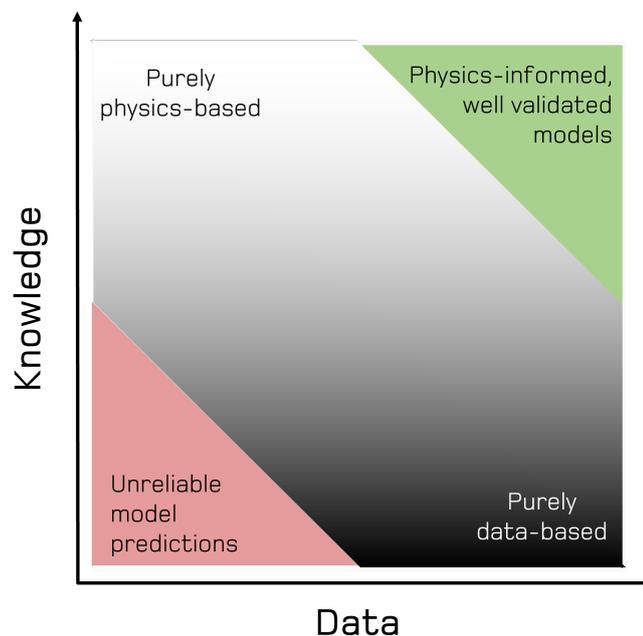


Figure 1.2: The knowledge vs data axes for a given modelling scenario. The knowledge scale ranges from zero (e.g. unknown variables) to complete understanding (e.g. closed form solutions available). The data scale measures the proportion of total possible conditions that have been observed by the model [9].

There are many possible ways to construct a grey-box model, an overview for some of which can be found in [9–12]. Figure 1.3 presents a range of existing grey-box modelling approaches and highlights an important concept; the type of model construction affects the balance between physics and data within the model.

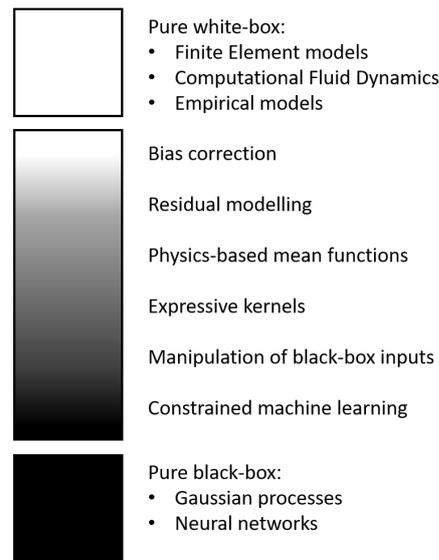


Figure 1.3: Potential Grey-box modelling approaches and their approximate location on a grey-box scale.

The concepts presented in Figure 1.1, Figure 1.2 and Figure 1.3 all represent important stages of the PIML modelling process. The identification of available physical knowledge, the relative balance between knowledge and measured data, how to design a model to reflect this relative balance and considerations of how physics may be integrated within a given machine learner are all important challenges. An understanding of the PIML model development process, and how some of these decisions may be made more easily is an important area of focus for the thesis.

1.2 Structural Health Monitoring

Physics-informed machine learning has shown to be a promising area for research and has the potential to be applied successfully to a wide range of tasks within structural dynamics. These may range from full structure response prediction, to health state estimation, to smaller subtasks such as load quantification. This thesis hopes to contribute toward the adoption of PIML in structural dynamics, with a focus on structural health monitoring for offshore structures and wave loading prediction. The nesting of subfields and contributions of the thesis are summarised in Figure 1.4.

Structural Health Monitoring (SHM) [13, 14] involves the collection, and utilisation, of data from a structure to estimate its condition (health-state) and inform

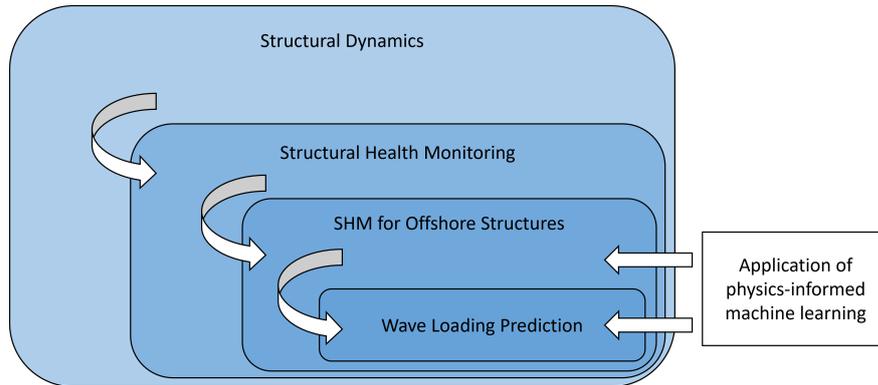


Figure 1.4: Graphical summary of the thesis contribution: The development and application of physics-informed machine learning models for structural dynamics, with a focus on wave loading prediction and structural health monitoring for offshore structures.

operational strategies. Outcomes from an SHM system might range from continued operation, for a healthy structure; to the scheduling of repair work, for a damaged structure; to full decommissioning, for a structure deemed at the end of its Remaining Useful Life (RUL). A distinguishing feature of SHM is that the data collected must be used, applied, or acted upon in some way; the installation of a measurement system and recording of data alone would not qualify as SHM [7].

The objectives for SHM are concisely defined by Rytter’s hierarchy [15], which breaks down the challenge of health-state estimation into a series of ‘levels’. Each level builds upon the work of those previous, with the difficulty increasing down the list. Rytter’s hierarchy is summarised as:

1. Detection - Is there any damage present?
2. Localisation - Where is the damage?
3. Assessment - What is the type and severity of the damage?
4. Prognosis - What is the remaining useful life of the structure?

Worden and Dulieu-Barton [16] extended the work of Rytter, splitting the assessment level into two steps: classification determined the type of damage present (e.g. cracks, corrosion, delamination etc.), whereas quantification determined the severity (e.g. crack length).

The potential advantages from the successful implementation of an SHM system in industry are vast and can be broadly considered to fall within two main categories: humanitarian and economic. The first of these includes the protection of human lives, which in the author's opinion, is the most important motivator for SHM systems. The early identification of structural damage, and repair/decommissioning of structures where necessary, could help avoid many life-threatening disasters caused by bridges [17] and aircraft [18, 19]. For structures including oil rigs, a failure may also have a lasting impact on the surrounding ecosystem [20], further motivating the need for an SHM system.

Although there will be an upfront cost associated with any SHM system, the price of which will vary by scale, application and sensor type [21, 22], a successful implementation aims to repay this with profit. A transition from Time Based Maintenance (TBM), where repairs and inspections are typically carried out periodically, to maintenance strategies informed by SHM has several potential financial advantages. Firstly, the quantity of unnecessarily scheduled inspections could be reduced, leading to large savings in applications such as offshore, where each visit to a structure is very expensive [23]. The ability of an SHM system to detect potential damage early and more quickly inspect or repair the damage can also reduce costs. The earlier a damage source is detected, the less chance the damage has to worsen (e.g. crack growth) before action is taken.

1.2.1 SHM for Offshore structures

Many engineering structures within the North Sea have already exceeded or are close to their initially specified 20-25 year design lives [24]. The financial incentive for continued operation of structures, beyond their design lives, brings attention to the need for accurate prediction of remaining fatigue life. Safety concerns around the operation of ageing structures are a key priority and confidence within prognosis is, therefore, paramount. For modern installations, particularly large-scale, expensive projects, minimising the risk of investment through informed decision-making [25, 26] is a key driver for SHM.

Offshore engineering structures operate in harsh, highly variable environments, in which characterisation of the conditions, and how these interact with the structure is challenging. Variation of weather, temperature, salinity and sea state place a significant strain on the required amount of data to be collected. For the occurrence of events including storms and freak waves [27], having measured data for a similar such event becomes increasingly unlikely. The ability of purely data-based models to predict on unseen data therefore becomes a key area of concern, and an important

area for PIML models to aim to improve.

An important factor impacting the nature of offshore SHM implementation is the type of structure being monitored; a distinction may be drawn between offshore wind turbines, and oil and gas platforms [28]. A key difference to highlight is the type of loading experienced: for a wind turbine, large aerodynamic loads due to wind and operational loads due to the rotating blades are major considerations. For an oil or gas platform, the priorities are instead focussed on large static loads, due to the weight of platform decks, and dynamic loading due to waves [29]. Although wave loading occurs on all offshore structures, its importance is typically higher for oil and gas platforms, and a secondary factor to wind loads for offshore wind turbines.

Within research communities, the study of Computational Fluid Dynamics (CFD) has dominated the quantification of wave loading forces [30, 31]. Within industrial applications however, the high computational resource requirements of CFD and difficulty of model validation for structures in complex environments, has led to a preference, in some industries at least, for more simplistic empirical methods. A common example of one such method is Morison's Equation [32], which offers an empirical solution for wave loading with minimal computational resources.

1.2.2 Wave loading prediction

The cyclic loading of an offshore structure due the motion of waves is a driving factor of fatigue and therefore has a significant impact of the useful remaining life of the structure. The quantification of wave loading, through either measurement or prediction, presents its own set of challenges. The direct measurement of wave loads acting on offshore structures is rare, and where attempted, it often requires the development and installation of bespoke systems. Even when measurements may be available, these are generally at point locations and do not provide access to a distributed load over the structure. The prediction of wave loads across a structure, using data readily available from other sensors has the potential to provide access to wave loads where they can not be measured and reduce the cost of implementing additional measurement equipment. This thesis will focus on the utilisation of sensor data commonly available on offshore structures, including incoming wave height, often measured using wave radars [33, 34].

The modelling of waves and prediction of wave loads acting on structures is challenging and forms an extensive field of research [35, 36]. The harsh offshore environments, highly variable conditions and complexity of phenomena such as

vortex shedding and breaking waves makes the validation of physics-based models difficult. Attempts to represent the underlying physics have led to the development of Computational Fluid Dynamics (CFD) models which have been shown to be effective in a range of wave loading prediction tasks [31, 37, 38]. A key challenge facing physics-based approaches, including CFD, is that as the complexity of the phenomena being captured grows, so too does the required model fidelity. For a case such as breaking wave dynamics [39], the resulting CFD model is computationally expensive and requires extensive resources (time, money and technical expertise) to validate.

An ability to model processes without complete understanding of the underlying physics has been a key driver of the development of data-based models. Here the relationship between variables may be learned directly, without prior knowledge of how a process may behave. Within the field of wave loading quantification, Neural Networks [40], Gaussian process NARX models [41] and Bayesian Regression [42] have shown to be helpful tools for capturing the non-trivial relationship between flow conditions and wave force. Although effective when operating within the realm of previously observed conditions, a tendency to extrapolate poorly, often exhibiting unexpected behaviours, and a lack of insight in to how a model operates has posed an obstacle to their adoption within an industrial setting.

Within industrial applications, empirical solutions including Morison's Equation [32] are often popular; they offer a computationally inexpensive means of achieving approximate wave loading solutions. They have been shown to be effective within a variety of applications including wind turbines [43, 44] and Spar platforms [45]. Although the limitations of such methods are understood (e.g. Morison's Equation valid for unidirectional waves [46]), they have been widely implemented for many years, and as such provide a degree of trust within model operation when used within appropriate conditions. An increasing concern, threatening a continued reliance on existing empirical methods is the growth of modern offshore wind turbine sizes. From 2010 to 2023, the average rotor diameter of a newly installed offshore wind turbine more than doubled, increasing from 126m to 280m [47]. As turbines, and their towers, grow, a necessary assumption of slender structural members becomes less valid and the results of empirical models less accurate.

At the outset of this thesis, the adoption of PIML within SHM applications was very rare, with purely data-based methods dominant within the previous decade of literature. For offshore structures, and the field of wave loading prediction specifically, the author was unaware of any existing PIML approaches. This presented a key area of novelty in which to make research contributions.

1.3 Contribution of this thesis

There are two aims of this thesis:

1. *To develop novel grey-box models for structural dynamics that combine physics and data-based learning with the intention to exploit the benefits of either approach used independently.*
2. *To explore and formulate the relationship between the type of prior physical knowledge possessed by a user and how best to incorporate it within a model.*

The first aim will be the initial focus of the thesis, forming a building block from which to tackle the second. A range of models, developed throughout the thesis and applied to relevant tasks within structural dynamics are used to study what types of model are most appropriate for a given modelling scenario. A framework to aid with the selection of model structure will form the concluding chapter of the thesis.

1.3.1 Outline of the thesis and contribution of each chapter

An outline of the content within the upcoming chapters of the thesis is as follows:

- Chapter 2 presents literature from the field of physics-informed machine learning and addresses how these methods change the typical workflow of a purely data-based approach. Current gaps in research fields, in which to focus PIML model development efforts, are highlighted.
- Chapter 3 provides necessary background theory and the authors interpretations on Gaussian process regression, the machine learner of choice for this thesis. Relevant physical topics including linear wave theory and empirical wave loading prediction are also introduced.
- Chapter 4 details the primary datasets used for model development throughout the thesis: the Christchurch Bay Tower (CBT) project, a full scale structure in a real offshore environment, and a monopile wave tank experiment, designed and carried out by the author at the Laboratory for Verification & Validation (LVV).

-
- Chapter 5 develops the first grey-box model to predict wave loads on an offshore structure. The widely used Morison's Equation is incorporated within a GP-NARX residual model, with a focus placed on improving performance in extrapolation on the CBT dataset.
 - Chapter 6 develops wave loading prediction models that utilise only incoming wave height as an input, a commonly available measured variable on many offshore structures. Experimental data are used to validate and include aspects of linear wave theory within Gaussian process models.
 - Chapter 7 focusses on physics-informed kernel design for Gaussian processes. Circumstances of how one might incorporate partial understanding of a process are addressed.
 - Chapter 8 explores how the method of incorporating prior physical knowledge within a Gaussian process affects the structure and capabilities of the final model. A case study using a consistent form of prior knowledge is used to directly compare results.
 - Chapter 9 presents a framework for prior knowledge inclusion within Gaussian processes. A tool, in the form of a flowchart, aims to help a user decide how best to incorporate a given piece of prior physical knowledge within a model. The success of the flowchart outcomes are tested using findings from throughout the thesis.
 - Chapter 10 concludes the thesis and discusses the direction of future work.

INTEGRATING PHYSICS AND DATA

This chapter introduces the reader to a range of modelling approaches that combine physical knowledge and machine learning within the literature. Methods are presented in order of progression through a typical black-box machine learning pipeline, from the initial definition of a problem to a complete functioning model, addressing how physics may be incorporated at each stage. The aim is to place the reader within the shoes of a prospective modeller, whom might wish to implement such methods, which will be beneficial throughout the thesis. The gaps in implementation, for which many of the methods includes structural dynamics, are addressed, and highlighted as areas of focus for the work of the thesis.

2.1 Terminology

The concept of combining physics with data-based learning is known by many names and is not exclusive to the term ‘grey-box model’ [11, 48], although that will be the preferred term throughout this thesis. Approaches that combine physics with data-based learners have been referred to as ‘hybrid models’ [49–51], ‘physics-informed’ [12, 52, 53], ‘physics-guided’ [54, 55], ‘physics-inspired’ [56–59] and ‘semi-physical’ [60, 61] amongst other terms [62–64]. The balance between physics and data within a model is not necessarily tied to the language used to describe it; for example, linguistically, one might expect a ‘physics-guided’ model to be more reliant on physics than one which is ‘physics-inspired’, but this need not be the case. The terminology used to define a model is primarily dependent on the preferences of

the author, rather than the specific meaning of language used, with a selection of valid terms possible for equivalent models.

The field of integrating physical knowledge within machine learners has grown rapidly in recent years and continues to do so. Benefits including a reduced reliance on data collection, reduced computational demands, increased capability to extrapolate and providing interpretability into model results are key drivers of this growth. The return of search results for academic articles and book chapters containing field relevant terms is shown in Figure 2.1. In the decade 2013-2023 the sum of results returned for all search terms grew by a factor of more than 10. Although search results are a proxy for the true underlying quantity of relevant research, their sharp growth highlights the increasing interest.

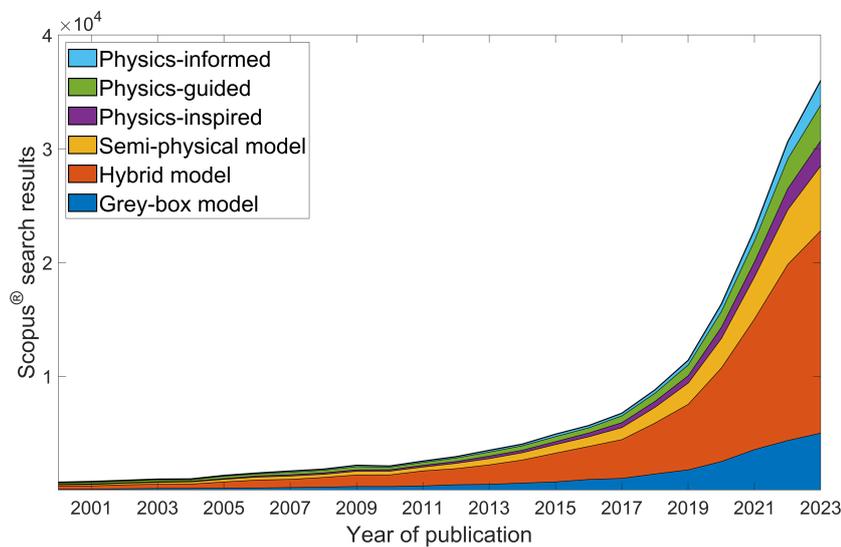


Figure 2.1: The growth of terms used to describe models combining physics with machine learning from the year 2000 to 2023. A Scopus[®] search was used to count results for academic articles and book chapters within a given publication window. In each case, terms were searched with ‘AND’ ‘Machine learning’ to alleviate the impact of search results that may refer to something unintended e.g. a model of a hybrid vehicle.

This chapter presents an overview of relevant literature, categorising similar methods of introducing physics within a machine learner and discussing their effect on modelling capabilities. More in depth analysis of selected methods, and how they tension against the models developed in this thesis, is saved for later chapters, with a focus here on introducing the reader to a variety of key approaches within the field.

2.2 A machine learning pipeline

There are a wide array of techniques to integrate physical knowledge within a machine learner, the ordering and presentation of which can be shown in a variety of ways. Karniadakis *et al.* [12] focus on the primary benefits and challenges of physics-informed machine learning and how specific methods might help potentially tackle these. Von Reuden *et al.* [10] separate methods by prior knowledge representation e.g. algebraic equations, logic relation etc., and explain how each type is most commonly integrated within a machine learner. Cross *et al.* [65] look at techniques specifically within the field of Structural Health Monitoring (SHM) and discuss the balance between physics and data within model structures. Here, the literature of physical knowledge integration techniques is presented narratively, as they would appear through the progression of a typical machine learning pipeline. The deviations from a black-box approach and at what stage of the modelling process physical knowledge is introduced will be highlighted. A simplified diagram of a black-box machine learning pipeline is presented in Figure 2.2. Throughout the chapter, modifications to this pipeline, caused by the incorporation of physical knowledge, are highlighted in blue using similar diagrams.

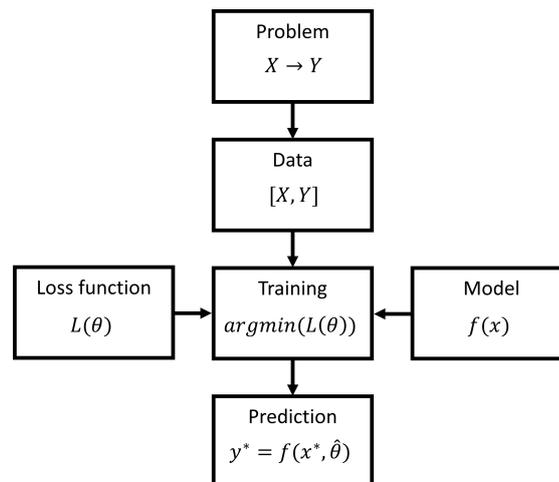


Figure 2.2: Typical black-box machine learning pipeline.

A typical machine learning pipeline¹ aims to progress from some defined problem $X \rightarrow Y$, the relationship between a set of inputs X and a target Y , to a model capable generating predictions y^* for a new set of inputs x^* . To achieve this, a set of data $[X, Y]$ is used to maximise model performance according to a given loss

¹References made here to a typical machine learner will most accurately apply to regression tasks, the primary focus of ML implementation within this thesis.

(or objective) function, via determination of an optimal set of model parameters $\hat{\theta}$. These learned parameters may then be used within the model to generate a prediction y^* for a new set of inputs x^* . The inclusion of physical knowledge is possible throughout each stage of this process, which will now be presented along with appropriate examples from the literature.

2.2.1 Problem focussed knowledge integration

The first stage of a typical machine learning pipeline is the definition of a problem $X \rightarrow Y$, the relationship between a set of inputs X and a target Y to be learned. Methods focussed on using physical knowledge to assist with the definition of $X \rightarrow Y$ are highlighted in Figure 2.3.

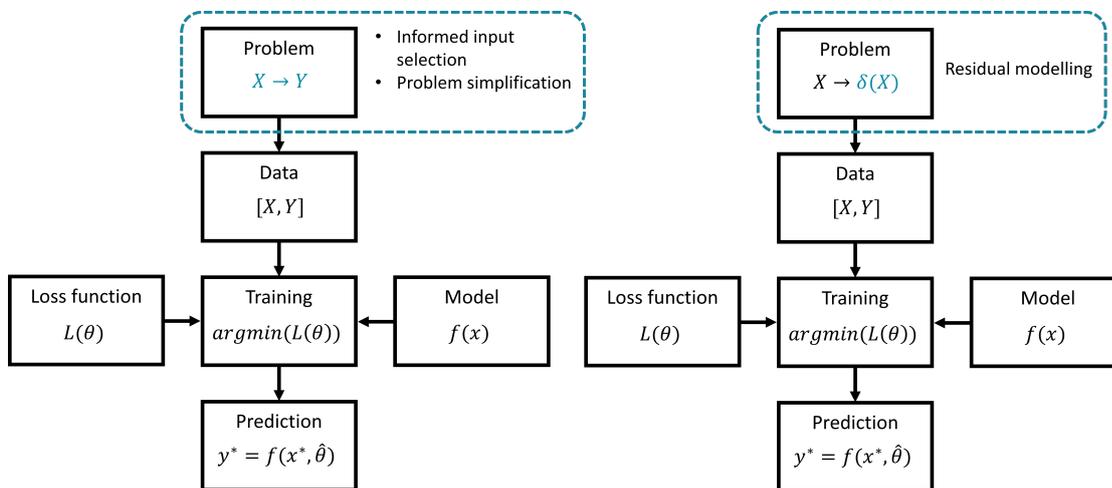


Figure 2.3: An overview of problem focussed physical knowledge integration methods within a typical machine learning pipeline. Informed input selection and problem simplification can change how the inputs X and target Y are specified. Residual modelling replaces the target Y with the residual of a physics-based model $\delta(X)$.

Informed input selection

The selection of appropriate inputs is a highly important step of the modelling process, which if done incorrectly may harm prediction quality with little potential scope to remedy through further modelling decisions [66, 67]. Within a black-box pipeline, inputs are selected via a range of means including statistical methods [68] and information theory [69]. Emphasis is placed on correlation and shared information with the target, diversifying selected features and minimising the total

number of selected features [70, 71].

Informed input selection focusses on the use of prior knowledge to aid with the selection of model inputs, either through accompanying or replacing data-based methods. Duarte [72] addresses the relative balance of prior knowledge and data for selecting inputs (relevant wavelengths) for spectroscopy. As experimental conditions became less favourable, reducing the signal strength of relevant spectra, the reliance on prior knowledge over data-based input selection methods increased.

A key benefit of informed input selection is model parsimony, the improvement (or maintaining) of performance whilst minimising the total number of selected inputs. This leads to more simplistic models, that are more convenient to work with and easier to interpret [73]. When selecting a reduced number of inputs from a larger feature set, the use of domain knowledge can help to ensure relevant inputs aren't excluded, thereby achieving a higher performance at smaller selected feature numbers. Des Touches [74] used sparsity-enforcing priors to improve input selection stability and encourage the selection of expert-determined features within the field of Chemometrics. Groves [75] used prior knowledge of airline passenger preferences, such as number of stops on a journey, to limit the potential feature set size for airline ticket price prediction before model inputs were selected.

The interpretability of model inputs and justification for their selection is important for inspiring confidence within model results. The distinguishing of correlation vs causation is important and well understood within modelling [76]. The reliance on external input of field knowledge promotes confidence within a models' construction. Where physical justification of selected inputs can be provided [77], or the decision process itself made more interpretable [78], further confidence within the model can be achieved.

Problem simplification

Problem simplification blurs the line between physics-informed machine learning and machine learning best practise. Here, the term defines the use of expertise or physical understanding to simplify a problem before it is passed to a machine learner, the aim being to achieve an improved performance on the easier task. Within the field of fluid mechanics for example, a high importance is placed on the understanding and definition of a problem, with attempts to simplify being common before any data is presented to a machine learner [53, 79]. A common form of this is dimension reduction and the treatment of 1D flows [80].

Residual modelling

Residual modelling replaces the target Y of the machine learner with the residual

of a physics-based model $\delta(X) = y(X) - f_{phy}(X)$, often referred to as model discrepancy [81–83]. This produces model predictions of the form:

$$y^*(X) = \underbrace{f_{phy}(X)}_{\text{Physical model}} + \underbrace{\delta(X)}_{\text{Machine learner}} + \underbrace{\epsilon}_{\text{Noise}} \quad (2.1)$$

The field of residual modelling acknowledges that physical models, even high fidelity FE and CFD models, are not able to fully capture a process. The role of the machine learner here is to capture phenomena not represented within the physical model, for example the friction of bearings within a swinging pendulum [84] or changing hydro-mechanical behaviours of soil due to internal erosion [85]. The nature of residual model structures does not place a requirement for any particular type of machine learner, with the most appropriate choice dependent on the modelling problem. For example Chen [86] uses a selection of nonlinear AutoRegressive (AR) techniques to capture the residuals of numerical wind models, with AR models being well suited to capturing the non-stationarity of wind speeds [87, 88]. Sirignano [89] used a Deep Neural Network (DNN), primarily due to its highly flexible modelling capabilities, in combination with reduced order Partial Differential Equations (PDEs) for large-eddy simulation of turbulence. The DNN had two roles here; to capture phenomena such as the complexities of turbulence about which little is known and to model phenomena more efficiently than computationally expensive available physical modelling techniques.

Utilising Gaussian Processes (GPs) to model residuals presents a promising framework with residual-based GPs offering a range of potential advantages: being a Bayesian technique they provide a quantification of uncertainty when estimating missing physics and the use of flexible kernels allows the modelling of phenomena about which little may be understood. The similarities shared between residual models and incorporation of physics within the GP mean function is an interesting research area, with discussion of this presented in Chapter 4. GP residual models are used within a large range of applications including modal identification of bridges [90], damage detection of multi-storey structures [83], liquid production rates of wells [91], molecular interactions [92] and spatial prediction of salinity curves [93].

Xing [94] extends the idea of residual modelling via introduction of a sum of GPs (termed ResGP), where the role of each individual GP is the capture the difference of successive increasing model fidelities. The advantage of this model structure is that training data from lower fidelity models can be used to train the early GP components of the sum to reduce reliance on higher fidelity models, improving

performance where high fidelity training data may be sparse. This is particularly helpful in cases of more computationally expensive high fidelity models. The key challenge of ResGP is that access to a range of increasing model fidelities is not always available, and when it is, the validation efforts of ensuring all the models used are working as intended are increased.

The development of residual models for wave loading prediction was identified as an area in which to focus efforts of this thesis. As discussed further in the upcoming Chapter 3, Morison's Equation [32] is a widely used empirical solution for the approximation of wave loads, and is known to simplify the true interaction between a structure and fluid. To overcome this, the development of extensions and corrections to Morison's Equation, to improve its capabilities, presents its own niche research field [95–99]. As yet, the nature of these extensions has primarily been the further inclusion of empirical terms, with the use of a machine learner to perform an equivalent role unexplored. Efforts of the author into this area are presented in Chapter 5 and published in [100].

2.2.2 Data focussed knowledge integration

Upon the definition of a problem $X \rightarrow Y$, machine learners will require data for use during their training. Methods focusing on the generation and modification of data before it is passed to the machine learner are outlined in Figure 2.4.

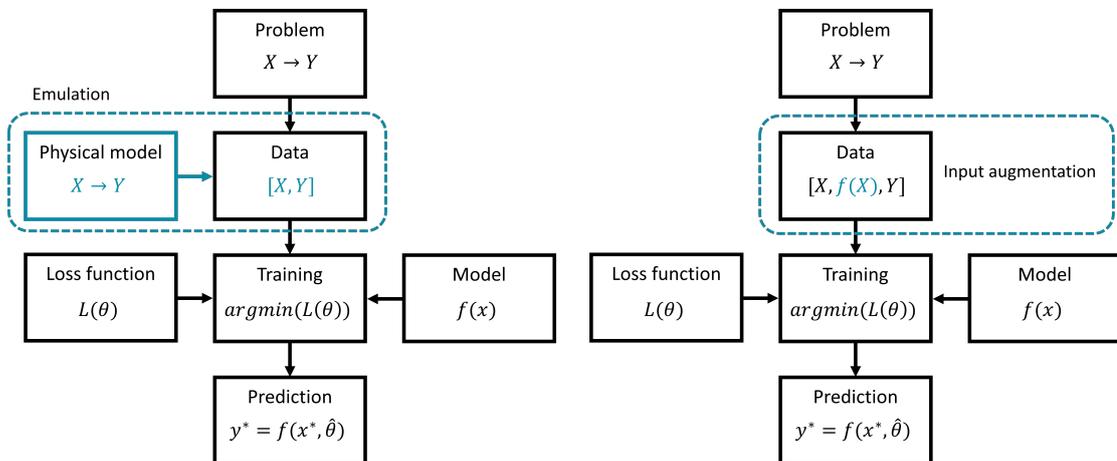


Figure 2.4: Data focussed integration of physical knowledge within a typical machine learning pipeline. Emulators use physical models to generate training data $[X, Y]$. Input augmentation uses physical knowledge to generate additional inputs $f(X)$ for a model.

Emulation

The availability of suitable training data is a key challenge for the development of machine learning models. In many cases it is not possible to measure the breadth of conditions required to adequately train a model. The high variability of real world environments, the rarity of conditions such as freak events [101] and expense of collecting large amounts of data are key drivers for the use of physically derived models and simulations. The field of emulation involves the use of physical models to generate training data for use within a machine learning model, with the machine learner then aiming to replicate the predictions of the physical model [102].

A key driver for the development of emulators is to achieve equivalent prediction quality of physical models at a fraction of the computational cost. This is of particular benefit within the fields such as Computational Fluid Dynamics (CFD) [103, 104] and Finite Element (FE) modelling [105–107], where expensive computational costs and long runtimes discourage running many simulations. Overcoming computational expense makes emulators useful within a variety of fields from infectious diseases [108] to the formation of galaxies [109].

A popular application for emulators is Sensitivity Analysis (SA) [110], which aims to determine how the output of a model varies with respect to changing inputs or parameters. Typically, SA will require a large number of model evaluations, resulting in high computational demands for expensive models. Using an emulator can greatly reduce this cost. SA is particularly important in applications where slight changes in input have the potential for large changes in output; this includes the modelling of natural disasters such as landslides [111] and Tsunamis [112], and environmental modelling [113]. The Gaussian Emulation Machine for Sensitivity Analysis (GEM-SA) [114] is a tool developed by Kennedy and Petropoulos, leveraging the flexibility of GP regression, to enable SA within a variety of applications.

There are many instances where a particular region of training data is known to be difficult to attain and the use of a physical model can be used to supplement existing measured data to train a data-based model. Omigbodun [115] used computational models of lung cancer to supplement measured patient data. The availability of healthy lung data is more widely available within a clinical setting, as of course would be hoped for, and therefore a focus is placed on the capability of models to generate additional data for cancer-positive patients.

Input augmentation

The use of physical knowledge to transform inputs before they are passed to a machine learner is useful to both replace and use alongside (augment) existing measured data. The work of Lindskog and Ljung [60] focussed on the transformation of black-box model inputs through physical insight. The aim here is to reduce the complexity of the relationship that must be learned by the black-box component, such that the chances of a superior model fit are maximised. An advantage of this approach is that standard black-box model structures need not be modified, only their inputs. This approach was used by Sohlberg [116] to achieve increased power generation and reduced water level variation at a hydroelectric power station. They used knowledge of river flow physics and turbine flow-power curves to transform the measured inputs of the model to an estimate of turbine flow.

Alongside potential benefits in performance, transformed additional inputs may offer a degree of interpretability in to a model that was not previously offered if the input has some physical meaning. Fuentes [117] approximated the dynamic pressure acting on an aircraft wing, which provides a measure of kinetic energy within a fluid, as an input for in-flight strain predictions. Rogers [118] used a physical approximation of the current water height within a series of cascaded tanks as an additional model input for estimation of water level in the lower tank. Here, the additional input provided a secondary estimation of the model target which could be compared with the final model prediction.

Brahma [119] focussed on physical transformations of inputs to reframe extrapolation tasks as ones of interpolation. Where prediction tasks may originally lie outside the extreme limits of observed data (i.e. extrapolation), within a transformed input space this is not necessarily the case. For the prediction of diesel engine emissions, extrapolation tasks in the original input space, e.g. cases of high fuel mass injection, were reframed as interpolation tasks in terms of transformed inputs, including Peak Mean Effective Pressure (PMEP) and Fuel-Oxygen Ratio, leading to increased model performance. Although effective in the cases shown, a consideration here is that through input space transformation, the reverse case would also be possible, with a transition from an interpolation to an extrapolation task occurring.

The development of novel input augmentation methods may stem from a number of possible avenues: the input variable itself may be new, the method by which an input is generated may be new, or specific advantages may be achieved through the generation of particular inputs (e.g. performance, reduced reliance on measurement equipment). For the work of this thesis, the development of input augmentation methods for offshore structures was identified as a key area of potential benefit. Its application to wave loading prediction is explored in Chapters 5 and 6.

2.2.3 Training focussed knowledge integration

The training process involves the determination of an optimal set of parameters (or hyperparameters) $\hat{\theta}$ for a selected model. These are obtained via treatment as an optimisation problem, with the performance according to a specified loss function maximised for a given set of training data $[X, Y]$. Methods integrating physics within the training of models are summarised in Figure 2.5.

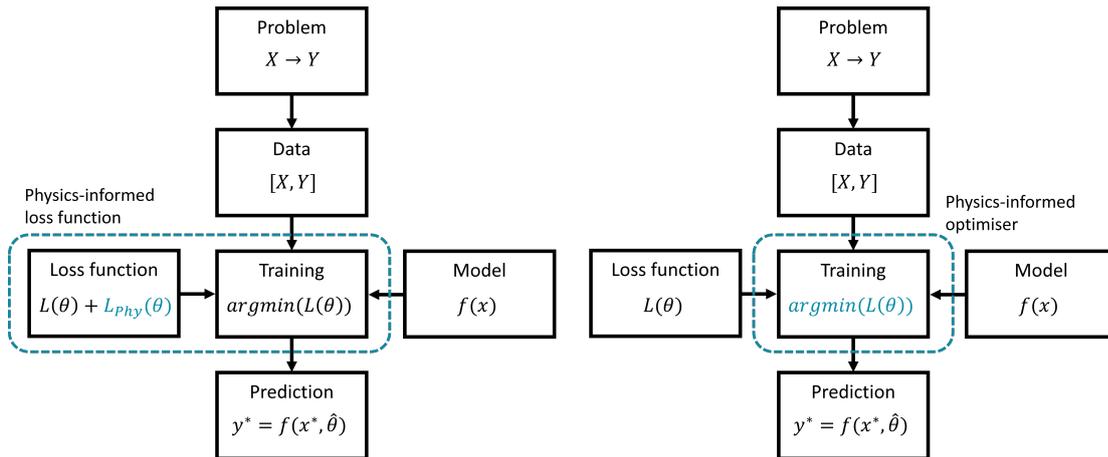


Figure 2.5: Training focussed integration of physical knowledge within a typical machine learning pipeline. Physics-informed loss functions typically include a term $L_{phy}(\theta)$ used to encourage optimisers to converge towards physically desirable solutions. Physics-informed optimisers use physical knowledge to improve optimisation stability and encourage faster convergence for a given task.

Physics-informed loss functions

The loss function provides a measure of model performance and, when minimised, determines the optimum set of parameters for use within a given model. There are a variety of common themes for terms within loss functions including: encouraging mean model fit, maximising likelihood and penalising model complexity [120, 121]. Physics-informed loss functions are used to encourage a model to converge towards a desirable behaviour, constraint, or governing equation, typically via inclusion of an additional term to reward its presence. This generally leads to a loss function of the form:

$$L(\theta) = \underbrace{L_{fit}(\theta)}_{\text{Encourage model fit}} + \underbrace{L_{complexity}(\theta)}_{\text{Penalise complexity}} + \underbrace{L_{phy}(\theta)}_{\text{Reward physical behaviour}} \quad (2.2)$$

Physics-informed loss functions are most commonly used in Neural Networks [122–124], although are possible to use within other types of machine learner [125]. They are used within a range of fields, including, solid mechanics [126], image classification [127], lake temperature modelling [55] and nonlinear diffusivity [128].

Laubscher [129] further separates the loss function with application to heat transfer of flow over in-line tubes, introducing specific physical terms such as adherence to initial and boundary conditions. The relative importance of terms is addressed through a parameter β , to be tuned by the user, allowing particular aspects of physics to be weighted more highly if the user is confident of their knowledge.

Physics-informed optimisers

Although less common than modification of the loss function, it is possible to integrate physical knowledge in to other aspects of the optimisation process. It is well known that it is not possible for any particular optimiser to outperform another across all problems [130]. When working with a specific problem however, it is possible to tune performance, in terms of convergence speed and stability, through the introduction of physical knowledge. A common example of this is the use of physics to constrain an optimisation problem and reduce the potential search space [131, 132]. With application to wind turbines, knowledge of parameters such as drag coefficients and turbine hub height was used to improve robustness and power generation in urban environments [133].

Xie [134] used knowledge of the dynamics of a robotic arm, such as number and connectivity of linkages, in order to speed up optimisation for online model predictive control. Termed ‘optimistic exploration’, a selection of linear dynamics models were repeatedly assessed, with poor performance encouraging the switching of model choice. The role of prior knowledge here was to limit the potential choice of new models and allow the construction of potential models in least squares form, allowing for their computation fast enough to perform online control.

The work of Ha [135] focusses on integrating the prior knowledge and intuition of design engineers within topological feature optimisation. Here a human is included within the loop and asked to judge the quality of designs after a number of iterations. Their input comes in the form of selecting Regions of Interest (ROI), where they believe design changes are most needed for the next set of optimiser iterations. This allows the temporary reduction of search space and unnecessary additional calculation. A challenge of this approach, acknowledged by the authors, is that the identifying of suitable ROI can be difficult, even for experienced design engineers. The quality of structural performance achieved was also highly tied to identification of ROI, placing risk on correct decision-making.

The utilisation of physical knowledge at the training stage of a modelling process poses a valuable tool within the field of structural dynamics. Generally, research efforts focussing on the loss function [136–138] are much more common than any other kind. Although the efforts of this thesis do not focus on the specific development of a physics-informed optimiser, consideration of loss functions, bounding of search spaces and selection of optimiser start points are all influenced by the author’s knowledge of various prediction tasks during model development. The development of task-specific optimisers, designed to work well for a given physical application remains an open and interesting research area in offshore structural dynamics.

2.2.4 Model focussed knowledge integration

The broadest category of physical knowledge integration within machine learners presented here is the direct modification of model architectures. These methods consider approaches that effect how the model itself is constructed, selected or defined to in someway impose physical understanding. Although highlighted within the model definition stage of the pipeline in Figure 2.6, the nature of these methods will also impact both how these models are trained and how they generate predictions.

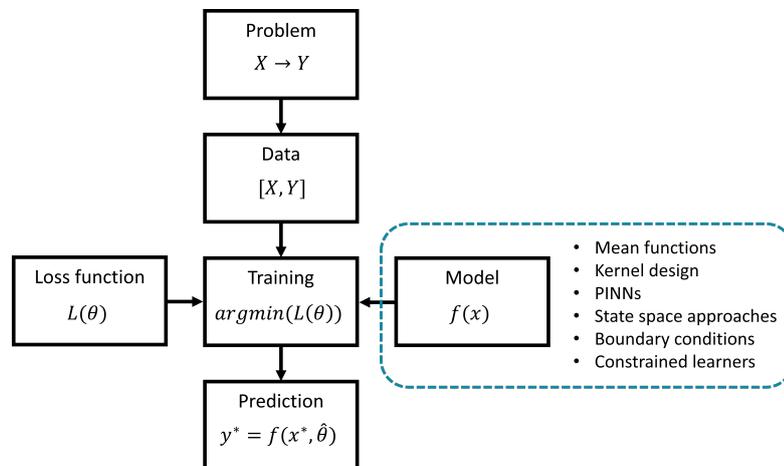


Figure 2.6: Methods focussing on the integration of physical knowledge within the selection or construction of model architectures.

Mean functions

The mean function of a Gaussian process dictates general trends and behaviours within predictions and can be a useful means to integrate prior knowledge within

a GP [139]. As earlier mentioned, the use of mean functions overlaps with the field of **residual modelling**, with an in depth discussion on this presented in Chapter 8. A useful property of mean functions is that they are not restricted to any particular functional form; this makes them popular within a range of applications. Engineering relevant examples include: crack growth [140, 141], tool wear [142, 143], estimation of Remaining Useful Life (RUL) [144], bridge response prediction [8], tracking control of robotic arms [145] and modelling of wind turbine power curves [146, 147]. The potential for mean functions to assist with predictions in regions of sparse data coverage, a key area for expected benefit over purely data-based approaches, is presented in [100] and [107].

For the fields of offshore SHM and wave loading prediction, the potential advantages of GPs with physics-based mean functions are yet to be explored. Also considering the reasons highlighted in the earlier discussion of residual modelling, the incorporation of physical knowledge in GP mean functions is therefore an area of focus for the thesis. The work of Chapter 5 includes the development of models with physics-informed mean functions for wave loading prediction. The capability of models to extrapolate, a key expected benefit, forms a specific study.

Kernel design

The covariance function (kernel) of a Gaussian process is responsible for the family of functions from which predictions may be drawn, with commonly used kernels enforcing properties such as smoothly varying functions, periodicity and localised behaviours [139, 148]. Through careful design of the kernel, it is possible to mimic physically desirable behaviours within drawn functions; such examples include the representation of a physical process [149, 150] or the enforcement of axial and rotational symmetries [151–154]. Being more restricted in their construction, having to satisfy several constraints, the design of kernels presents a more challenging means for the embedding of physics within Gaussian processes. A more detailed discussion on integration of physical knowledge through kernel design is presented in Chapter 7 with a few key examples shown here.

The work of Duvenaud [155, 156] attempts to tackle the challenge of kernel selection via automation of the process. A small number of candidate kernels are composed, through either summation or multiplication, to create more complex customised kernel structures. To find an appropriate kernel structure with feasible computation effort, the problem was framed as a greedy search where only the next best kernel, and how it should be composed, is considered. Even with this consideration however, the computational demands are high, with the potential number of candidate models growing very rapidly with search tree depth. Building on this work, Duvenaud [156] and Lloyd [157] improve the interpretability of models through the

generation of natural language descriptions alongside selected kernel components. For example, a kernel structure of $K_{SE} \times K_{Per}$ (a Squared Exponential multiplied by a periodic kernel) would generate the accompanying description ‘*A periodic function whose shape changes smoothly*’. Although less helpful for experienced users of GPs who might infer similar descriptions from the mathematical expression themselves, these automatic interpretations are particularly useful for users less familiar with the behaviours present within common kernels, providing a concise way to interpret the model that may be understood by a layperson.

Alvarez [158] introduced the Latent Force Model (LFM), which derives a GP kernel from a second-order Ordinary Differential Equation (ODE) for a dynamical system. The LFM models the response of the system as a finite combination of latent functions (unknown input forces) convolved with Green’s functions (impulse response functions). The unknown input forces are modelled as GPs in time, and since they are transformed via a linear operator, the response of the system is also a valid GP. This allows important aspects of the systems physics to be directly encoded within the kernel, improving extrapolation capabilities of the model. The major challenge of the LFM comes from the computational expense of evaluating the cross correlations between the response and output, which are typically obtained through numerical methods. Research efforts to alleviate this computational expense have included state space representations of the LFM [159–162], where the states of the system may be more efficiently estimated through Kalman filtering [163]. Guarnizo [164] used Random Fourier Features (RFF) to obtain approximate solutions to the cross covariance terms rather than relying on more expensive numerical solutions. With application to SHM, Petersen [165] used a GP-LFM for the estimation of wind loading on the Hardanger Bridge. Although capable of working with only access to acceleration measurements, the use of measured wind velocities and wind tunnel test data were used to augment the GP-LFM, helping to improve performance whilst restricting the growth of model structure complexity.

The field of physics-informed kernel design spans a large variety of methods, with many kinds of potential advantages, including increased performance [149, 150], reduced demand for data collection [151–154] and insight in to model behaviour [156, 158]. The author argues that there are many more potential ways to design a kernel than there would be to include knowledge within the inputs or mean function of a GP. High variability within methods, and a current lack of literature applied to the field of offshore SHM, make kernel design another region of focus for the thesis. The work of Chapter 7 explores how one might use kernel structures to represent partial understanding of a physical system.

PINNs

Although this thesis focusses on the integration of knowledge within Gaussian processes, for the capabilities outlined in the theory of Chapter 3, Physics Informed Neural Networks (PINNs) form a large and varied research area that should not go unmentioned². For the interested reader, more detailed overviews of PINNs are presented in [12, 166, 167], with a brief outline of key themes given here. As with many of the earlier presented methods, PINNs are a recent and rapidly growing research field, with the introduction of the term ‘*Physics Informed Neural Network*’ widely credited to Raissi [52] in 2019. Raissi used PINNs for two distinct tasks: the efficient solving of physically derived Partial Differential Equations (PDEs) and the discovery of new PDEs to explain physical processes e.g. pressure fields within flow past a cylinder. This was achieved through modification of the loss and activation functions.

The physical knowledge integrated within PINNs commonly takes the form of PDEs, with engineering relevant examples including the wave equation [168, 169], Navier-Stokes equations [170, 171], von Mises elastoplasticity [172] and the Buckley-Leverett equation [173, 174]. As earlier stated in the section on training focussed knowledge integration, a common means of including physics in PINNs is through the loss function, where for a PDE this will generally take the form:

$$L(\theta) = \underbrace{L_{Res}(\theta)}_{\text{PDE residuals}} + \underbrace{L_{BC}(\theta)}_{\text{Boundary conditions}} \quad (2.3)$$

By minimising the PDE residual, the network is encouraged to converge towards physically meaningful solutions whilst the introduction of a boundary term can be used to penalise solutions where they may be broken. References [168–174] involve modification of the loss function similar to this way. An important property of loss based approaches is that the adherence to physics is rewarded, rather than enforced, meaning that outside the scope of training data, incorporated physics is not guaranteed to hold.

Physical Activation Functions (PAFs) provide another means of including physics in PINNs, via the mimicking of a physical process within the expression of the activation function, for example, the dynamic behaviour of a harmonic oscillator [175] or solute transport in porous media [176]. A key distinction between PAFs

²At the outset of the thesis, a narrower scope, focussing on the integration of knowledge within a consistent type of machine learner was viewed as a more likely direction to generate valuable contributions.

and loss based knowledge inclusion is that here the functional form is enforced, rather than rewarded during training. This is a stricter enforcement of physics, with the network constrained to functions of the specified form [175]. The design of PAFs presents a promising, but challenging research area, with the adoption of loss based approaches currently much more widely used.

State space approaches

State Space (SS) representations of systems provide an effective modelling framework to account for varying levels of prior knowledge of a systems' behaviour. Simpler systems, that are typically more well understood, may be suited to the linear case (the Kalman filter [163]), where solutions are available in closed form. Non-linear extensions to the Kalman filter [177–179] allow for an increased capability to model more complex systems. In cases where a higher degree of flexibility is desirable, for example the non-linear dynamics of a cart-pole system [180, 181], Gaussian Process State Space Models (GPSSMs) incorporate a GP to capture some desired nonlinear aspect of system behaviour, commonly in place of the transition or measurement matrices [182–184].

Lindsten [185] used a GPSSM to model a Wiener system, where a state space model was used to capture the behaviour of the underlying linear system and a GP was used to capture the effects of the static nonlinearity. The model outperformed standard existing approaches (ADM [186] and PEM [187]) at the expense of higher computational demands. This was due to the GP evaluation scaling with $\mathcal{O}(N^3)$ and the requirement of particle-Gibbs with ancestor sampling to obtain a solution. Rogers [188] used the same approach to model a known nonlinearity present within a Duffing oscillator, which was later extended [189] for the case of an unknown nonlinearity. Here, without any a priori knowledge of the form of the nonlinearity, the states and parameter estimates of the system along with an estimate of the nonlinearity were extracted.

For SHM applications, there are a few possible avenues through which novel state space contributions may be made: the continued development of alternative non-linear Kalman filter extensions, capable of representing real world structures more accurately; achieving increased computational efficiency, allowing for utilisation of datasets from large monitoring campaigns; and the incorporation of physical knowledge alongside existing state space approaches. The work of this thesis is primarily centred around the use of ocean state measurements (wave heights, flow conditions etc.), rather than response measurements from the structure, and as a result, state space methods are not a major focus. It is still acknowledged as an area of potential future contributions however.

Boundary conditions and constrained learners

The inclusion of boundary conditions or constraints within a model represents a scenario where some aspect of a systems' behaviour is known for a given condition, often in the form of a physical law e.g. conservation of energy [190, 191]. The use of constrained learners is useful to ensure predictions obey known restrictions, whilst maintaining a degree of model flexibility to capture the specifics of complex phenomena. With application to SHM, Jones [192] used Neumann boundary conditions to improve the localisation of acoustic emissions in instances of low data coverage. Particular improvements were seen in locations further from sensor locations but close to where constraints were enforced. Mclean [193] used Beta likelihoods within a GP to constrain wind turbine power curve predictions within physical limits i.e. below the maximum power rating. This not only improved performance, but also ensured predictive distributions were physically meaningful, inspiring a higher degree of trust from operators.

Overlapping with the field of kernel design, Wahlström [194] developed curl and divergence free kernels for the modelling of magnetic fields, which Jidling [195] extended to a more general linearly constrained case. The cross covariance terms within Multiple Output Gaussian Processes (MOGPs) provide a useful way to integrate known constraints between outputs. Later work of Jidling [196] predicted strain fields with a MOGP and enforced equilibrium constraints for linearly elastic materials through manipulation of the cross-covariance terms within the kernel. Cross [197] introduced pseudo-measurements of cantilever beam boundary conditions along with derivative constraints within the MOGP cross covariance terms. This constrained the position and slope of the beam to zero at the wall, improving performance in this region.

The enforcing of symmetries is a popular type of constrained learner, with it being an effective way to reduce the required amount of data to learn a task. Reviews of symmetry enforcement within wider ML applications are available in [198] and [199]. As earlier stated, the use of Gaussian process kernel design is one way in which symmetries may be enforced [151–154]. Meyer [200] utilises data embeddings to exploit symmetries that may be more complex, and therefore challenging to represent. An example of *tic-tac-toe* game states is presented to introduce methods, where rotating the board by a multiple of 90° or reflecting along an axis will produce equivalent game states. In this example, the total number of observable game states, and therefore the effective demand for data, is reduced by a factor of 8.

For contributions to the field of offshore SHM, the constraining of GPs through the modification of the kernel is an area for potential model development. Though examples exist for SHM [193, 196, 201] and structural dynamics [197] more generally,

there are currently far fewer examples for offshore structures. Constraints of sea bed foundations, principles of fluid mechanics and structural response could all be potential candidates for inclusion.

2.2.5 Output focussed knowledge integration

Within a black-box pipeline, the output of models is typically a prediction of the target variable of interest Y . It is possible however to use the outputs of machine learners alongside, or to assist with, the use of physical knowledge. This presents a modification to the end of a typical pipeline, with these methods summarised graphically in Figure 2.7.

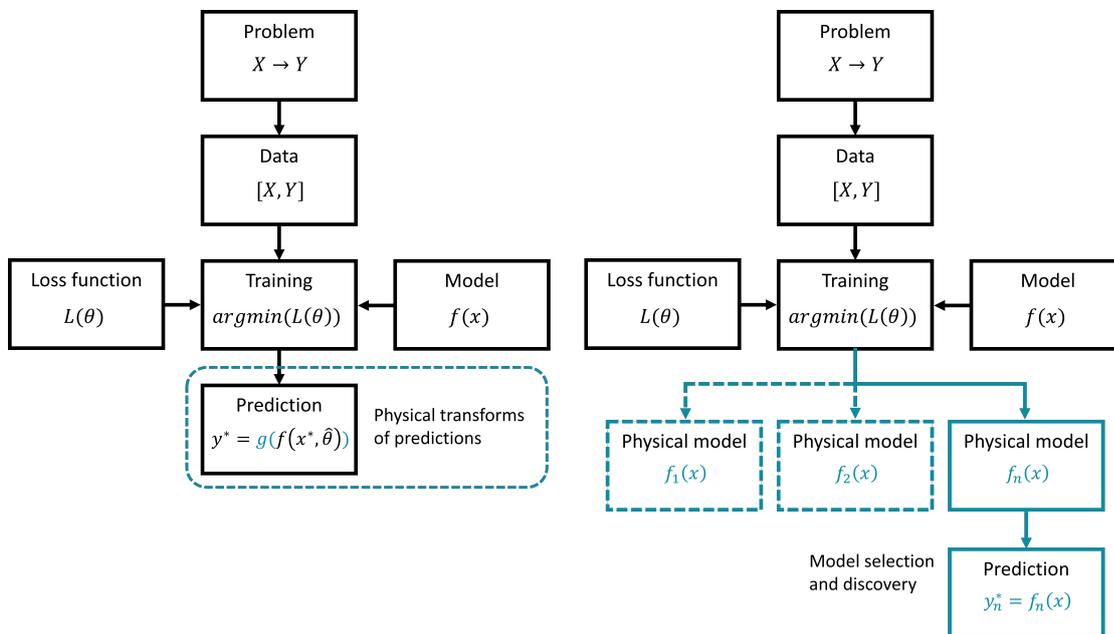


Figure 2.7: Output focussed integration of physical knowledge within a typical machine learning pipeline. Physical transforms of predictions use knowledge to modify the output of machine learners in some way, whilst model selection uses the outputs of machine learner to assist with the selection or discovery of physical model structures.

Physical transforms of machine learning outputs

In some cases, a machine learner may be used to approximate part of a process, with further modification of the output occurring using physical knowledge to obtain the parameter of interest. An example of this, sometimes referred to as

virtual sensing [202–205], is the use of machine learners to estimate parameters that may be difficult (or impossible) to measure, which are then used as inputs to physical models. Specific to SHM, Gibson [206] used a GP to predict strain within an aircraft wing during flight, with samples from the GP posterior then individually passed through a rainflow count to obtain a distribution of fatigue damage. Using sample draws from the GP, which inherently contain a larger number of stress cycles than a GP mean, was found to estimate a larger accrual of fatigue damage.

Putz [207] used knowledge of dependencies between parameters such as aspect ratio, area and slope when generating input data for microstructure modelling using a GAN. This ensured generated grains followed physically representative statistical distributions which is important for accurate definition of metallic microstructures. In this application, the parameter of interest is not the grains themselves, but the properties of materials they form, and how they change throughout manufacturing processes such as casting and welding.

Model selection and discovery

When opting to use a physical model, deciding how to best represent the physical system is a challenging and important task. For many scenarios, such decisions will require access to high levels of subject expertise, leading to a high level of cost. The use of machine learners to assist with how to physically represent a system is a potential way to alleviate this reliance on experts, or to provide a tool to aid with their decisions. A popular form of this is the use of a machine learner for the discovery of physically representative PDEs [52, 208–210]. A key advantage of this approach is model parsimony, with the discovered PDEs being a simpler, more interpretable way to represent the problem than directly using the network.

Singh [211] looks at the selection of the most appropriate model form for the modelling of turbulent flows over airfoils using a neural network. The task is framed as a selection of inverse problems, representing a series of model discrepancies, which when solved and combined may be used to reconstruct the full flow. Through learning of a model form rather than direct prediction of the flow, the model is able to achieve improved predictions for airfoil predictions not shown to it during training.

With application to sound synthesis, Gabrielli [212] used a neural network to aid with the selection of physical model parameters. Here, the structure of the physical model aimed to mimic a pipe organ, with three main parts: an exciter, to model the wind jet; a resonator, acting as a wave guide; and a noise model, to simulate air noise modulated by the wind jet. Determination of suitable parameters was non-trivial and had a significant effect on properties such as pitch and timbre of

sounds.

In the field of structural dynamics, the work of Lai [213] developed methods for the identification of state-space representations of dynamic systems. Termed physics-informed neural Ordinary Differential Equations (Neural ODEs), two components within a hybrid ODE representation are utilised to represent the system: a physics-based term, encapsulating a users prior knowledge of a system, and a discrepancy term, offering increased flexibility, captured by a feed-forward neural network. Through sparse identification, retrieving closed form expressions for terms represented by the network, an additional level of insight into the ‘known’ and ‘unknown’ parts of system behaviour could be gained. The isolation of non-linear effects was identified as a promising avenue for progression.

2.3 Grey-box models for structural dynamics

This chapter has explored a range of ways in which to incorporate prior knowledge within a machine learner and discussed how methods introduce differences in methodology at each stage of a typical black-box modelling pipeline. Examples from the literature covered a variety of applications from heat transfer [129], to galaxy formation [109], to the modelling of magnetic fields [214]. The novelty of this thesis comes from the development and application of grey-box models to the field of structural dynamics, where the integration of physics within machine learners is currently much rarer. A specific focus is placed on the wave loading of offshore structures, a field where the benefits of physics-informed machine learning have yet to be exploited.

Throughout this chapter, the research gaps specific to PIML methodologies were highlighted as areas in which to focus efforts of model development. Where possible, avenues through which potential contributions might be made were identified (e.g. generation of new input variables, non-linear Kalman filter extensions etc.). Although not possible to tackle all identified areas of novelty, the categorising of PIML methods was very helpful for guiding the efforts of the thesis. As seen within the thesis outline of Chapter 1 and upcoming GP theory of Chapter 3, many of the PIML models developed aim to target a specific gap in a particular research field.

BACKGROUND THEORY AND INTERPRETATION

This thesis focusses on the integration of data-based and physics-based methods within structural dynamics applications, with the aim of this chapter to provide an overview of key theory relevant to each. First, Gaussian process models, the machine learning technique of preference used throughout later work, are introduced with a focus placed on the interpretation of model components and important behaviours exhibited during their use. With the wave loading of offshore structures a reoccurring application, a chosen selection of key topics from a physics-based perspective are then presented, namely: linear wave theory, the representation of sea states and empirical wave loading prediction.

3.1 Gaussian processes

This section aims to introduce Gaussian process models, explain important aspects of how they operate and motivate their use within engineering regression tasks. Although there are many ways to interpret a Gaussian Process (GP), the most intuitive of which will vary by reader, only the authors favoured interpretation is presented here. Extensive derivations of GP theory are widely available within the literature [139, 215, 216], and as such, a focus is instead placed on the interpretation of key equations and use of illustrated examples.

A family of functions

A Gaussian process (GP) is comprised of a set of random variables, any finite number of which, share a joint Gaussian distribution. A popular view of a GP is as a distribution over functions, where each individual draw from the GP represents a realisation of one of the infinitely many potentially generated functions. A GP, and the family of functions it may generate, is fully defined by its mean function and covariance function (kernel). For a real process $f(X)$, a GP is expressed

$$f(X) \sim \mathcal{GP}(m(X), K(X_i, X_j)) \quad (3.1)$$

where X is an $N \times D$ matrix; N being the number of datapoints and D being the dimension of the input space. For a pair of inputs X_i and X_j , the mean function $m(X)$ and covariance function $K(X_i, X_j)$ are defined

$$m(X) = \mathbb{E}[f(X)] \quad (3.2)$$

$$K(X_i, X_j) = \mathbb{E}[f(X_i)f(X_j)] - \mathbb{E}[f(X_i)]\mathbb{E}[f(X_j)] \quad (3.3)$$

Sampling a function from a GP is equivalent to making a draw from the multivariate normal distribution $\mathcal{N}(m(X), K(X_i, X_j))$. An example of 20 draws from a GP with a Squared Exponential (SE) kernel is shown in Figure 3.1. A SE kernel produces smoothly varying functions and its flexibility makes it a popular choice for many modelling tasks.

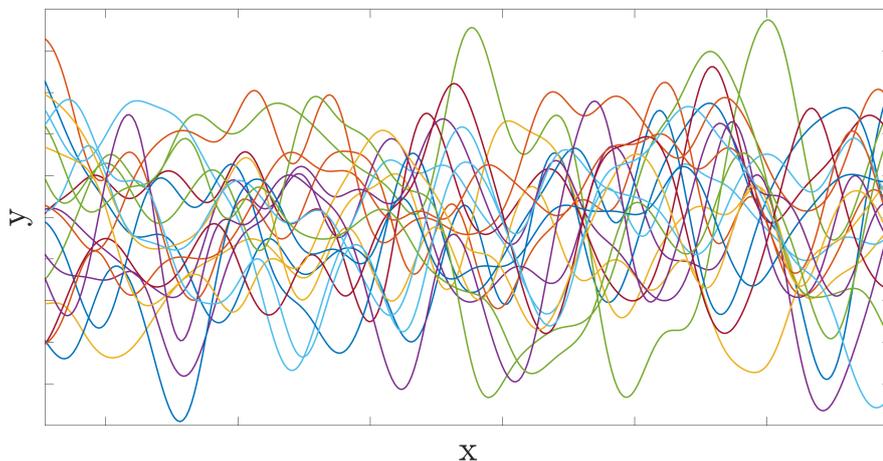


Figure 3.1: 20 prior draws from a GP with a Squared Exponential (SE) kernel. Coloured lines represent individual draws.

Conditioning on data

A highly useful property of a GP, particularly for regression tasks, is an ability to be conditioned on data. Fitting within a Bayesian framework, when a GP is shown a set of observations (evidence), it is able to update the prior distribution of functions (beliefs), such that they align more closely with observed data. This allows information and structure present within datasets to be captured and utilised by a GP. The term conditioning stems from the use of conditional distributions, $p(A|B)$, the probability of event A , given that event B has already occurred. In an engineering regression context, this is often framed as predicting how a system is likely to behave, given that we have observed some previous behaviour.

To introduce an example of conditioning, following the notation of Rasmussen [139], consider a noisy process $y = f(X) + \epsilon$, where measurements y represent noisy observations of the underlying process $f(X)$. Assuming access to N measured pairs of inputs and outputs, y will be an $N \times 1$ vector and X will be a $N \times D$ matrix, where D is the dimension of the input space. This represents a typical scenario of a real world engineering regression task where data might come from a sensed structure. Commonly referred to as a training set, a set of inputs with corresponding targets $[X_i, y_i, i = 1, \dots, N]$ is used in the conditioning step, to encode the relationship between X and y . Generally, a prediction is then made on a test set $[X_{*i}, y_{*i}, i = 1, \dots, N_*]$, that the GP has not previously seen.

For a prediction y^* at new set of test points X_* , the covariance matrices for the training points $K(X, X)$, the test points $K(X_*, X_*)$, training-test points $K(X, X_*)$ and test-training points $K(X_*, X)$ are required. For predictions based on noisy observations, with an assumed zero-mean Gaussian noise of variance σ_n^2 , the problem may be formulated:

$$\begin{bmatrix} y \\ y^* \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} m(X) \\ m(X_*) \end{bmatrix}, \begin{bmatrix} K(X, X) + \sigma_n^2 \mathbb{I} & K(X, X_*) \\ K(X_*, X) & K(X_*, X_*) \end{bmatrix} \right) \quad (3.4)$$

Within a machine learning context, the conditional distribution of the test target y^* , given the test inputs X_* , observed training inputs X and training targets y is of key interest. This is often referred to as the posterior or predictive distribution.

$$p(y^* | X_*, X, y) = \mathcal{N}(\mathbb{E}[f(X_*)], \mathbb{V}[f(X_*)] + \sigma_n^2 \mathbb{I}) \quad (3.5)$$

where the predictive mean $\mathbb{E}[f(X_*)]$ and variance $\mathbb{V}[f(X_*)]$ are expressed¹:

$$\mathbb{E}[f(X_*)] = m(X_*) + K(X_*, X)(K(X, X) + \sigma_n^2\mathbb{I})^{-1}(y - m(X)) \quad (3.6)$$

$$\mathbb{V}[f(X_*)] = K(X_*, X_*) - K(X_*, X)(K(X, X) + \sigma_n^2\mathbb{I})^{-1}K(X, X_*) \quad (3.7)$$

These expressions are among the most important for those interested in using GPs for regression, and their structure offers insight into the modelling capabilities of a GP. The predictive mean and variance may be interpreted as follows:

$$\begin{aligned} \mathbb{E}[f(X_*)] = & \underbrace{m(X_*)}_{\text{Predictive mean follows mean function...}} \\ & + \underbrace{K(X_*, X)(K(X, X) + \sigma_n^2\mathbb{I})^{-1}(y - m(X))}_{\text{...unless observed data suggests otherwise.}} \end{aligned} \quad (3.8)$$

$$\begin{aligned} \mathbb{V}[f(X_*)] = & \underbrace{K(X_*, X_*)}_{\text{A maximum of test set covariance...}} \\ & - \underbrace{K(X_*, X)(K(X, X) + \sigma_n^2\mathbb{I})^{-1}K(X, X_*)}_{\text{...with observations shrinking final result.}} \end{aligned} \quad (3.9)$$

To observe the effects of conditioning on data, Figure 3.2 shows 20 draws from the predictive distribution $\mathcal{N}(\mathbb{E}[f(X_*)], \mathbb{V}[f(X_*)] + \sigma_n^2\mathbb{I})$ of a GP conditioned on 3 data points. The same SE kernel was used as the prior draws presented in Figure 3.1. The sampled functions now pass within a noise level σ_n^2 of observed data.

Considering infinite samples

As earlier stated, a GP represents an infinite set of potentially generated functions, defined by its mean and covariance functions. While possible to draw and observe individual functions, what is often of more interest is the entire set of functions as a whole. However, naively considering increasing numbers of samples quickly becomes infeasible, both in terms of practicality and computation. A more convenient means of handling, is through the use of the predictive mean and confidence intervals. For

¹Many readers may be most familiar with a zero prior mean GP. Here the predictive mean reduces to $\mathbb{E}[f(X_*)] = K(X_*, X)(K(X, X) + \sigma_n^2\mathbb{I})^{-1}y$

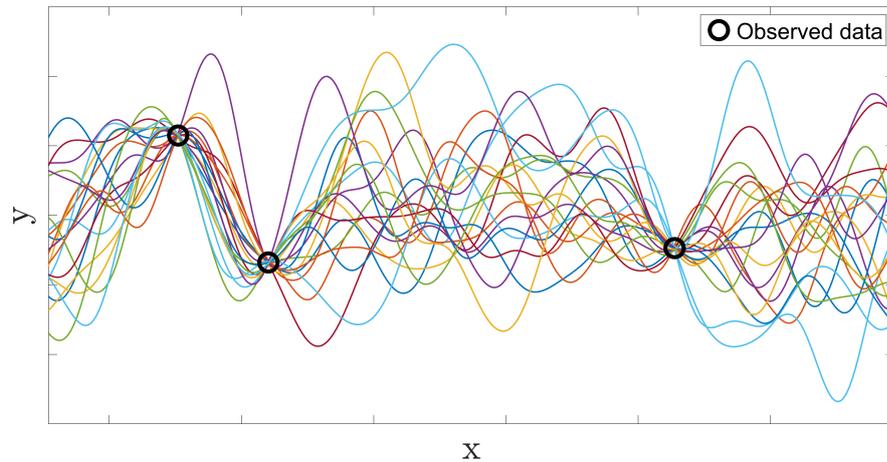


Figure 3.2: 20 draws from the predictive distribution of GP with a squared exponential kernel, conditioned on three observed data points. Coloured lines represent individual draws, with observations as black circles.

a prediction with a set of inputs X_* , a 3σ interval is calculated:

$$\pm 3\sigma = \pm 3\sqrt{\text{diag}(\mathbb{V}[f(X_*)]) + \sigma_n^2 \mathbb{I}} \quad (3.10)$$

A confidence interval provides information about a region through which a desired fraction of functions pass and provides a useful method of visualising a GP. Increasing numbers of draws from the same predictive distribution, along with a plot of the predictive mean and confidence intervals is shown in Figure 3.3. All plots are produced from a GP with a SE kernel, conditioned on the same three data points.

Learning from data

One of the most attractive, but computationally intensive properties of a GP is an ability to encode information present within datasets. This is particularly helpful when phenomena present within data are not fully understood e.g. response of composite structures, bolted joints or extreme weather conditions. Although both are closely involved, it is important to distinguish between conditioning and training in this context:

Conditioning: The calculation of the conditional distribution $p(y^*|X_*, X, y)$; the probability of an output for a new set of inputs, given that a set of data has already been observed.

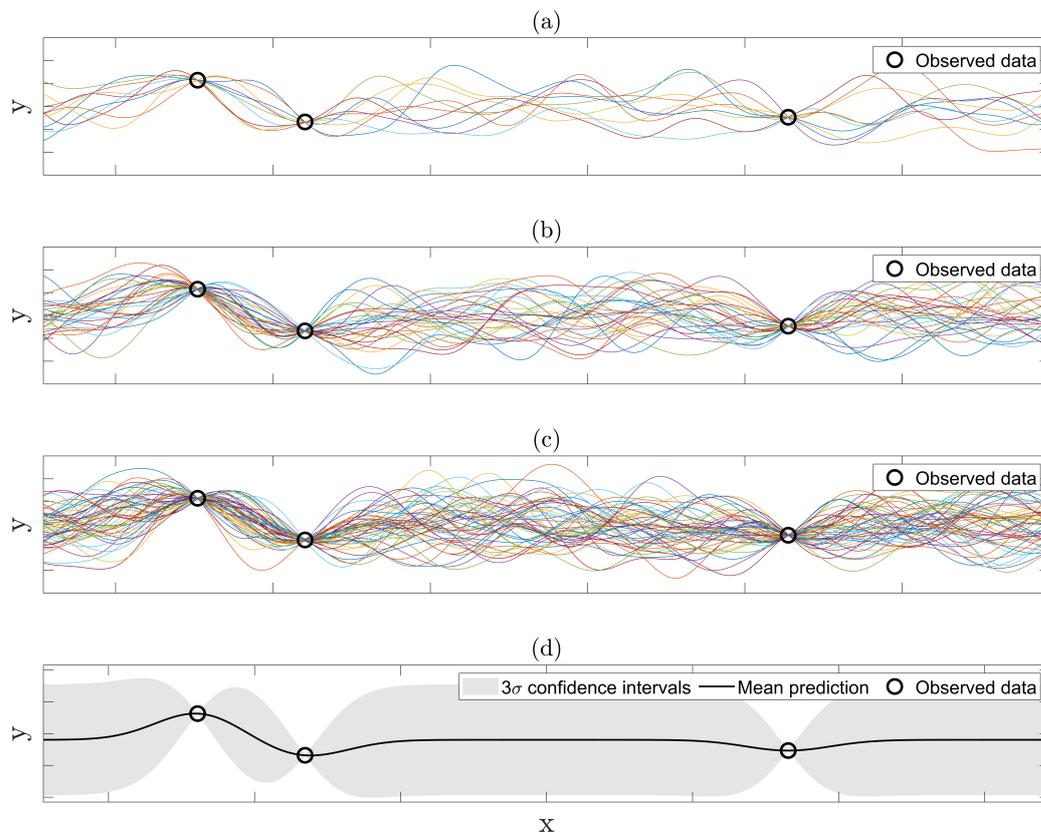


Figure 3.3: Draws from the predictive distribution of GP with a squared exponential kernel, conditioned on 3 observed data points. Increasing numbers of draws are plotted: (a) 10, (b) 30, (c) 50, with the predictive mean and confidence intervals shown in (d). Coloured lines represent individual draws, with observations as black circles.

Training: The optimisation of hyperparameters within the GP for a given set of training data. A selected cost function is minimised in order to maximise performance on the observed data.

Although referred to as non-parametric models, in that model complexity is not constrained by a fixed number of parameters, a GP still requires the determination of hyperparameters. Hyperparameters control the behaviours of functions generated by a GP and their effect will vary depending on the selected kernel. Here a squared

exponential (SE) kernel is considered.

$$K_{SE}(X_i, X_j) = \sigma_f^2 \exp\left(-\frac{1}{2}(X_i - X_j)^T \Lambda^{-1}(X_i - X_j)\right) \quad (3.11)$$

where σ_f^2 is the signal variance, dictating the magnitude of generated functions, and Λ is the matrix of length scales such that $\text{diag}(\Lambda) = [l_1^2, l_2^2, \dots, l_D^2]$ for a D dimensional input. The lengthscale(s) within a SE kernel controls the distance from which observed data may influence the prediction of a GP and the ‘wiggleness’ of generated functions. Short length scales allow more sudden changes of direction, but will cause predictions to quickly return to their prior when moving away from observations. Long length scales will produce slowly varying functions that return to their prior very gradually when moving away from observations. Subplots highlighting the effects of varying hyperparameters within a SE kernel are shown in Figure 3.4.

The kernel hyperparameters, along with a noise variance σ_n^2 are used to construct a vector θ , to be passed as an argument to an optimiser. For a SE kernel with $D = 1$:

$$\theta = \{\sigma_f^2, l^2, \sigma_n^2\} \quad (3.12)$$

These parameters are typically optimised over the Negative Log Marginal Likelihood (NLML) of model predictions on a training set:

$$-\log p(y|X, \theta) = \frac{1}{2}y^T(K + \sigma^2 I)^{-1}y + -\frac{1}{2} \log |K + \sigma^2 I| + \frac{n}{2} \log(2\pi) \quad (3.13)$$

The predictions of models trained on increasing numbers of data points are shown in Figure 3.5. The model trained on only five data points (a), exhibits behaviours of a model with a long lengthscale and does a poor job of capturing the true underlying function. Any machine learning model, not just GPs, can only be expected to capture information presented to it during training and an inability here to capture the complexities of unobserved data is to be expected. As the training data supplied to the model increases, so too does model performance. The 50 data point model (c), would represent a scenario of very high data coverage in which data-based models are known to work well. Within engineering applications however, while desirable, such high data-coverage levels across all observable conditions is generally unattainable. Highly variable environments, extreme weather and changing structural conditions make data coverage a challenging task.

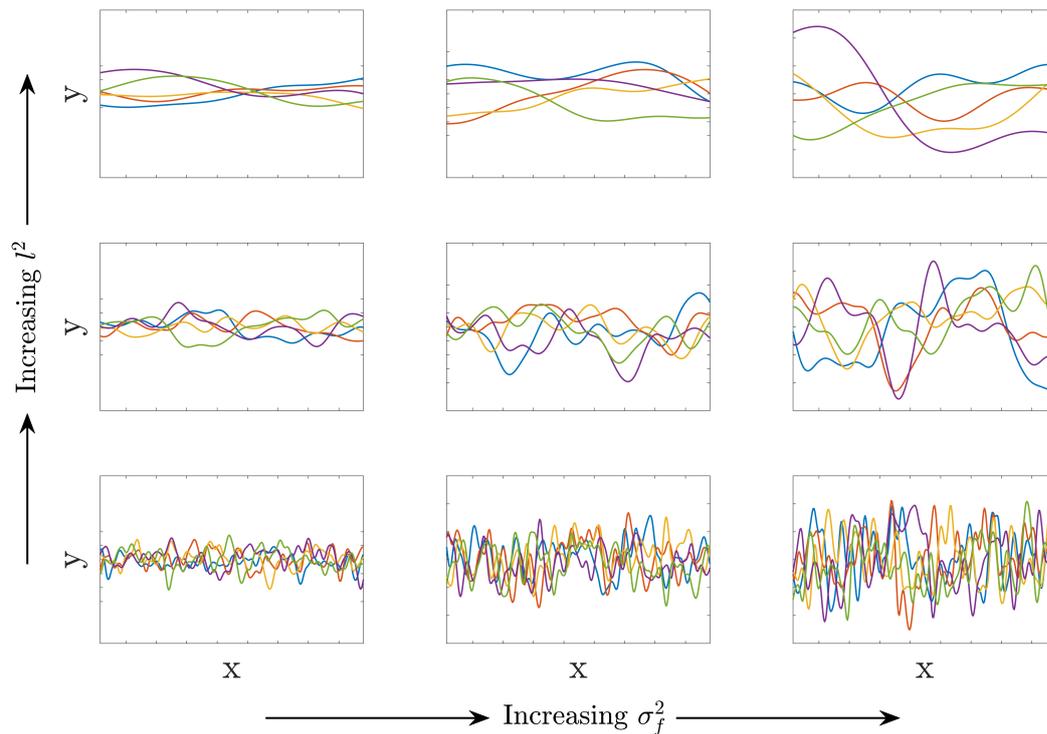


Figure 3.4: Draws for a GP with a SE kernel highlighting the effects of varying hyperparameters. The signal variance σ_f^2 controls the amplitude of generated functions, whilst the lengthscale l^2 is responsible for how smoothly a function varies. Shorter lengthscales produce more ‘wiggly’ functions.

Quality as well as quantity

Intuitively, as a GP is presented with an increasing quantity of training data its ability to capture the complexity of an underlying function increases. However, another important factor could be described as the quality of data. Here, the following interpretations are proposed:

Data quantity: The number of data points N supplied to a GP during training.

Data quality: The degree to which the training data supplied to a GP is helpful in capturing the behaviour of a desired underlying function.

There are many factors that may effect the perceived quality of training data including noise levels, relevance of measured variables and distribution of data.

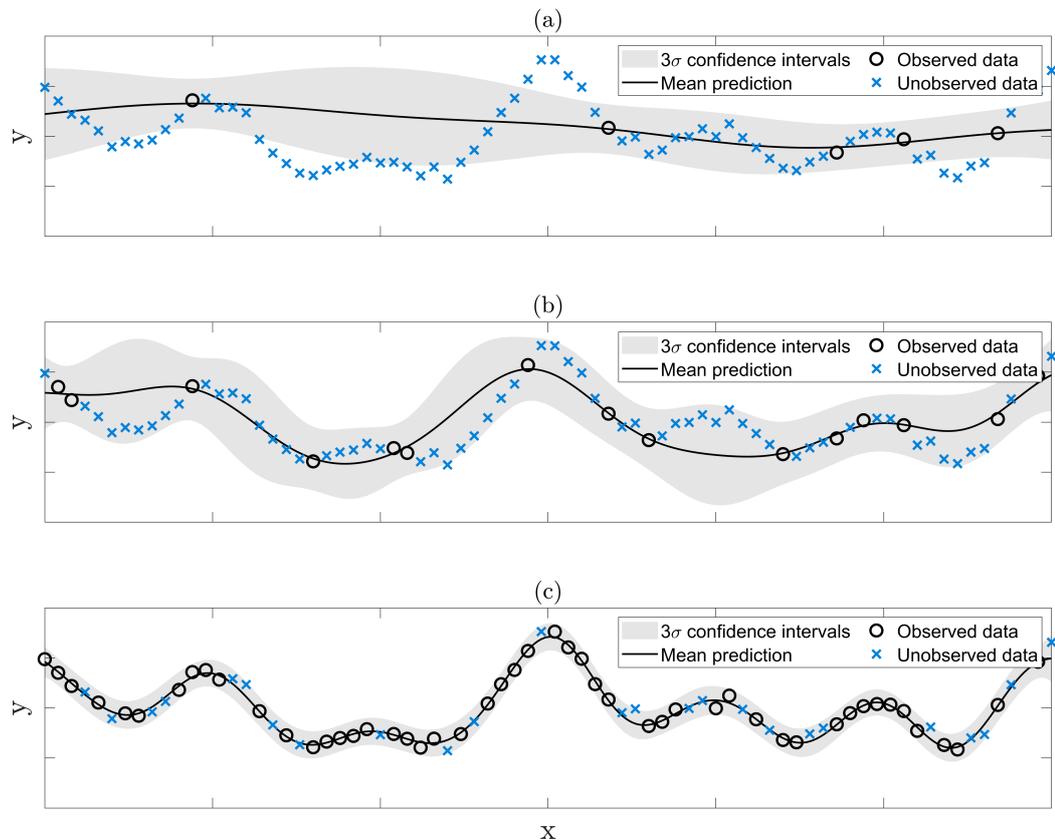


Figure 3.5: Predictions of a GP with a SE kernel with 5 (a), 15 (b) and 50 (c) training points. Observations are shown as black circles, with data not shown to the GP during training shown as blue X's.

The determination of data quality is highly application specific and an important step within GP model construction. For purposes of highlighting an example, a variation of data distribution for 20 training points is shown in Figure 3.6.

The distribution of training data has a significant effect on the properties of a constructed model. Figure 3.6 (a) shows a model trained on evenly spaced points, where a reasonable capture of the mean behaviour of the function is achieved. However, due to a lack of closely spaced data points, the noise variance σ_n^2 is visibly underestimated. Accurate estimation of the noise will generally require at least some training data to be close together in the input space. Figure 3.6 (b) shows a model trained on randomly scattered data points, with an improvement in the estimation of noise present within the signal due to the presence of closely spaced training points. However, due to the randomness of data location, regions with

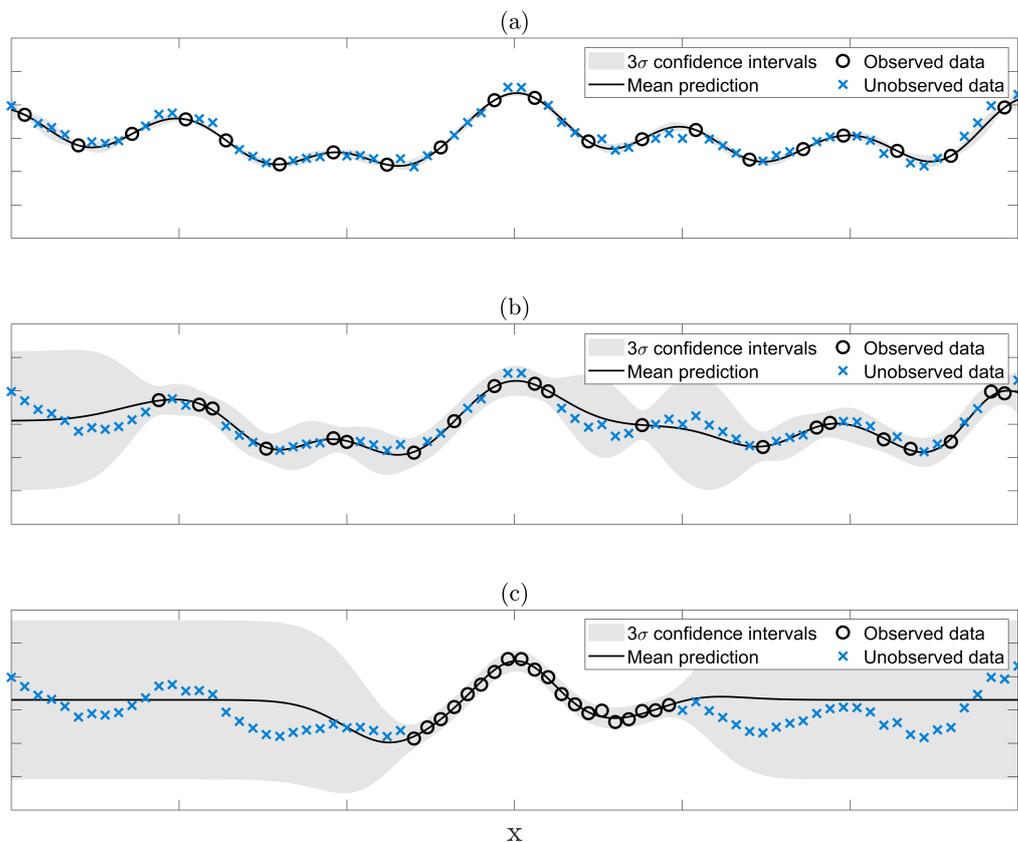


Figure 3.6: Predictions of GPs with varied distributions of 20 training data points: (a) Evenly spaced points, (b) Randomly located points and (c) Consecutive points.

missing training data also cause to the estimation mean function behaviour to suffer. Figure 3.6 (c) shows a model trained on consecutive data points. The model performs well, both in terms of mean function behaviour and noise estimation, within the observed data and poorly outside.

Although limited examples are highlighted here through the use of a toy function, the phenomena observed are an important consideration when constructing GP models. The impact of training data quality, which will be specific to the modelling scenario, on final model performance should be well understood. The conditions in which a model will perform well, along with expected limitations should be used to advise how a model may be used most effectively.

Physical knowledge in Gaussian processes

A running theme throughout this thesis is the incorporation of physical knowledge within GP models to improve their capabilities, in terms of performance, interpretability, ability to work with restricted data or other benefits. There are many ways in which this is possible, a few of which are summarised here.

Beginning with a reference case, a **black-box** model is a term used to refer to when no prior knowledge is integrated within a machine learner. Typically, for a GP, this would involve the use of a zero mean function and a flexible kernel, producing a model of the form:

$$y \sim \mathcal{GP}(0, K(X, X)) \quad (3.14)$$

where $K(X, X)$ is a kernel of choice, often a Squared Exponential or Matérn kernel. Throughout the thesis, black-box models are utilised at a start point from which to develop PIML model architectures. They also serve as a performance benchmark with which to compare the results of potential PIML model structures. Generally, for a successful PIML model, the hope would be to outperform (or provide some other benefit) over an equivalent black-box case.

Input augmentation uses physical knowledge to generate input(s) for the GP. This can be of particular help when desired inputs may be difficult or expensive to measure. Input augmentation has a model structure of:

$$y \sim \mathcal{GP}(0, K([X, f_{Phy}(X)], [X, f_{Phy}(X)])) \quad (3.15)$$

where $f_{Phy}(X)$ is an additional input generated using physical knowledge in some way. The use of input augmentation is explored within the work of Chapters 5 and 6 in the context of wave loading prediction. Aspects of linear wave theory and Morison's Equation, detailed in the upcoming Section 3.2, were both utilised to generate inputs for the GP.

Constrained learners use physical knowledge to restrict how the predictions of a GP may behave. This can take many forms, from known behaviours under specific conditions e.g. boundary constraints [196, 197], to the constraining of predictions between limits [193]. There are multiple avenues through which a GP may be constrained, including the kernel, the loss function and through the use of pseudo-measurements. Here, to represent a GP, constrained by physical knowledge in some way, set notation is used to convey which variable has been constrained:

- For outputs y , constrained to a physically meaningful subset y_{Phy}

$$\{y \sim \mathcal{GP}(0, K(X, X)) \mid y \in y_{Phy}\} \quad (3.16)$$

- For inputs X , constrained to a physically meaningful subset X_{Phy}

$$\{y \sim \mathcal{GP}(0, K(X, X)) \mid X \in X_{Phy}\} \quad (3.17)$$

- For hyperparameters θ , constrained to a physically meaningful subset θ_{Phy}

$$\{y \sim \mathcal{GP}(0, K(X, X)) \mid \theta \in \theta_{Phy}\} \quad (3.18)$$

where the hyperparameters θ control the properties of the kernel $K(X, X)$.

Due to the nature of physical knowledge utilised in the thesis, the development of constrained learners was not a major focus. Discussions on how to define the types of prior knowledge possessed by a user, including those used in the thesis, and how this affects the most appropriate PIML model structure is presented in Chapter 9.

A **residual model** is one way of incorporating physical knowledge within the **mean function** of a GP; it replaces the target of the machine learner with the residual of a physics-based model $\delta(X) = y(X) - f_{Phy}(X)$, often referred to as model discrepancy [81–83]. Where data coverage for a problem may be sparse, a residual model is able to rely upon a physical model $f_{Phy}(X)$ to increase the quality of predictions. A GP residual model is of the form:

$$y \sim \mathcal{GP}(f_{Phy}(X), K(X, X)) \quad (3.19)$$

The development of a residual model is presented in Chapter 5, where Morison’s Equation was used to assist the predictive capabilities of a GP-NARX. Discussions on how the capabilities of residual model structures differ from alternative approaches is explored in Chapter 8.

Kernel methods look to encode physical knowledge within the covariance function of a GP. The kernel of a GP determines the nature of functions that may be used by a GP to make predictions. For a process $f(X)$, the calculation of its covariance $\text{Cov}(f(X), f(X))$ may be used to enforce physically derived behaviours within the

kernel:

$$\begin{aligned} K_{Phy}(X, X) &= \text{Cov}(f(X), f(X)) \\ K_{Phy}(X, X) &= \mathbb{E}[f(X)f(X)] - \mathbb{E}[f(X)]\mathbb{E}[f(X)] \end{aligned} \quad (3.20)$$

Used alone in a GP, assuming a zero-mean, a **physics-informed kernel** would lead to a model structure of:

$$y \sim \mathcal{GP}(0, K_{Phy}(X, X)) \quad (3.21)$$

As later discussed in the thesis, a useful property of physics-informed kernels is their ability to be combined with other kernels. Circumstances where a user might have partial physical knowledge of a system and how this can be represented through the structure of kernels is presented in Chapter 7.

Along with an introduction to GPs, and interpretations of model operation, the secondary role of this chapter is to provide key background of the types of physical knowledge utilised in the thesis. In keeping with the theme of offshore structures, the following section introduces the reader to linear wave theory, the representation of sea states and empirical methods of wave loading prediction.

3.2 Wave theory, sea states and wave loading

The prediction of wave loading on offshore structures is a critical step within the accurate estimation of their remaining life and it forms the application for many physics-informed models developed later in the thesis. Here, aspects of physical theory are introduced, which are later used in models and experimental design.

3.2.1 Linear waves

The simplest case of wave theory considers a two-dimensional, small amplitude gravity wave. A linear wave of amplitude A , wavelength λ , propagating within a fluid of depth d is shown in Figure 3.7. It is assumed that the fluid is incompressible, inviscid (negligible viscosity), without surface tension and without any temperature or salinity gradients [35, 36].

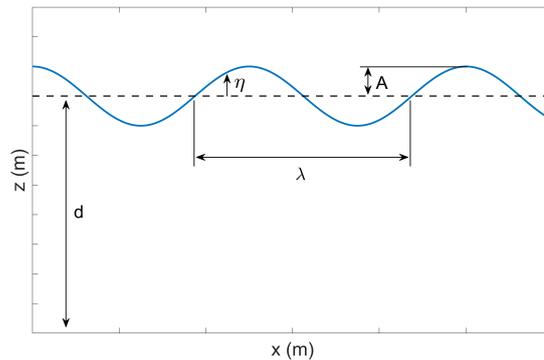


Figure 3.7: Diagram of a linear wave. The free surface elevation (η), wave length (λ), amplitude (A) and depth (d) are highlighted.

The determination of expressions for many important physical quantities of linear waves including free surface elevation, particle velocities and accelerations, pressure and the dispersion relation relies on a velocity potential solution for the Laplace equation:

$$\frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial z^2} = 0 \quad (3.22)$$

The velocity potential ϕ is a scalar function of space and time, the negative derivatives of which are equal to the fluid velocity component for a given direction:

$$U_x = -\frac{\partial \phi}{\partial x}, \quad U_z = -\frac{\partial \phi}{\partial z} \quad (3.23)$$

The solution velocity potential requires the imposing of boundary conditions, both at the sea bed ($\frac{\partial \phi}{\partial z} = 0$ at $z = -d$) and at the free surface. An important assumption made here is the adoption of small amplitude wave theory ($A \ll \lambda$ and $A \ll d$). This allows for the neglect of non-linear terms within the free surface boundary conditions:

$$\frac{\partial \eta}{\partial t} + \frac{\partial \phi}{\partial x} \frac{\partial \eta}{\partial x} - \frac{\partial \phi}{\partial z} = 0 \quad \text{at } z = \eta \quad (3.24)$$

$$\frac{\partial \phi}{\partial t} + \frac{1}{2} \left[\left(\frac{\partial \phi}{\partial x} \right)^2 + \left(\frac{\partial \phi}{\partial z} \right)^2 \right] + g\eta = 0 \quad \text{at } z = \eta \quad (3.25)$$

Without adoption of small amplitude wave theory and neglect of corresponding

nonlinear terms, the solution for velocity potential becomes complex, presenting its own area of study and shall not be further investigated here. For the interested reader, see the work of [217–219]. So long as wave amplitudes remain small, this simplification remains a valid means through which to derive the solution for velocity potential:

$$\phi = \frac{Ag}{kc} \frac{\cosh(k(z+d))}{\cosh(kd)} \sin(kx - \omega t + \Phi) \quad (3.26)$$

where $k = \frac{2\pi}{\lambda}$ is the wave number, $c = \frac{\omega}{k}$ is the wave speed and Φ is the phase. The solution for velocity potential allows for the expression of a variety of useful quantities, a selection of which is given in Table 3.1. Access to closed form expressions for quantities including free surface and particle velocities is a key advantage of utilising linear wave theory within wave loading applications.

Table 3.1: Quantity expressions for linear waves [36].

Quantity	$f(X)$
Dispersion relation	$\omega^2 = gk \tanh(kd)$
Free surface	$\eta = A \cos(kx - \omega t + \Phi)$
Particle velocities	$U_x = \omega A \frac{\cosh(k(z+d))}{\sinh(kd)} \cos(kx - \omega t + \Phi)$ $U_z = \omega A \frac{\sinh(k(z+d))}{\sinh(kd)} \sin(kx - \omega t + \Phi)$
Particle accelerations	$a_x = \omega^2 A \frac{\cosh(k(z+d))}{\sinh(kd)} \sin(kx - \omega t + \Phi)$ $a_z = \omega^2 A \frac{\sinh(k(z+d))}{\sinh(kd)} \cos(kx - \omega t + \Phi)$

The theory for linear waves is useful both for the study of sinusoidal waves, and also to provide a platform from which to construct more complex wave forms.

3.2.2 Representation of sea states

An irregular wave is constructed via a sum of linear waves. Assuming superposition, the free surface of an irregular wave is expressed

$$\eta(x, t) = \sum_{i=1}^N A_i \cos(k_i x - \omega_i t + \phi_i) \quad (3.27)$$

where N is the number of linear waves, A_i is the amplitude, k_i is the wave number, ω_i is the frequency and ϕ_i is the phase of wave i respectively [35, 36]. The complexity of irregular waves can grow quickly with increasing N , with an example of irregular wave construction for $N = 3$ shown in Figure 3.8.

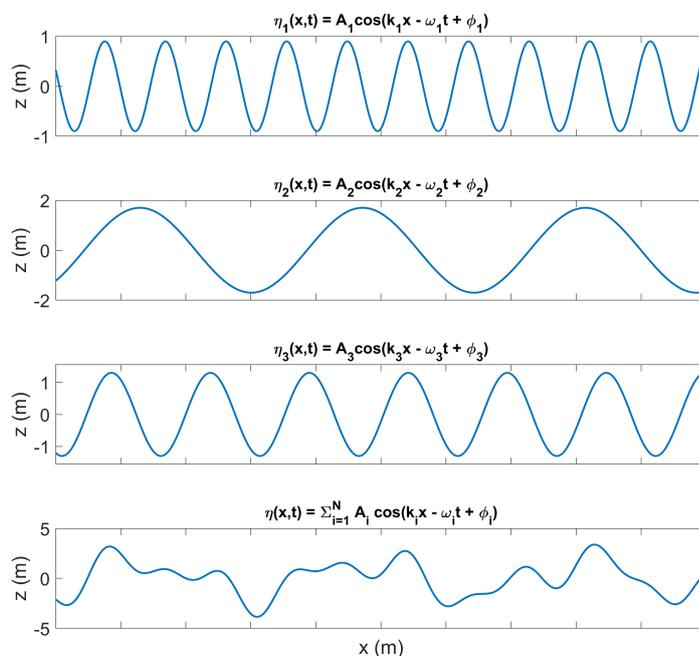


Figure 3.8: Construction of an irregular wave of the form $\eta(x, t) = \sum_{i=1}^N A_i \cos(k_i x - \omega_i t + \Phi_i)$ where $N = 3$.

Irregular wave construction forms the basis for which ocean states may be represented, with the A_i , k_i and ω_i determined via a representative spectra rather than randomly assigned. The ability to represent sea states is useful to measure,

compare and approximate real world conditions that an offshore structure might be exposed to. A popular example is the Pierson-Moskowitz [220] spectrum:

$$S(\omega) = \frac{\alpha g^2}{\omega^5} \exp\left(-\beta \left(\frac{\omega_p}{\omega}\right)^4\right) \quad (3.28)$$

where ω is angular frequency, g is acceleration due to gravity, ω_p is peak frequency, $\beta = 0.74$ and additionally

$$\alpha = 0.076 \left(\frac{U_{10}^2}{Fg}\right)^{0.22} \quad (3.29)$$

$$\omega_p = 22 \left(\frac{g^2}{U_{10}F}\right)^{\frac{1}{3}} \quad (3.30)$$

$$\sigma = \begin{cases} 0.07 & \omega \leq \omega_p \\ 0.09 & \omega > \omega_p \end{cases} \quad (3.31)$$

where U_{10} is the wind speed 10m above the surface, F is the fetch (the distance over which wind velocity remains constant) and σ determines peak sharpness either side of the peak frequency ω_p .

An extension to Pierson-Moskowitz spectrum, the JOint North Sea WAve Project (JONSWAP) [221], introduced the peak enhancement factor γ to better represent of waves in the North Sea. The JONSWAP spectrum is expressed:

$$S(\omega) = \frac{\alpha g^2}{\omega^5} \exp\left(-\beta \left(\frac{\omega_p}{\omega}\right)^4\right) \gamma^a \quad (3.32)$$

where a typical peak enhancement factor of $\gamma = 3.3$ is often used with an exponent of

$$a = \exp\left(-\frac{(\omega - \omega_p)^2}{2\sigma^2\omega_p^2}\right) \quad (3.33)$$

The peak enhancement factor is an important parameter within a JONSWAP spectra and effects the maximum spectral densities that are achieved. A JONSWAP spectra with $\gamma = 1$ (i.e no peak enhancement) is equivalent to the Pierson-Moskowitz spectrum. A comparison of the JONSWAP and Pierson-Moskowitz spectra is shown

in Figure 3.9.

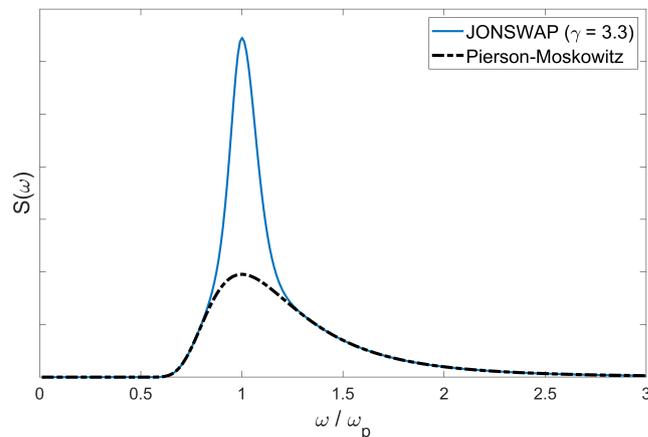


Figure 3.9: Comparison of the JONSWAP and Pierson-Moskowitz spectra. The introduction of the peak enhancement factor γ causes a sharper peak and higher maximum spectral density within the JONSWAP spectra.

By varying the parameters within the spectra, it is possible to define and generate a variety of ocean states. Where the variation of ω_p is intuitive, having an effect on the dominant frequencies present with the wave, for the variation of γ it can be helpful to visualise with the examples shown in Figure 3.10. In general, the higher the γ value, the narrower the peak within the spectrum and the higher the maximum value within the spectrum achieved. For the relative wave profiles, this produces higher peak wave amplitudes with a narrower band of frequency content within the wave.

3.2.3 Empirical wave loading prediction

Empirical methods of wave loading prediction offer a balance between predictive performance, computational resource requirements and ease of model validation; they are popular within many industrial applications [44, 45]. For the modelling of wave loads on slender members, which many offshore structures are comprised of, Morison's Equation has been the most widely used such method since its introduction in 1950 [32].

For a stationary, rigid, slender, cylinder of diameter D positioned within waves of

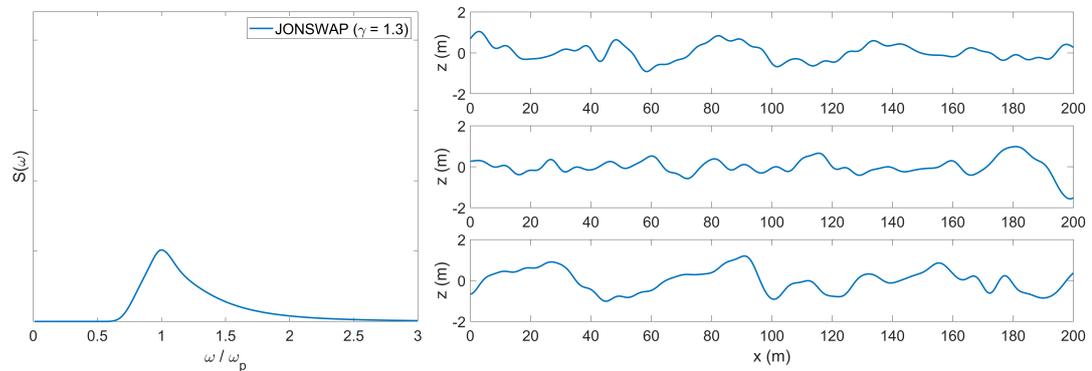
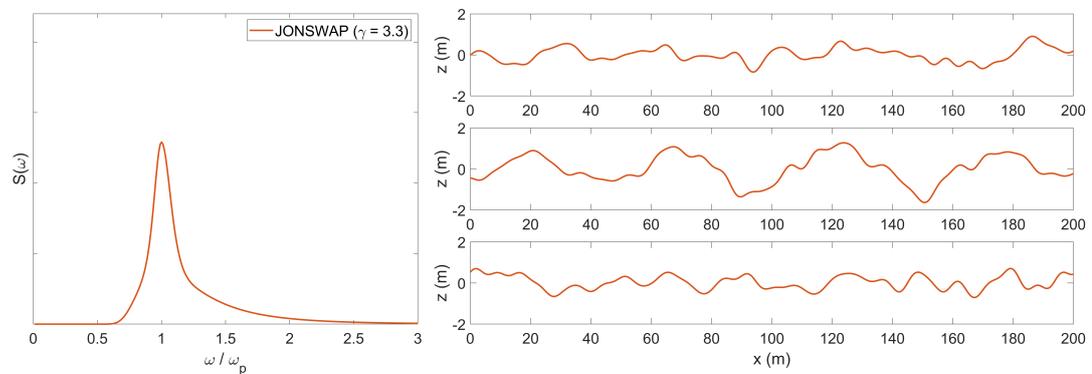
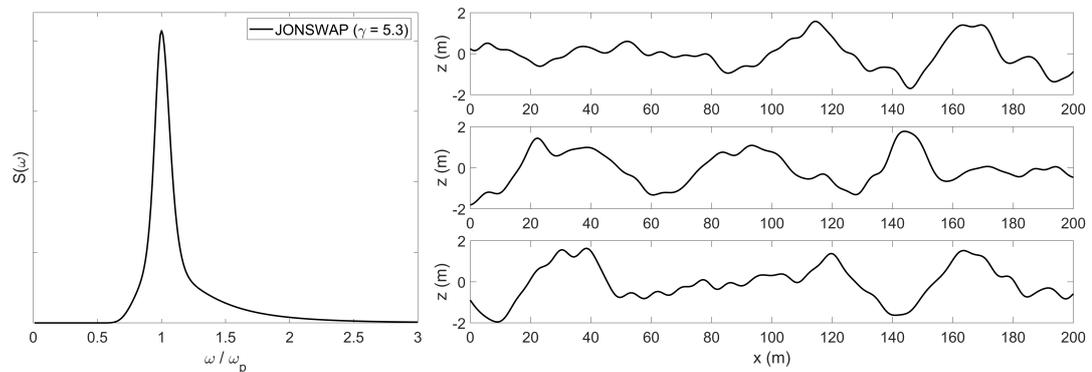
(a) Free surface realisations from a JONSWAP spectra with $\gamma = 1.3$.(b) Free surface realisations from a JONSWAP spectra with $\gamma = 3.3$.(c) Free surface realisations from a JONSWAP spectra with $\gamma = 5.3$.

Figure 3.10: A comparison of free surface realisations from JONSWAP spectra with varying γ values. A wind speed of $U_{10} = 12\text{ms}^{-1}$ and fetch of $F = 100\text{km}$ were kept constant for all waves.

velocity U and acceleration \dot{U} , the force per unit axial length F is expressed:

$$F = \underbrace{\frac{1}{2}\rho DC_d U|U|}_{C'_d} + \underbrace{\frac{1}{4}\pi\rho D^2 C_m \dot{U}}_{C'_m} \quad (3.34)$$

where ρ is the fluid density, C_d is the drag coefficient and C_m is the inertia coefficient. The dimension specific terms may be grouped to form two constants C'_d and C'_m relating to the drag and inertia forces of the wave [36, 222].

An important consideration when using empirical methods, as with any engineering model, is understanding the limitations and assumptions made within the construction of the model. To achieve their computational efficiency, empirical methods often rely on strong simplifying assumptions. For Morison's equation these include:

- The waves are not affected by the presence of the submerged members. For a cylindrical structure, the wavelength and water depth should far exceed the diameter [44].
- Flow should be unidirectional [46].
- The wave force may be separated into a velocity-dependant drag term and an acceleration-dependant inertia term, simplifying the wave-structure interaction [46].
- The considered waves are surface waves and unbroken [32].

Extensions to Morison's Equation

Morison's Equation is generally well regarded within the literature [46, 223, 224] with Sarpkaya stating 'it is unlikely that an entirely new equation will ever replace it' [225]. Research efforts focus mainly on the development of extensions and modifications to Morison's Equation, rather than a competing alternative.

The work of Chakrabarti [95], Sundar [96] and Anandkumar [97] modified Morison's equation to calculate forces on inclined cylinders. This is particularly useful in the case of offshore jacket structures, helping to maintain the unidirectional flow assumption relative to diagonal struts.

The Lighthill correction [226] is expressed

$$F = C_m^* \rho V \dot{U} + \frac{1}{2} \rho A U^2 C_d \quad (3.35)$$

where C_m^* is the ideal inertia coefficient, V is the volume of the body and A is the projected area. Lighthill's correction has the advantage of now requiring only a singular model parameter, C_D , however further studies have shown that performance declined in comparison to Morison's Equation outside inertia dominated flow regimes [224, 227].

Sarpkaya [98] introduced a parameter $\Delta = \frac{\pi^2}{K}(C_m^* - C_m)$ such that

$$F = -C_d |\cos(\omega t)| \cos(\omega t) + \frac{\pi^2}{K} C_m \sin(\omega t) - \frac{\Delta |\Delta|}{C_d} \sin(3\omega t) \quad (3.36)$$

where K is the Keulegan Carpenter number. Although the model was able to reduce the residuals of Morison's Equation for cases of sinusoidal flow, this would not be a valid assumption for an offshore structure in a general sea state.

Drawing from methods within the system identification community, Stansby [99] and Worden [228] proposed the Morison Duffing equation

$$\hat{F}_i = -\alpha_1 \ddot{F}_i - \alpha_2 \dot{F}_i - \alpha_3 F_i |F_i| + \beta_1 U_i |U_i| + \beta_2 \dot{U}_i \quad (3.37)$$

where the higher order wave force terms, \ddot{F}_i and \dot{F}_i , aimed to capture the history effect of vortex shedding. The model was able to improve wave force classification in to drag, inertia and history components, however the prediction errors remained inline with Morison's Equation.

3.3 Concluding remarks on theory

In this chapter, introductions to relevant theory from both a data-based and physics-based perspective were presented to the reader. Gaussian process models form the data-based component of later developed models and were therefore a major focus. Instead of a traditional mathematical introduction, for which many high quality examples exist within the literature [139, 215, 216], model interpretation through figures, particularly during their training stage was the main aim here.

From a physics-based perspective, linear wave theory, the representation of sea states and empirical wave loading prediction methods were introduced. These are all widely used aspects of physical theory used within offshore applications. It was important to introduce readers to the types of physical knowledge used in the models developed within the thesis. For readers already familiar with physical

methods used within offshore applications, this chapter serves to collate the authors understanding of a selection of key methods.

DATA AND EXPERIMENTATION

The physics-informed models presented throughout this thesis are primarily developed using two datasets: a subset from the Christchurch Bay Tower (CBT) project, a heavily instrumented column structure within a real offshore environment; and an experimental study of a monopile within a wave tank designed and carried out by the author. The experiment is designed to study the wave loading of monopiles in a range of representative sea states, steady flows and sinusoidal waves in a controlled lab environment. Its primary motivation comes from only having access to a narrow range of conditions from the CBT dataset. This chapter aims to outline each dataset and provide detail of experiment design.

4.1 The Christchurch Bay Tower

The Christchurch Bay Tower (CBT) was an offshore test facility, constructed specifically for the study of wave loading and effects of current and tides on cylindrical offshore structures. This provided a dataset capturing a real sea state environment with valuable measurements of the wave load, allowing for the validation and performance measurement of modelling approaches [229]. The project design and data collection was lead by British Maritime Technology (BMT) throughout the 1980s, with data collected in 1987 [230, 231].

The structure is comprised of a large central column, 2.8m in diameter, and a smaller column, 0.48m in diameter, each equipped with an array of sensors. Perforated ball Velocity Meters (PVMs), pressure transducers, force sleeves and wave buoys were

used to create a 41 channel dataset [230]. Although this is a historic dataset, the sensor network here is more densely populated than those that might be employed on offshore structures today due to the CBT being constructed specifically as a test facility. A schematic of the structure is shown in Figure 4.1.

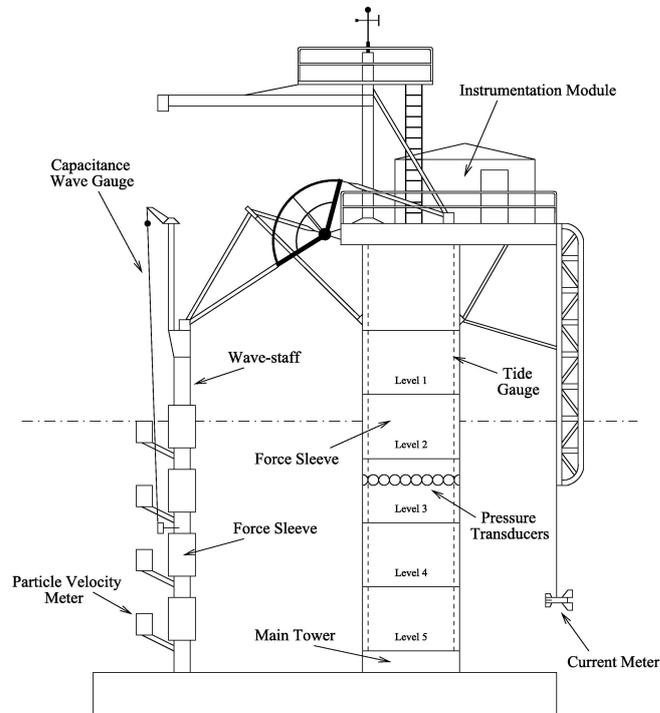


Figure 4.1: Schematic of Christchurch Bay Tower [232].

Whilst it would be desirable, the author did not have access to the full CBT dataset. Access to a single 17 minute run was kindly provided by Keith Worden from use in earlier work [232]. To the best of the authors knowledge, much of the remainder of the dataset has been lost over time. Access to only a narrow range of wave conditions was a major motivator for the development of an experiment, as discussed later in this chapter, in which a range of conditions could be generated and controlled.

Three subsets of 1000 data points, sampled at 13.25Hz, were selected from the 17 minute run for use as training, validation and test sets. A region of 3000 points was selected where the ratio of the x-velocity to y-velocity of the wave was a maximum to ensure that the flow was primarily unidirectional. This was then split in to three sequential subsets as seen in Figure 4.2. The study here focuses on data from the small column where the assumptions of Morison's equation around slender members are more likely to be valid.

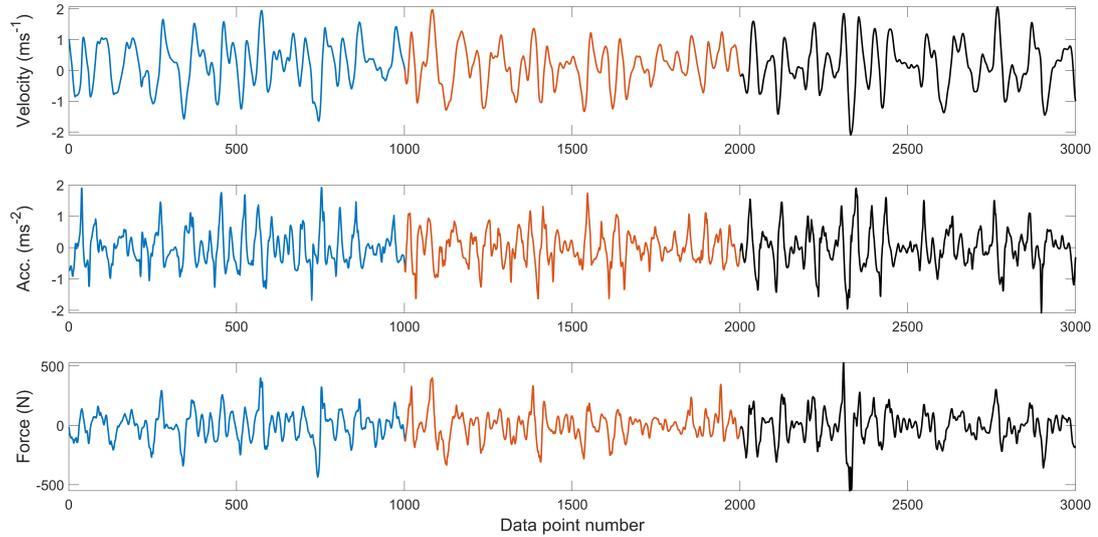


Figure 4.2: The inline water particle velocity, acceleration and wave load of the selected subset of the CBT dataset. The training, validation and test sets are shown in black, red and blue respectively.

The data used was taken from one of the high intensity measurement runs of the complete CBT dataset in which the Reynolds number (Re) $> 1 \times 10^5$ and the Keulegan–Carpenter number (Kc) was in the range $17 < Kc < 26$ [230]. For these flow conditions, Clauss [233] suggests a drag coefficient of $C_d = 0.6$. For the inertia coefficient, DNV-RP-C205 [234] also considers the effect of surface roughness, which for a heavily instrumented cylinder gives $C_m = 1.2$. The presence of the force sleeve, accelerometers and numerous pressure transducers along the cylinders length significantly increase the surface roughness.

In the literature, some of the earliest analysis is available from Bishop throughout the data collection phase of the project. He derived wave force coefficients for the tower both before [231] and after [229] the repair of concrete foundations. The initial gravity base of the structure experienced rocking due to being too small, leading increased erosion of surrounding seabed and eventually failure. This led to a pause of the project and an increase of base weight from 200 to 800 tonnes before measurements continued [229]. Later work of Burrows [235] and Najafian [230] studied the applicability of Morisons Equation across the range of conditions within the dataset, with both reporting good agreement. Burrows [235] reported that no simple extension to the equation, to improve performance in the presence of history effects and vortex shedding, could be found.

The CBT dataset is used in both Chapter 5 and Chapter 7 to develop PIML models for wave loading prediction. Chapter 5 incorporates physical knowledge within the inputs and mean function of a GP whilst Chapter 7 focusses on physics-informed kernel design.

4.2 Monopile Experiment

Although access to a section of the CBT dataset was very useful for the observation of wave loads acting on a real world offshore structure, the narrow window of conditions available to the author meant that the development of models able to operate across a range of wave states was challenging. An ability of a model to operate well in conditions it hasn't previously seen is a key driver of grey-box research. This therefore motivated the design and completion of an experiment, with the aim of providing access to measurements of the wave load acting on a structure across a wider range of conditions.

The wave loading of a monopile structure was performed in the wave tank at the Laboratory for Verification and Validation (LVV)¹. This provided an environment in which conditions could be more tightly controlled by the user, ensuring through experimental design that a wide variety of wave states could be measured. Conditions including sinusoidal waves, steady flow and representative sea state spectra were generated within the tank and data from the structure and tank recorded. The structure and wave tank were heavily instrumented with accelerometers, strain gauges, a force collar, wave gauges and a velocimeter, providing a 2048Hz, 27 channel dataset. This section aims to detail the methods and motivation of the experiment along with specific detail of wave load measurement and wave generation.

4.2.1 The structure

A monopile structure with a fixed base was used within the experiment, allowing for the study of wave loads on vertical cylinders, a popular application within literature [224, 236, 237]. Monopile structures including offshore wind turbines are a common industrial application, with the potential extension of findings to slender cylindrical members relevant for jacket structures. The structure, before it was placed within the wave tank, is seen in Figure 4.3.

¹More information on LVV facilities available at lvv.ac.uk



Figure 4.3: The monopile structure before it was placed within the wave tank. The fixed metal base and crane used to load variable masses to the top of the structure can be seen.

The structure was comprised of a 2990mm length of PVC pipe with an outer diameter of 90mm and wall thickness of 5.4mm . The relatively low stiffness material in combination with long pipe length helped to keep the natural frequency of the structure low, both to mimic the low natural frequencies of real offshore structures and to allow the study of excitation close to this frequency via waves within the tank. Keeping the pipe length long rather than trimming also aimed to align the ratio of the structure above and below the waterline with some real structures. When within the 1m water depth of the wave tank, this would result in approximately one third of the structures length being submerged; which is within the range of deeper water offshore wind turbines found in the literature [238–242].

In its default state, the first natural frequency of the structure was measured at 2.46Hz . This was significantly higher than the wave frequencies that could be generated within the wave tank and therefore studying the behaviour of the structure around resonance would be challenging. It was therefore desirable to lower the first natural frequency of the structure within the bounds of wave frequencies that could be generated. The adopted means of natural frequency reduction was

through the addition of mass to the top of the structure; this had the additional benefit of mimicking the mass of a rotor and nacelle at the top of a wind turbine.

To calculate the required additional mass, the monopile was first modelled as a cantilever beam following the approaches of Blevins [243] and Rao [244] from which the natural frequency is expressed:

$$\omega_n = (\beta_n L)^2 \sqrt{\frac{EI}{\rho AL^4}} \quad (4.1)$$

where β_n is a boundary condition specific constant, L is the length of the beam, E is Young's modulus, ρ is material density and A is cross sectional area. The cantilever approximation was found to be in good agreement with measured results, with the first four modes shown in Figure 4.4. This was to be expected due to the homogenous material, constant cross section and sturdy fixed clamp at the base of the structure.

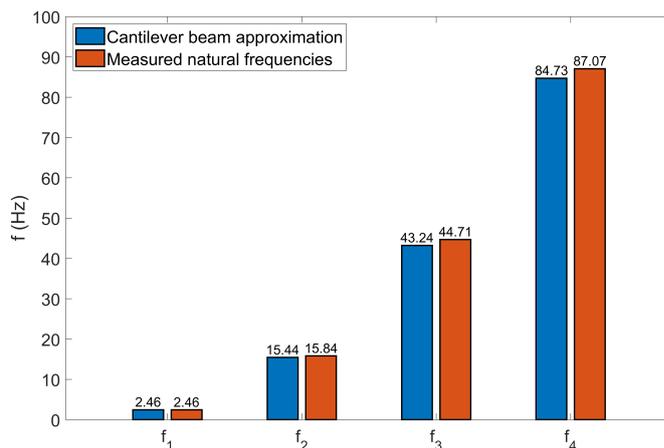


Figure 4.4: Comparison of the first four measured natural frequencies, in Hertz ($f_n = \frac{\omega_n}{2\pi}$), with the calculated result from a cantilever beam approximation.

To calculate the change in natural frequency due to the addition of mass, the effective mass m_{effn} of the beam was calculated as if all mass was placed at the beam tip.

$$m_{effn} = \frac{3EI}{L^3\omega_n^2} \quad (4.2)$$

This was then able to be summed with added mass and the new natural frequency ω_n^* calculated.

$$\omega_n^* = \sqrt{\frac{3EI}{L^3(m_{eff_n} + m_{add})}} \quad (4.3)$$

To validate the approach, a series of increasing masses were placed on the structure and natural frequencies measured. Figure 4.5 shows the FRFs for the structure with varying levels of added mass. The optimum operating range for the wave tank is 0.7-1.0Hz, and a desired first mode natural frequency of 0.8Hz was therefore chosen. Measured results were found to be in close agreement with calculated changes in natural frequency. An estimated 8.43kg of added mass was required to reduce the first natural frequency to 0.8Hz, as highlighted in Figure 4.6. The attached mass along with a safety tether in case of it falling from the structure can also be seen installed on the structure in Figure 4.10.

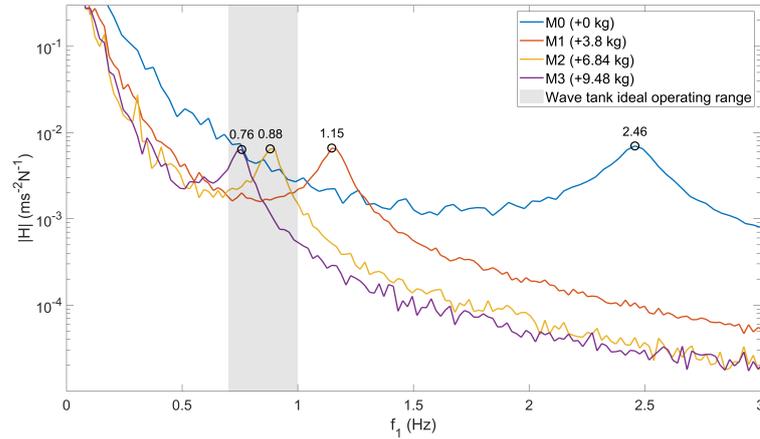


Figure 4.5: FRFs of the monopile with increasing amounts of added mass. The optimum operating range of the wave tank is highlighted in grey.

The main requirement of securing the structure in the tank was a solution that would last for the full duration of the experiment without any degradation or loosening. This would provide a consistent set of boundary conditions for all wave run measurements. Were the structure to separate from the tanks base, this would require draining of the tank to perform repairs and rerunning of measurements, therefore causing delays. For this reason, a fixed metal baseplate was used to securely clamp the structure, with high strength suction cups used to attach this to the glass floor of the wave tank. The fixed baseplate installed within the tank before it was filled can be seen in Figure 4.7.

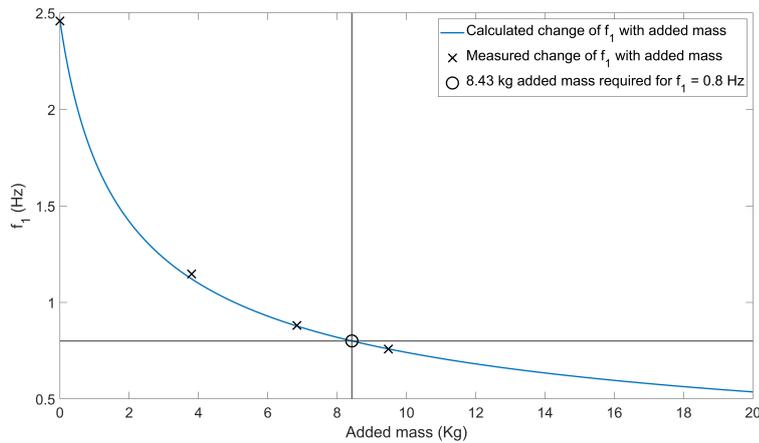


Figure 4.6: Measured and calculated changes of the first natural frequency of the monopile with increasing quantities of added mass. The desired natural frequency of 0.8Hz required an estimated 8.43kg of added mass.



Figure 4.7: Metal baseplate with suction cups used to secure the structure to the wave tank. The photo was taken before the tank was filled with water.

The accelerometers and strain gauges were both installed on the structure in pairs; two single-axis accelerometers measured a_x (inline acceleration) and a_y (lateral acceleration) at six locations along the structures length. Strain gauges fixed to opposite sides of the column were used to measure compressive and tensile strain, allowing the calculation of inline bending strain. Unlike the accelerometers,

the strain gauges were fully waterproofed, enabling some to be placed below the water-line. The locations of the accelerometers and strain gauges are shown in Figure 4.8.

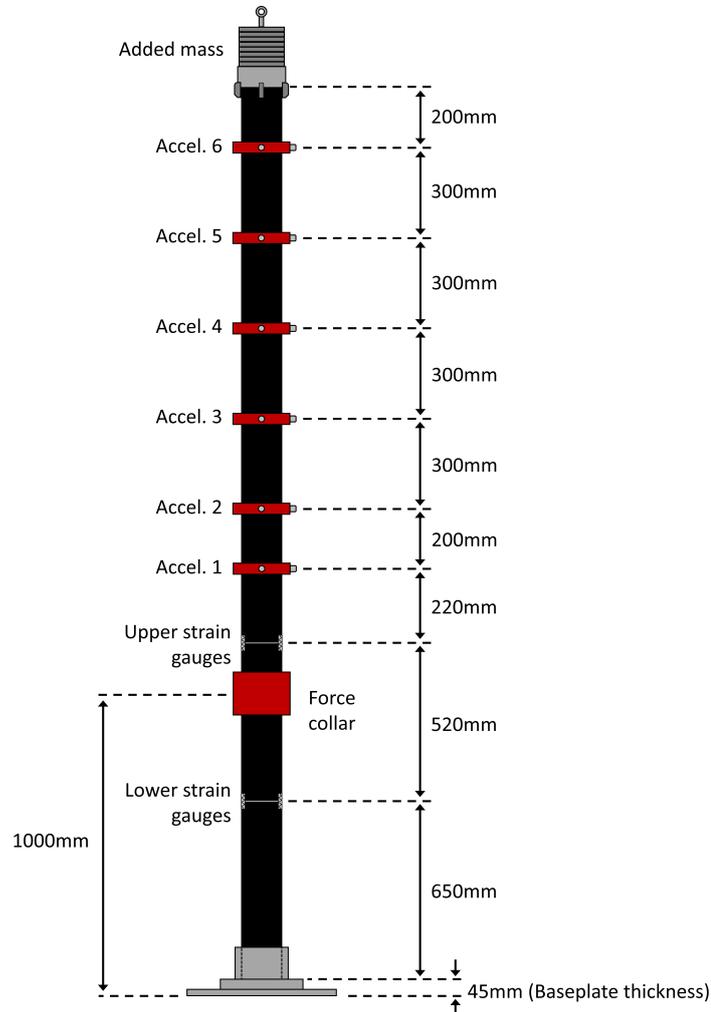


Figure 4.8: Locations of installed sensors on the monopile structure.

4.2.2 The wave tank

The wave tank had an operational volume $10.74m$ in length, $0.5m$ in width and $1m$ in depth. Although the tank had paddles capable of generating waves in either direction, one set of paddles was kept slack to absorb incoming waves and reduce reflection of waves back along the tank as much as possible. A series of

two wave gauges were used to capture the free surface of passing waves, whilst an Acoustic Doppler Velocimeter (ADV) was used to measure water particle velocity and acceleration. An ADV relies on the transmitting, reflection and receiving of high frequency sound waves on particulates suspended within the water. This therefore required the seeding of the wave tank. Dimensions of the wave tank and experimental setup are shown in Figure 4.9.

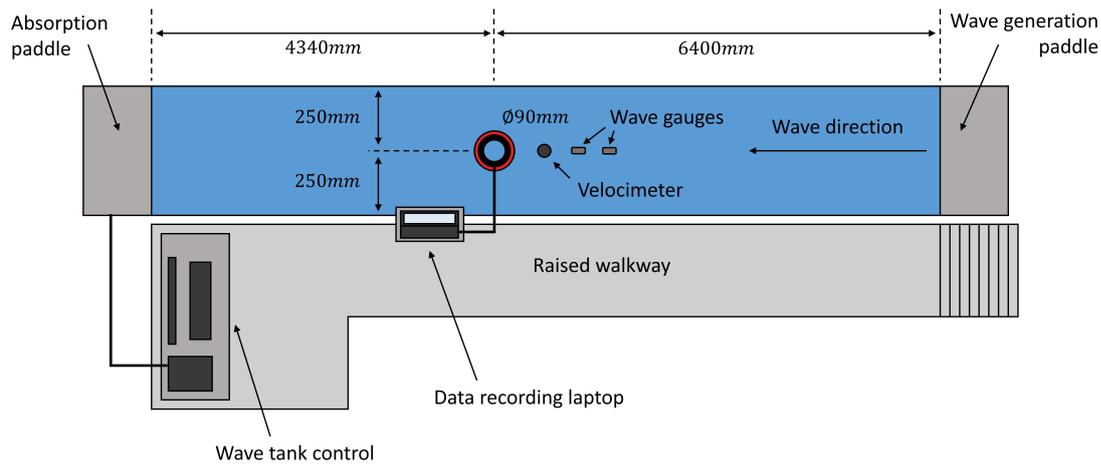


Figure 4.9: Dimensions of the wave tank and experimental setup.

The generation of waves was done using Edinburgh Designs Wave Synthesiser and Runtime [245]. This allowed the definition of sinusoidal waves via frequency and amplitude, along with irregular waves based on representative spectra. The wave tank was capable of generating waves within the frequency range $0.5 \leq \omega \leq 1.5 Hz$, up to a maximum amplitude of approximately 280mm (peak to trough). The structure was instrumented with accelerometers, strain gauges and a force collar. Along with the measurement equipment within the wave tank, this would provide access to incoming wave heights, flow conditions, wave load, structure response and strain. A photograph of the experimental setup is shown in Figure 4.10.

4.2.3 Measurement of the wave load

The wave load was the most challenging variable to measure and required the design and development of bespoke equipment. The difficulty and expense in measuring wave loads is a driving factor for the field of wave loading prediction. There exists a range of methods by which to measure wave loads within the literature: Bachynski [246] used a leaf spring and dynamometer at the base of the monopile to allow for the measurement of moments, Esandi [247] used a fixed cross beam connected to

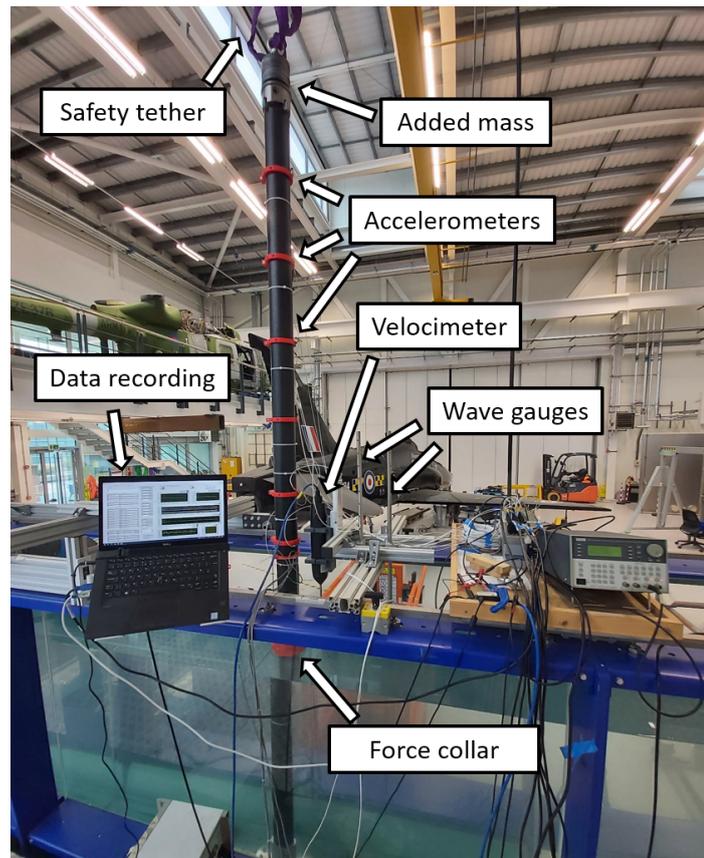


Figure 4.10: Monopile structure and measurement equipment setup within the wave tank.

the monopile via load cells and the Christchurch Bay Tower (CBT) project [230] used a force collar. Here, a force collar was used for two main reasons: it would not impede the motion of the monopile, as would be the case when connecting to a fixed cross beam; and the force collar could be positioned at any desired vertical height, providing a local measurement of wave load.

The force collar used for this work was designed by Robin Mills and manufactured by Robin Mills and Michael Dutchman. The author wishes to express thanks here for both their efforts. The force collar consisted of two concentric rings; an inner ring fixed to the structure and an outer ring fixed to cover plates. The rings were connected via hinged strain gauges, such that when a force is applied to the cover plates, a series of stains are produced when the load is transferred to the inner ring. With three strain gauges, there was sufficient information to resolve for forces in x and y directions. A diagram of the force collar is shown in Figure 4.11.

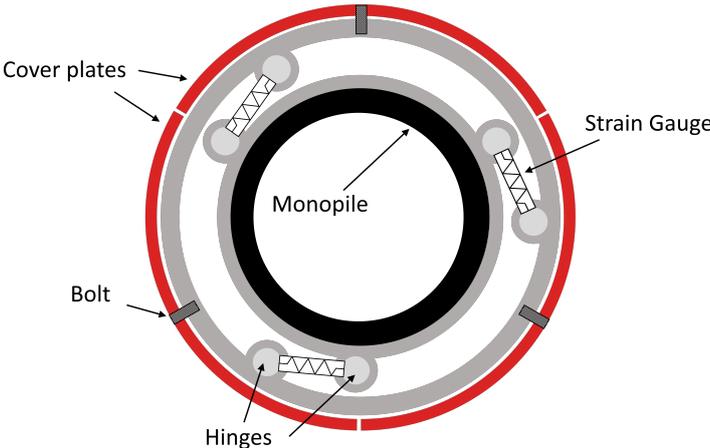


Figure 4.11: Force collar diagram with key components highlighted.

The force collar was aligned such that the x direction corresponded to an inline force, with its centre positioned at the waterline. The area close to the free surface is typically where the maximum wave forces are experienced and is therefore the region of interest of this study. The force collar was first installed on the structure without cover plates and a calibration check was performed by Robin Mills in dry conditions under known forcing. The wave tank was then filled and cover plates installed. Images of the installed force collar are shown in Figure 4.12.

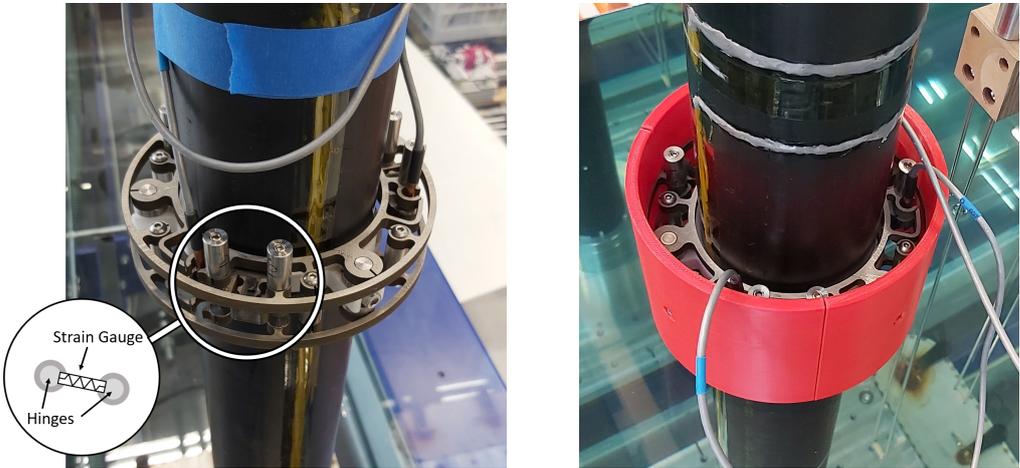


Figure 4.12: Images of the force collar installed on the monopile without the cover plates (left) and with the cover plates (right). The position of the hinged strain gauge connecting the inner and outer rings is highlighted.

4.2.4 Measurement of flow conditions

The instantaneous velocity and acceleration of the flow was measured using an Acoustic Doppler Velocimeter (ADV), a popular method used within civil and geographical applications. An ADV relies on the transmitting, reflection and receiving of high frequency sound waves on particulates suspended within the water. This therefore required the seeding of the wave tank. A diagram of an ADV in operation is shown in Figure 4.13.

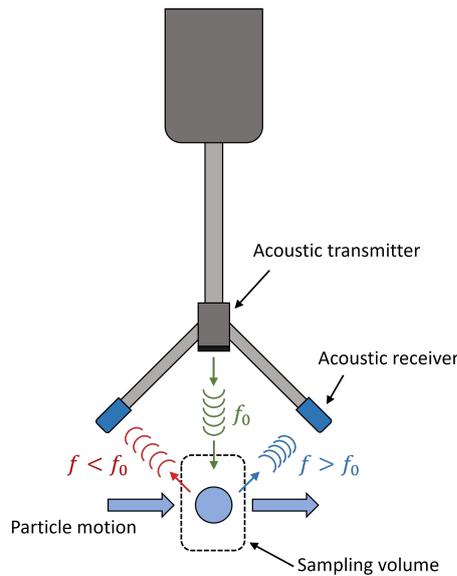


Figure 4.13: An Acoustic Doppler Velocimeter (ADV) in operation. The relative change between the source frequency f_0 and received frequency f at is highlighted.

An ADV works via calculation of Doppler shift, the change in frequency observed due to relative motion of an observer and source. Here, the motion of a passing particle (seeding) produces a frequency shift in the transmitted and observed frequencies. The observed frequency (f) is expressed

$$f = \left(1 + \frac{\Delta v}{c}\right) f_0 \quad (4.4)$$

where Δv is the relative motion between the receiver and source, c is the wave propagation speed and f_0 is the source frequency. Assuming a stationary source, the particle velocity can then be derived. The presence of multiple receivers allows for the calculation of velocity components in multiple directions.

The data from an ADC can be subject to a high level of noise and the presence of spikes. This is due to the reliance on a consistent stream of particles passing within the control volume. The numerical differentiation required to obtain acceleration measurements was highly susceptible to both spikes and noise and the use of filters was therefore a necessity. A 102 point median filter was first used to remove spikes within the data and a 4Hz stopband Butterworth filter was then used to reduce the impact of noise. In both cases the filter parameters required significant tuning: the median filter was increased until all spikes were successfully removed and the Butterworth stopband frequency reduced until successful numerical differentiation could be achieved. A single set of filter parameters was kept constant for all test cases.

4.3 Test cases

The primary motivation for this experiment was to provide access to a range of controlled conditions and wave states with which to test the development of wave loading prediction models. The range of test cases for in the experiment are detailed here with the development of models using this data shown in Chapter 6.

4.3.1 Sinusoidal waves

Sinusoidal waves are the most simplistic wave type that may be generated within the wave tank and present a test case suitable for the study of linear wave theory and initial testing of new methods. A series of 153 sinusoidal waves were generated through variation of frequency and amplitude to create a 9×17 grid. The amplitude was varied from 0.01m to 0.05m in 0.005m increments, whilst the frequency was varied from 0.5Hz to 1.3Hz in 0.05Hz increments. The wave tank was programmed to run batches of wave states, with each individual sine wave case run for a period of two minutes. A 30 second settling period was used to separate wave states; this allowed any leftover waves from previous states to decay before the next run started. A time series segment of selected variables for a sinusoidal wave is shown in Figure 4.14.

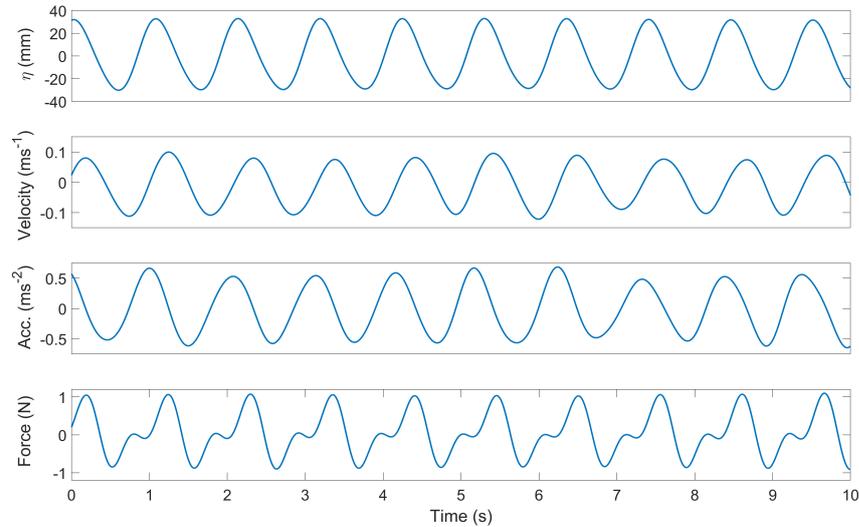


Figure 4.14: Measured surface elevation, inline water particle velocity and acceleration, and inline wave force over a 10 second time period for a sinusoidal wave. Note that unlike the other variables, the wave load here is not sinusoidal; this is due to the interaction between the wave and the structure.

4.3.2 Representative sea states

The generation of sea state spectra is of particular interest when studying offshore structures. The prediction of wave loads and structural response within these conditions is more challenging and provides a more representative measure of model performance. Although within a lab environment, which does not capture the true variability of a real sea, sea spectra are a useful means through which to first develop model structures.

By varying the parameters within the JONSWAP [221] spectra, it is possible to define and generate a variety of ocean states. For this experiment, a series of JONSWAP waves were generated by varying both γ and ω_p , thereby creating a matrix of test conditions. Based on the capabilities of the wave tank, ω_p was varied from $0.7Hz$ to $1.1Hz$ in $0.1Hz$ increments and γ was varied from 1.3 to 5.3 in 1.0 increments. This was deemed to give a suitably fine grid of 25 test conditions. The JONSWAP spectra variation is shown in Figure 4.15.

Unlike sine waves, which have a repeating profile, the surface of a JONSWAP wave

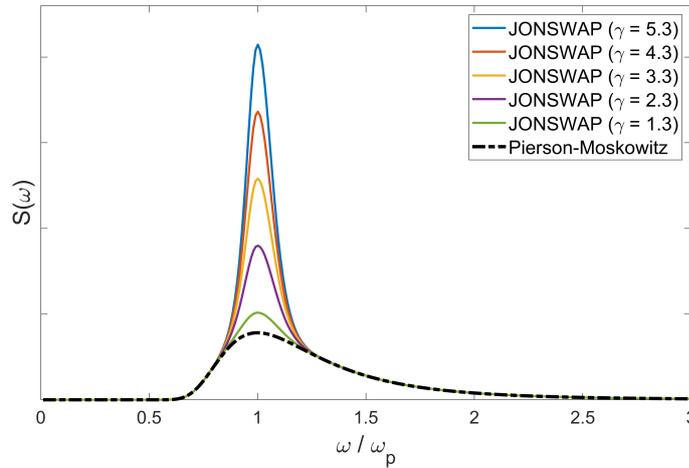


Figure 4.15: JONSWAP spectra highlighting the effect of γ variation for the generated test cases. All JONSWAP wave states had a runtime of 10 minutes.

varies over time, even when the spectra used to generate it remains constant. This behaviour can be seen in the time series in Figure 4.16. For this reason the wave state duration was increased to 10 minutes for each individual wave state. In the case that a longer continuous length of wave state measurement might be required, some additional longer runs of 1 hour each were also recorded.

4.3.3 Steady flow

Structures subject to steady flow can be of interest when studying phenomena such as tidal flows, currents and the potential identification of vortex shedding. To record the force due to the flow and the response of the structure for a range of flow speeds a ramp up in flowrate was performed. The wave tank had capability to adjust steady flow as a % of maximum, with a ramp from 0% to 80% performed over a period of 40 minutes. Beyond 80% flowrate, the flow became very turbulent and measurements of flow conditions became challenging. Plots of the inline water particle velocity and wave force are shown in Figure 4.17. A key observation here is that for a linearly increasing flow velocity, a square relationship can be seen in the inline force. This is to be expected due to the expression for drag load caused by a moving fluid ($F_D = \frac{1}{2}\rho U^2 C_D A$) [222].

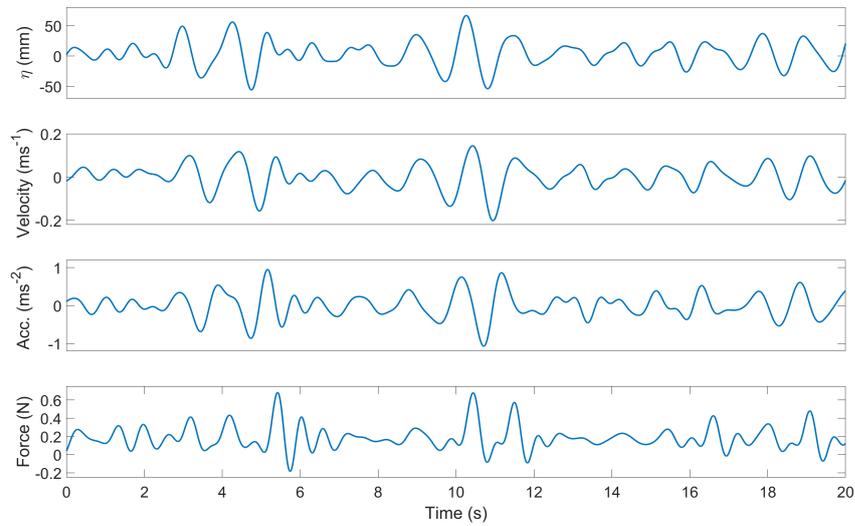


Figure 4.16: Measured surface elevation, inline water particle velocity and acceleration, and inline wave force over a 20 second time period for a JONSWAP wave. Note the more complex wave profile than the previously shown sinusoidal wave in Figure 4.14.

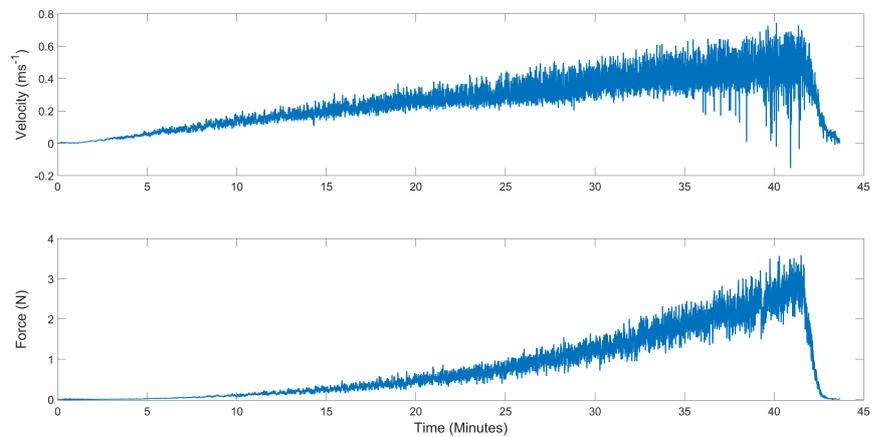


Figure 4.17: Inline water particle velocity and wave force during the steady flow ramp. Wave tank flowrate was linearly increased from 0% to 80% over a 40 minute duration.

4.4 Data within the thesis

This chapter presented the main datasets used to develop models within this thesis, one of which was a monopile experiment designed and completed by the author. Although both useful for the study of offshore structures, each dataset will serve its own distinct role. A section of data the Christchurch Bay Tower project provided access to measurements from a real world structure in an offshore environment. The experimental dataset provided a wider range of conditions that are more tightly controlled by the user, allowing for the evaluation of model performance in conditions they hadn't previously seen.

The experiment itself is a novel contribution here, with measurements of wave load on structures very rare. This is primarily due to the difficulty and expense of creating the bespoke equipment required. An effort was made to measure more variables and generate more conditions than would be needed for this thesis, with the hope that this dataset will be useful for the future work the author and others. The PIML models in Chapter 6 are developed using the experimental dataset, showcasing how both Morison's Equation and linear wave theory may be incorporated within a GP-NARX framework.

A FIRST GREY-BOX MODEL FOR OFFSHORE STRUCTURES

This chapter presents the first development, to the author’s knowledge, of a grey-box model for the prediction of wave loads on a real offshore structure. It is developed using the Christchurch Bay Tower dataset detailed in Chapter 4. The widely used Morison’s Equation, an empirical wave loading solution, is incorporated within an autoregressive form of Gaussian process Regression (GP-NARX). Important details of implementation for GP-NARX models are addressed including lag selection, uncertainty propagation and cost function selection, all of which have an impact on final model performance. A specific study focuses on the ability of different grey-box model architectures to extrapolate in circumstances of low data coverage. This is a key expected benefit of incorporating physical knowledge within a machine learner. The majority of the work in this chapter has been published in [100].

5.1 Development of a grey-box model

The work in this chapter was the first attempt of the author at grey-box model development and was completed at a time when the field was much smaller than today. At the time, and arguably still, there was not a defined methodology for the creation of grey-box models. This therefore required thought on how one might sensibly define and create an effective model. To provide structure to the task, the answering of three questions was proposed:

1. What type of physics is to be included in the model?
2. What type of machine learner should be used?
3. How should the physics be integrated within the machine learner?

The ordering of the questions is also important here, with the most suitable way to construct the model highly dependent on the chosen types of physics and machine learner. In the view of the author, it is therefore sensible to first define the physics and data-based approaches before deciding how they should be combined.

This section details the modelling methodology proposed for the prediction of wave loading and is structured to answer the questions above. Morison's Equation forms the basis for the white-box model construction (with Bayesian linear regression for parameter estimation), whilst a GP-NARX model is used for the black-box. Methods of combining the approaches to form grey-box models are then presented.

5.1.1 White-box

Within research communities, the study of Computational Fluid Dynamics (CFD) has dominated the quantification of wave loading forces [30, 31]. Within industrial applications however, the high computational resource requirements of CFD and difficulty of model validation for structures in complex environments, has led to a preference, in some industries at least, for more simplistic empirical methods. A common example of one such method is Morison's Equation [32], which offers an empirical solution for wave loading with minimal computational resources. It has been used in a range of applications including wind turbine design [44] and characterising dynamic behaviour of offshore spar platforms [45].

Further details on the application of Morison's Equation and the assumptions behind its construction are given in Chapter 3. Here, the earlier defined form is adopted:

$$F = \underbrace{\frac{1}{2}\rho DC_d U|U|}_{C'_d} + \underbrace{\frac{1}{4}\pi\rho D^2 C_m \dot{U}}_{C'_m} \quad (5.1)$$

where ρ is the fluid density, D is the cylinder diameter, C_d is the drag coefficient and C_m is the inertia coefficient. The dimension specific terms may be grouped to form two constants C'_d and C'_m relating to the drag and inertia forces of the wave.

This leads to the simplified version of Morison's equation used within this work:

$$F = C'_d U|U| + C'_m \dot{U} \quad (5.2)$$

The construction of the white-box model presented will rely on the estimation of the grouped parameters within Morison's equation. The simplicity of this approach will aid in minimising the complexity of the final combined grey-box models presented in later sections, although a non-dimensional form can be readily used.

Parameter estimation and model prediction with Morison's equation is achieved via Bayesian Linear Regression (BLR). Approaching the regression in a Bayesian manner provides distributions over the parameter estimates and confidence intervals for the predictions, which can then be compared with the Gaussian process models used later.

Implementation of Morison's Equation as a BLR begins with gathering model inputs U, \dot{U} and grouped constants C'_d, C'_m ; they may then be expressed as a single design matrix X along with model parameters β :

$$X = [U|U|, \dot{U}] \quad (5.3)$$

$$\beta = \begin{bmatrix} C'_d \\ C'_m \end{bmatrix} \quad (5.4)$$

Morison's equation can then be expressed in matrix form:

$$F = X\beta + \varepsilon \quad \text{where} \quad \varepsilon \sim \mathcal{N}(0, \sigma_n^2 \mathbb{I}) \quad (5.5)$$

A Bayesian linear regression can be set up for the model.

$$p(F|X, \beta, \sigma_n^2) = \mathcal{N}(X\beta, \sigma_n^2) \quad (5.6)$$

In order to retrieve the desired posteriors over the parameters for Morison's Equation β and noise variance σ_n^2 , it is necessary to place priors over the parameters. Here a Normal-Inverse-Gamma semiconjugate prior is used:

$$p(\beta) = \mathcal{N}(m_\beta, \sigma_\beta^2) \quad (5.7)$$

$$p(\sigma_n^2) = \mathcal{IG}(a, b) \quad (5.8)$$

A Gaussian prior over the parameters for Morison's Equation allows for a positive or negative mean m_β with a given variance σ_β^2 . The selection of appropriate C'_d and C'_m priors can be made using the dimension specific terms for the structure and standards relating to flow specific drag C_d and inertia C_m coefficients such as

DNV-RP-C205 [234]. An Inverse Gamma (\mathcal{IG}) prior can encode belief about the noise variance through hyperparameters a and b , whilst restricting to only positive values.

The full joint posterior $p(\beta, \sigma_n^2|F, X)$ is unavailable in closed form and it is therefore necessary to calculate the conditional posterior for each parameter: $p(\beta|F, X, \sigma_n^2)$, $p(\sigma_n^2|F, X, \beta)$. The parameter posterior distributions were recovered via Gibbs sampling with 10,000 draws. This provided a computationally efficient method for accurate estimation of the conjugate conditional distributions [248]. Once the parameter distributions were retrieved, further sampling was used to recover the posterior distribution of the test set force prediction from Equation 5.6.

5.1.2 Black-box

Gaussian Process Regression (GPR), utilised here, is a non-parametric, flexible, Bayesian machine learning technique [139]. The return of confidence intervals with predictions, minimal requirement for prior knowledge and modelling capabilities under the presence of noise have lead to the popularity of GPR within a wide range of usage applications. These span from standard regression tasks [249], to image processing [250], to more engineering relevant examples [1, 214, 251]. A dynamic-variant of a GP regression model is employed here, namely a GP-NARX [41, 232, 252].

A Nonlinear AutoRegressive model with eXogenous inputs (NARX) is a function of previous signal values and additional (exogenous) inputs, in which both are fed through some nonlinear function $f(x)$.

$$y_t = f([u_t, u_{t-1}, \dots, u_{t-l_u}, y_{t-1}, y_{t-2}, \dots, y_{t-l_y}]) + \varepsilon \quad (5.9)$$

The previous signal values, y_{t-i} and exogenous inputs, u_{t-j} are considered up to l_y and l_u lagged time steps respectively. For the wave force estimation in this work, the exogenous inputs, u , considered are the velocity, U and acceleration, \dot{U} of the wave particles. There are several reasons why one might wish to use an autoregressive model for the prediction of wave loads: they are well suited to the capture of periodic behaviours, for example, dominant frequencies within waves; they have been shown to be effective in capturing higher order terms present within fluid phenomena such as vortex shedding [232]; and they are capable of capturing delays between measured signals. The perforated ball velocimeters on the CBT project, as with most flow measurement equipment, were offset from the main tower, thereby creating a delay between the measured velocities and wave load.

The nonlinear function, $f(x)$, in a NARX model is commonly fixed to be a polynomial [253, 254], but in the case of a GP-NARX, a Gaussian process (GP) is used. With a GP one avoids needing to fix the functional form explicitly, instead, the selection of a mean and covariance function defines a family of feasible functions that may explain the data. An overview of GP theory along with the author's interpretation of how a GP model operates is given in Chapter 3.

5.1.3 Grey-box

A grey-box model combines physics and data-based approaches with the aim of extracting benefits from each of the model types: structure, insight and extrapolative performance from the white-box component and flexibility and ability to model unknown phenomena from the black-box component. There are two potential architectures presented here, both of which combine the earlier discussed Morison's equation with GP and GP-NARX models.

Residual modelling

Perhaps the simplest approach to grey-box modelling is to sum the predictions of a white and black-box model. If the white-box takes a fixed form, this summation is equivalent to using the black-box to model the residual error between the white-box and any collected data. In the FE modelling community the practice of using a machine learner (often a GP) to model the residuals from an FE model is often referred to as 'bias correction', acknowledging that there is likely to be some error in the complex FE representation of the structure. It is possible to apply the approach to more simplistic base models such as Morison's Equation, with the modelling of polynomial regression residuals using a GP being explored as early as 1975 [255].

It is known that Morison's Equation simplifies the behaviour of wave loading, not accounting for effects such as vortex shedding or other complex behaviours [256] and will typically have residual errors in the region of 20% [257]. Through the use of a black-box component, these residual errors may be modelled and the result added to the white-box, producing a model of the form:

$$y_t = \underbrace{F_{mor}}_{White-box} + \underbrace{f([U, \dot{U}]) + \varepsilon}_{Black-box} \quad (5.10)$$

Here, the parameters in Morison's equation will be established via Bayesian linear regression. The GP or GP-NARX model is set up identically to those discussed within the black-box section, except that the target is now the residual error of Morison's Equation rather than the measured wave force itself. The intention is to capture the missing physics excluded by the simplifications present within Morison's Equation.

The magnitude of the black-box term can be interpreted as being the extent to which the data confirms the prediction of Morison's Equation. In regions of high uncertainty, outside the observed training data, the GP will revert to its prior of zero with the overall model therefore outputting Morison's Equation. An equivalent view of this is the usage of a white-box mean function within the GP or GP-NARX [139].

Input augmentation

An alternative means by which physics and data-based approaches may be combined is via the inclusion of the white-box model output as an additional input to the black-box. Models involving the manipulation of the non-parametric components of black-boxes using physics have been termed 'type B' [258] grey-box models within the nonlinear system identification community, whilst the transformation of model input data through physical insight has been referred to as 'semi-physical' modelling [60]. This approach was used in [116] to model and optimise hydroelectric power generation. Input augmentation has been shown to offer performance increases over white and black-box approaches in the context of a nonlinear cascaded tanks system, particularly in the case of an extended physical model [118].

Here, the result of Morison's Equation is used along with the originally included water particle velocity U and acceleration \dot{U} as the input to the GP or GP-NARX. The model is of the form:

$$y_t = f(\underbrace{[F_{mor}, U, \dot{U}]}_{\text{White-box}}) + \varepsilon \quad (5.11)$$

Black-box

The output of the white-box model is carried forward to provide an input for the GP that is strongly linked to the physics of the problem. An advantage of this approach over residual modelling is the maintaining of the signal to noise ratio. This is particularly important in cases of high fidelity physics-based models, where the residual of the model will be small in comparison to the model noise.

5.2 GP-NARX implementation

There are several important distinctions between the implementation of GP-NARX models, compared with a classic static GP. Many of these are introduced due to the ability of a GP-NARX to utilise previous predictions as future model inputs. They are important considerations within the modelling process, to ensure the potential benefits of a GP-NARX may be maximised. Key decisions made for the implementation of GP-NARX models on the CBT dataset are detailed here.

5.2.1 One step ahead and model predicted output

There are two types of prediction that may be calculated from any autoregressive model form: One Step Ahead (OSA) and Model Predicted Output (MPO)¹. For OSA, previously measured values of the output signal are used as lagged inputs to the model, whilst MPO requires the feedback of the model prediction itself. An OSA prediction for a GP-NARX has the form:

$$y_t = f([u_t, u_{t-1}, \dots, u_{t-l_u}, y_{t-1}, y_{t-2}, \dots, y_{t-l_y}]) + \varepsilon \quad (5.12)$$

where $f(X)$ is a GP, $[u_t, u_{t-1}, \dots, u_{t-l_u}]$ are previously measured exogenous inputs and $[y_{t-1}, y_{t-2}, \dots, y_{t-l_y}]$ are *previous measurements* of the target variable. An MPO prediction for a GP-NARX has the form:

$$y_t = f([u_t, u_{t-1}, \dots, u_{t-l_u}, \hat{y}_{t-1}, \hat{y}_{t-2}, \dots, \hat{y}_{t-l_y}]) + \varepsilon \quad (5.13)$$

where $[\hat{y}_{t-1}, \hat{y}_{t-2}, \dots, \hat{y}_{t-l_y}]$ are now *previous predictions* of the target variable. This one change can have a large impact on the properties of a model, how best to select hyperparameters and prediction errors. The MPO performance will generally be worse than that of OSA due to the compounding of model errors, however, it is a much more representative measure of how well the model has captured the true dynamics of the process and therefore a more rigorous test.

The practical use of an OSA prediction occurs most naturally in a control setting, where continual measurements of the target of interest are available. In a Structural Health Monitoring (SHM) context, the assumption is that continual measurement of the wave force itself will not be available, meaning that an OSA prediction will not generally be useful here. This work, therefore, focuses on the MPO task.

¹In some communities these are referred to as the *prediction* and *simulation* tasks for OSA and MPO respectively.

The focus on an MPO necessitates careful attention to how the hyperparameters in the GP covariance function are optimised. In a standard static implementation of a GP, the hyperparameters, θ of the covariance function, which control things like the smoothness of predictions, are set by optimising a negative log marginal likelihood:

$$\hat{\theta} = \arg \max_{\theta} \{-\log p(y|X, \theta)\} \quad (5.14)$$

where \mathbf{y} are the set of targets in the training set, with corresponding inputs X . In this work a squared exponential kernel with Automatic Relevance Determination (ARD) is used. This allows an independent length scale for each input parameter and offers increased model flexibility when operating with multiple types of input parameter.

$$k(x_i, x_j) = \sigma_f^2 \exp\left(-\frac{1}{2}(x_i - x_j)^T \Lambda^{-1}(x_i - x_j)\right) \quad (5.15)$$

where σ_f^2 is the signal variance and Λ is the matrix of length scales such that $\text{diag}(\Lambda) = [l_1^2, l_2^2, \dots, l_D^2]$ for a D dimensional input. These parameters are typically optimised over the Negative Log Marginal Likelihood (NLML) of model predictions on a training set:

$$-\log p(y|X, \theta) = \frac{1}{2}y^T(K + \sigma^2 I)^{-1}y + -\frac{1}{2}\log|K + \sigma^2 I| + \frac{n}{2}\log(2\pi) \quad (5.16)$$

This optimisation doesn't reflect the dynamic nature of the GP-NARX and strongly favours the performance of OSA predictions if used. The cost function should always be aligned with the desired performance criteria of the model, in this case the MPO. Here, therefore, a more appropriate choice of cost function is the Negative Log Predictive Likelihood of the Model Predicted Output (MPO NLPL):

$$\hat{\theta} = \arg \max_{\theta} \{-\log p(y|\mathbb{E}(y^*), \mathbb{V}(y^*), \theta)\} \quad (5.17)$$

The NLPL of the MPO is calculated as a joint Gaussian likelihood of each measured data point y_t coming from the corresponding predictive distribution $y_t^* \sim \mathcal{N}(\mathbb{E}(y_t^*), \mathbb{V}(y_t^*))$ of the GP-NARX output. Importantly, it provides a measure of how well the GP-NARX is able to perform when relying on previous predictions that have been fed back for use as model inputs. The MPO NLPL may be derived

as:

$$\begin{aligned}
-\log p(y|\mathbb{E}(y^*), \mathbb{V}(y^*), \theta) &= \frac{1}{2}(y - \mathbb{E}(y^*))^T \mathbb{V}(y^*)^{-1}(y - \mathbb{E}(y^*)) \\
&+ \frac{1}{2} \sum_{t=1}^n \log(\mathbb{V}(y_t^*)) + \frac{n}{2} \log(2\pi)
\end{aligned} \tag{5.18}$$

For the implementation of a MPO NLPL cost function, the author suggests using an independent validation set, used only for the assessment of MPO performance. The process is defined in Algorithm 1.

Algorithm 1: MPO NLPL cost function for GP-NARX optimisation.

- 1 Calculate training set covariance matrix $K(X, X)$ for hyperparameters θ
 - 2 Initialise GP-NARX for validation set from $U_{t:t-l_u}$, $\dot{U}_{t:t-l_u}$ and $y_{t-1:t-l_y}$
 - 3 **for** $t = 1 : T$ **do**
 - 4 | Calculate: $p(y_t^*|U_{t:t-l_u}, \dot{U}_{t:t-l_u}, \mathbb{E}(y_{t-1:t-l_y}^*), \theta) = \mathcal{N}(\mathbb{E}(y_t^*), \mathbb{V}(y_t^*))$
 - 5 **end**
 - 6 $NLPL = - \sum_{t=1}^T \log \mathcal{N}(y_t|\mathbb{E}(y_t^*), \mathbb{V}(y_t^*))$
-

where $U_{t:t-l_u}$ and $\dot{U}_{t:t-l_u}$ refer to the lagged exogenous input vectors of velocity and acceleration, $y_{t-1:t-l_y}$ refers to the lagged vector of measured wave force and $y_{t-1:t-l_y}^*$ the lagged vector of predicted wave force.

5.2.2 Uncertainty propagation in a GP-NARX

One of the benefits of using a Bayesian regression approach, such as GPR, is access to the full posterior distribution and therefore availability of confidence intervals on any prediction made. This causes an issue, however, in a NARX MPO setting, as model predictions get fed back and used as model inputs at the next step. Previous uses of GP-NARX models have generally avoided the tricky issue of uncertainty propagation in an MPO setting.

Without alteration, the confidence intervals of the GP-NARX fail to account for the full uncertainty within the prediction. The MPO requires the feedback of the model prediction, for use as subsequent lagged output. This is typically taken as a point estimate of the expected value of the GP-NARX prediction, failing to acknowledge that the output of the model is in fact a distribution. This does not

account for the potential variation in the feedback of model outputs which would have a cumulative effect over time. This causes an uncertain input problem in the GP, not handled by the standard GP framework.

Uncertainty propagation within the GP-NARX can be achieved via the use of Monte Carlo sampling [232, 259]. Instead of feeding back the model output mean, a sample from the output distribution \hat{y}_t is used. This may be repeated for N samples to form a series of N potential realisations for the model output $y_{t*}^{(n)}$ from which more realistic posterior distributions may be estimated. The procedure for generating the Monte Carlo sampled Model Predicted Output (MC MPO) is summarised in the block diagram within Figure 5.1.

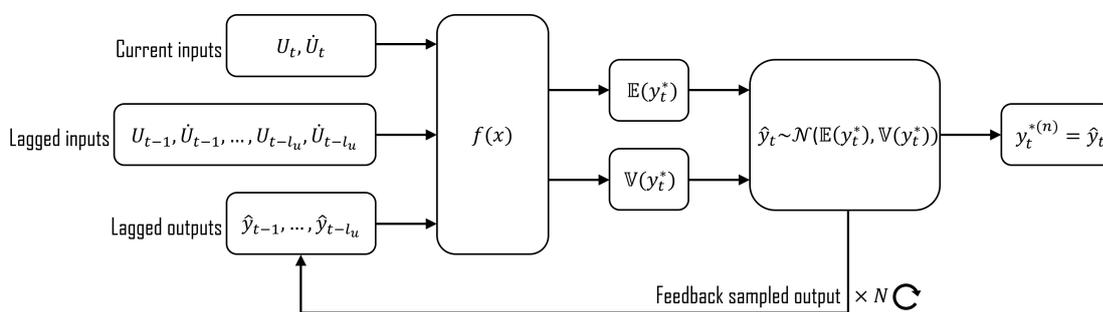


Figure 5.1: Block diagram of Monte Carlo uncertainty propagation within the GP-NARX.

5.2.3 Hyperparameter optimisation

Quantum Behaved Particle Swarm Optimisation (QPSO) [260] was used as a global, gradient-free method for the determination of hyperparameters in GP covariance functions, although it would be possible to use any other appropriate optimisation scheme. To ensure stable convergence, optimisation runs were repeated and the hyperparameters cross checked. The swarm size and cost function stability tolerance were adjusted accordingly until stability was achieved over 12 repeated optimisation runs. The required parameter settings for stable convergence of the GP and GP-NARX models are shown in Table 5.1.

In order to achieve stable convergence of hyperparameters, the GP-NARX required both a higher swarm size and a tighter cost function stability tolerance than the static GP. There were two major reasons for this: the increased number of hyperparameters introduced via additional length scales for lagged inputs and the increased complexity of the cost function. The additional hyperparameters

Table 5.1: QPSO parameter settings used for GP and GP-NARX optimisation.

Model type	Swarm size (n)	Cost function stability tolerance (t)
GP	200	1×10^{-3}
GP-NARX	1000	1×10^{-5}

increased the dimensionality of the search space whilst the propagation present within the calculation of the MPO NLPL cost function led to a high sensitivity to changes in hyperparameters.

5.2.4 GP-NARX lag selection

The selection of lag terms within a GP-NARX model has a significant effect on the structure and performance of the model. The number of lags included for the previous signal values and exogenous inputs can be optimised by considering l_u and l_y as hyperparameters. The optimal model may then be chosen via calculation of an appropriate model selection criterion. Two key challenges of lag selection addressed here are the high computational cost and evaluation of the most suitable model.

Computational cost

The primary difficulty faced within GP-NARX lag selection is the computational cost of the search. Even when considering only the maximum lags, rather than the full combinatorial problem, a significant number of models are required to be constructed and evaluated. Due to the long training time of the GP-NARX, particularly when using the MPO NLPL cost function, the computational cost of covering even moderate search spaces becomes an issue. A proposed solution is to perform the search on a computationally inexpensive AutoRegressive model with exogenous inputs (ARX) mimicking the structure of Morison's Equation and carry the lag selections forward for use in the full GP-NARX model. Although still autoregressive in nature and trained using the same datasets as the GP-NARX, it should be noted that the ARX is a linear model and will thus capture a reduced range of behaviours when compared with the GP-NARX. This will likely introduce slight deviation from the optimal lag selections; however, due to

Morison's Equation performing in line with best expectations, the result is expected to provide a sensible lag selection. This pragmatic decision was taken here to make a compromise between computation time and potential improvements in model performance.

An ARX model is considered for the lag selection search of the form:

$$y_t = \sum_{i=0}^{l_u} \alpha_i u_{t-i} + \sum_{i=1}^{l_y} \beta_i y_{t-i} + \varepsilon \quad (5.19)$$

where similarly to the GP-NARX, the previous signal values, y are the wave force and the exogenous inputs considered are the velocity, U and acceleration, \dot{U} of the wave particles.

$$U^T = [U|U|_t, \dot{U}_t, U|U|_{t-1}, \dot{U}_{t-1}, \dots, U|U|_{t-l_u}, \dot{U}_{t-l_u}] \quad (5.20)$$

$$\alpha = [C'_{d0}, C'_{m0}, C'_{d1}, C'_{m1}, \dots, C'_{dl_u}, C'_{ml_u}] \quad (5.21)$$

$$y^T = [y_{t-1}, y_{t-2}, \dots, y_{t-l_y}] \quad (5.22)$$

$$\beta = [C'_{y1}, C'_{y2}, \dots, C'_{yl_y}] \quad (5.23)$$

Model evaluation

Broadly, model selection criteria have two roles: to reward models that are of a 'good fit' to target data and to penalise model complexity. The aim is to have the most simplistic model that achieves good performance. The Akaike Information Criterion (AIC) [261] is a commonly used metric for the selection of models. In its least squares case, for a model with k parameters and n data points, AIC can be calculated as:

$$AIC = n \log \left(\sum_{i=1}^n \frac{(y_i - f(x_i))^2}{n} \right) + 2k \quad (5.24)$$

The first term, containing the sum of square residuals, has the role of minimising model error whilst the second term penalises models with a high number of parameters. The more favourable the model, the lower the AIC.

For instances of a low sample size relative to the number model parameters ($n/k < 40$) [262], it has been shown [263] that a second order variant (AIC_c) should be used. For a large sample size with respect to number of model parameters, the difference between AIC and AIC_c is negligible, however AIC_c provides better

model order choices at lower ratios [263]. AIC_c includes a bias correction term that favours more simplistic models and may be expressed in its least squares case as:

$$AIC_c = n \log \left(\sum_{i=1}^n \frac{(y_i - f(x_i))^2}{n} \right) + 2k + \frac{2k(k+1)}{n-k-1} \quad (5.25)$$

Due to two exogenous inputs being used within the model, the number of model parameters will grow very quickly within a search space of maximum lags considered. For example, by 10 output lags and 10 input lags, the number of model parameters will have already reached 30. This will exceed the condition ($n/k < 40$) for the sample size used of $n = 1000$. It is therefore more appropriate to use AIC_c for lag selection within the ARX model.

For model selection, it can often be preferable to consider multiple metrics and cross check results. This can be done as either a sense-check or for additional validation in the decision. An alternative to AIC is the Bayesian Information Criterion (BIC) [264]. It is similar in structure to AIC and may be expressed in its least squares case as:

$$BIC = n \log \left(\sum_{i=1}^n \frac{(y_i - f(x_i))^2}{n} \right) + k \ln(n) \quad (5.26)$$

The differing second term of the BIC takes in to account the sample size when penalising model complexity. At higher sample sizes, more complex models are penalised more heavily. For sample sizes of 8 and above, BIC will favour more simplistic models than AIC. Determining whether AIC or BIC is used depends heavily on model assumptions, approximations and inference and it is general good practise to consider both in model selection [265].

Viewed alone, model selection metrics have no physical meaning. The absolute values can differ greatly depending on the nature of datasets used and numbers of model parameters. It is instead more useful to consider the metric difference Δ_i . This is the difference between the metric score for a given model and the best performing model. Δ_i may be calculated as:

$$\Delta AIC_{ci} = AIC_{ci} - AIC_{cmin} \quad (5.27)$$

$$\Delta BIC_i = BIC_i - BIC_{min} \quad (5.28)$$

The best model will have a Δ_i of zero whilst all other models will be some positive value. The higher the Δ_i , the worse the model. Suggested levels of model support for given values of Δ_i are shown in *Table 5.2*.

Table 5.2: Perceived levels of model evidence for delta values. [262]

Δ_i range	Level of empirical support of model i
$0 < \Delta_i < 2$	"Substantial"
$4 < \Delta_i < 7$	"Considerably less"
$\Delta_i > 10$	"Essentially none"

A search space of up to 20 lagged time steps was considered for both the outputs and exogenous inputs. A heatmap of ΔAIC_c and ΔBIC values for both the OSA and MPO prediction of the ARX model is shown in Figure 5.2. The blue areas represent lower values of Δ_i and superior models whilst yellow indicates higher Δ_i values and therefore worse models.

The optimum lags were found to be $l_u = 1$ and $l_y = 3$ for both the OSA and MPO predictions of ΔBIC and the OSA prediction of ΔAIC_c . The MPO prediction of ΔAIC_c narrowly suggested $l_u = 1$ and $l_y = 4$ as optimal with $l_u = 1$ and $l_y = 3$ having a ΔAIC_c of 1.67. This still provided 'substantial' [262] evidence in favour of the lags $l_u = 1$ and $l_y = 3$ which were therefore selected for the model.

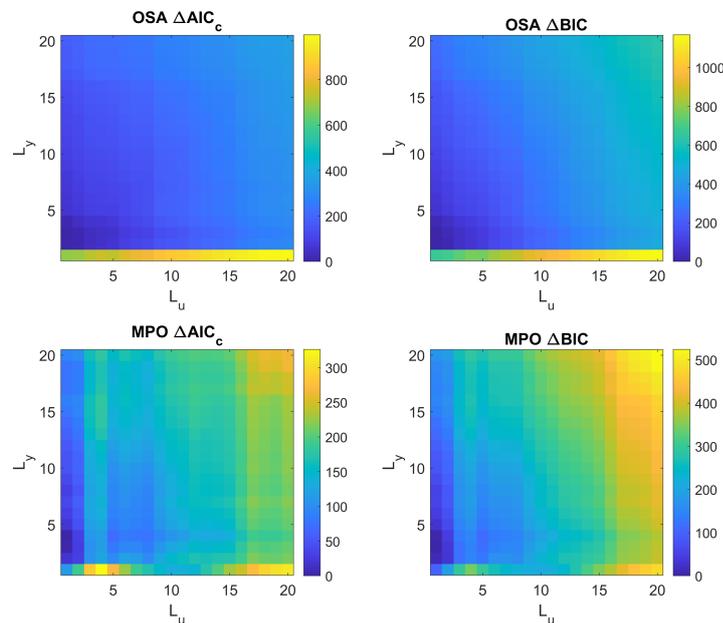


Figure 5.2: Heatmaps of lag selection metrics for the OSA and MPO predictions of the ARX model.

The results suggest that, uniformly, there is little benefit in considering more than one lag for the particle velocity and acceleration terms. Instead, the richer dynamics are expressed through the autoregressive terms for the force.

5.3 Results on the Christchurch Bay Tower

This section presents implementations of the proposed model architectures on a dataset collected from the Christchurch Bay Tower (CBT) [230] detailed in Chapter 4. This provided a test of model performance on a structure within a real sea environment. In all cases, the training, validation and test sets remain the same, and unless otherwise specified, all results correspond to performance on the unseen test set. All GP-NARX model variants included three lags of the wave load and one lag of the exogenous input terms (water particle velocity and acceleration).

5.3.1 Uncertainty propagation

The consideration of the uncertainty present within the feedback of GP-NARX outputs via implementation of MC MPO was found to contribute a significant amount to the overall uncertainty within predictions. Figure 5.3 compares GP-NARX MPO and MC MPO predictions, where one can see a significant difference between the widths of confidence intervals. The average increase in $\pm 3\sigma$ confidence interval width was 75.0%.

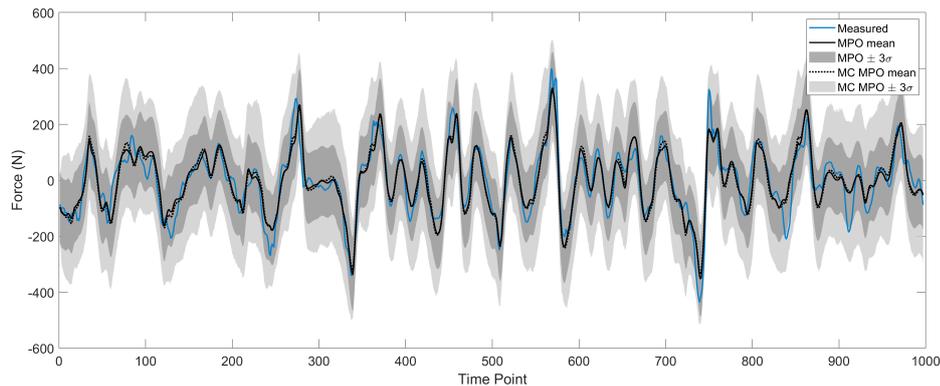


Figure 5.3: Comparison of black-box GP-NARX MPO and MC MPO predictions on the test set.

Although a perceived increase of model uncertainty may be argued a disadvantage, the underlying uncertainty present within the modelling processes has remained the same. What has instead changed is the proportion of the uncertainty that has been captured. A model should aim to be as realistic as possible about uncertainty within its predictions in order to prevent circumstances of ‘confidently wrong’ predictions. At around time points 840 and 910, instances can be seen of poorer model performance when a sudden downwards spike in the measured data occurs. The data under study here are selected from a time when the wave state was close to unidirectional. It is likely that these spikes occur at times when the wave direction changed. This results in both a data-based component that is unable to characterise the unseen conditions well and a physics-based component where the underlying assumptions are likely less valid. In the case of the MPO prediction, the confidence bounds are not wide enough to accommodate the measured data. For the MC MPO prediction however, the confidence interval width can be seen to increase significantly in these areas and is able to account for the true behaviour. A reduction in cases of ‘confidently wrong’ predictions is a major advantage of using MC MPO predictions.

In contrast to the confidence intervals, very little variation between the mean outputs can be seen. For the majority of the test set prediction, the MPO (solid line) and MC MPO (dotted line) are difficult to distinguish. For the purposes of uncertainty feedback, 10,000 Monte Carlo samples were used which were enough to ensure the convergence of NMSE within a 0.001% tolerance and average variance within a 0.01N tolerance. With the wave load exceeding a 200N magnitude in many areas, this level of precision was deemed acceptable. If the number of samples were to be increased, the mean output of MC MPO should continue to converge towards the MPO. This similarity in response means that the advantages in uncertainty representation are achieved without deterioration in performance of the mean output.

From here onwards, the MC MPO will be considered the primary prediction type of interest.

5.3.2 Model predictive performance

The performance of each model was assessed using the response prediction on an unseen test-set that was not used in the estimation of any parameters or hyperparameters. For the purposes of model comparison, two measures are used: the Normalised Mean Square Error (NMSE) to assess the performance of each models expected output and the Mean Standardised Log Loss (MSLL) to provide

a probabilistic measure. The NMSE is expressed:

$$NMSE = \frac{100}{n\sigma_y^2}(y_\star - y^*)^T(y_\star - y^*) \quad (5.29)$$

where n is the sample size, σ_y^2 is the signal variance, y_\star is the measured test signal and y^* is the model prediction. An NMSE of zero implies perfect prediction whilst an NMSE of 100 would be equivalent to predicting the mean of all observations.

To construct the MSL, one must first consider the negative log predictive likelihood of the model, $-\log p(y_\star|X_\star, X, y)$, where y_\star is the measured test signal, X_\star is the set of test inputs, X is the set of training inputs and y is the training target. Taking the negative here returns a loss rather than a utility which may be standardised by subtraction of the loss calculated when predictions equal the mean and variance of the training set. This returns a Standardised Log Loss (SLL):

$$SLL = -\log p(y_\star|X_\star, X, y) + \log p(y_\star; \mathbb{E}(y), \mathbb{V}(y)) \quad (5.30)$$

The SLL, and therefore the MSL, will be equal to zero for the baseline case of predicting with the training set mean and variance and increasingly negative for improved model predictions.

A comparison of metrics for the models and their various prediction types is shown in Table 5.3. Comparisons of the full test set posterior between the model with the lowest NMSE, the grey-box residual modelling GP-NARX, and other model types are shown in Figures 5.4-5.6. All results presented within this section relate to models constructed using the full training and validation sets.

A stand out observation from the results in Table 5.3 is the significant performance gap, in terms of both NMSE and MSL, between the GP-NARX OSA and all other prediction types. This is to be expected due to the nature of OSA predictions and inclusion of lagged measured outputs within the model inputs. The prediction of a single time step ahead is of very limited use in SHM applications, particularly in the case of high sample rates, thus the good performance is of little benefit.

The NMSE of the white-box linear regression was found to be 19.528% which is in line with the expected 20% [257] residuals of Morison's Equation. Although this is around 3-6% higher than the NMSE of other models, the result is achieved with reduced modelling complexity and computational burden. Considering the simplified version of Morison's Equation used had only the two model parameters C'_d and C'_m to model the relatively complex wave load, even moderate levels of model performance are commendable.

Table 5.3: Performance comparison of model types.

Model	Model type	NMSE (%)	MSLL
Morison's Equation	White-box	19.528	-0.813
Static GP	Black-box	16.433	-0.939
	Residual modelling	16.751	-0.914
	Input augmentation	15.627	-0.951
GP-NARX OSA	Black-box	2.702	-1.444
	Residual modelling	5.212	-1.012
	Input augmentation	2.947	-1.429
GP-NARX MPO	Black-box	14.773	-0.968
	Residual modelling	13.862	-0.872
	Input augmentation	14.072	-0.994
GP-NARX MC MPO	Black-box	14.643	-0.788
	Residual modelling	13.840	-0.835
	Input augmentation	14.088	-0.791

All models including a machine learning component were able to offer significant reductions in NMSE over the white-box model. The success of the grey and black-box models was to be expected due to the failure of Morison's Equation to account for complex behaviours present within wave loading such as vortex shedding [256]. The inclusion of the black-box component, whether this be a GP or GP-NARX, increased model flexibility, allowing the representation of such behaviours. This indicates that previously missed underlying structure within the wave force was then able to be captured.

For all grey-box models except the residual modelling static GP, modest improvements in NMSE over the equivalent black-box approach were observed, implying that the inclusion of physics through the white-box component was able to aid model performance. For the residual modelling static GP, the inability of the black-box component to model the dynamics in the white-box residuals led to a reduction in performance. The fact that the residual modelling GP-NARX was the best performing grey-box model suggests the presence of some structure within the residuals only able to be captured by the GP-NARX and not by the static GP. In terms of missed phenomena from the model, this would be a process captured well by an autoregressive model exhibiting features such as periodicity over a small time scale.

In general, a larger increase in performance can be seen between the grey-box and white-box models, than between the grey-box and black-box models. The primary reason for this is when the full training and validation sets are used, the model is deemed to be mostly interpolating (See Section 3.1.3). Black-box and grey-box models are expected to achieve a similar performance in interpolation, with a physics-based component being most useful to assist with extrapolation. The specific type of components used to construct a grey-box model will also affect the relative differences in performance. Here, Morison's Equation – an approximate wave loading solution, is combined with a GP-NARX – a relatively powerful black-box architecture. The computational balance is heavily weighted in favour of the GP-NARX and it is to be expected that the performance of the grey-box would be more similar to the black-box than the white box in this case.

The MSSL of Morison's Equation was -0.813 , which being in line with the MSSL of the residual modelling GP-NARX MC MPO of -0.835 , indicated a strong model performance. Although a similar MSSL results in a similar prediction likelihood, it does not describe other aspects of model performance. Morison's Equation achieved the result with the highest NMSE of all models and the residual modelling GP-NARX the lowest. The similar MSSL was achieved through the lower prediction variance of Morison's Equation which can be seen from the narrower confidence interval width within Figure 5.4. The trade-off between variance and NMSE would generally be preferred in favour of NMSE with wider confidence intervals better able to contain the measured result.

The earlier discussed effect of GP-NARX uncertainty propagation can be seen within the increase in MSSL between the MPO and MC MPO of the GP-NARX. The increased prediction variance caused by the feedback of output distribution samples reduced the likelihood of the prediction significantly. However, the proper treatment of uncertainty is important in preventing overestimation of prediction likelihood and overconfidence within predictions far from the observation.

5.3.3 Model performance during extrapolation

A major drawback of black-box models is poor performance in conditions outside those experienced within the training and validation sets. In an engineering context, this necessitates the collection of data across all possible operating conditions of the structure or system of interest. Such a demand may be extremely costly, or impossible in some cases. Improvements in extrapolative capabilities would relax the dataset coverage of conditions required for effective machine learning implementation, thereby reducing the associated monitoring efforts. This is where

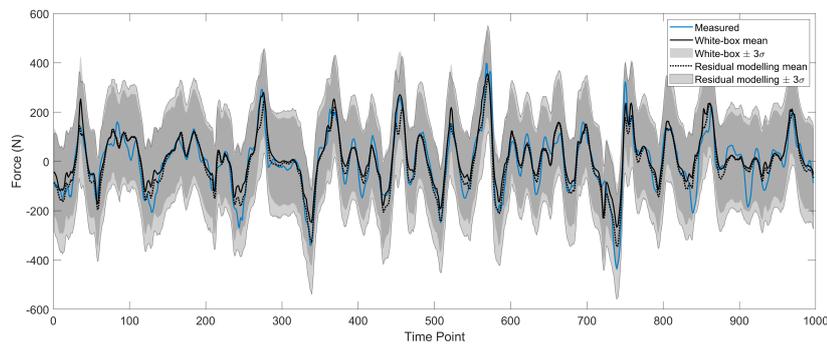


Figure 5.4: Test set prediction comparison between the white-box linear regression (solid line) and the MC MPO of the grey-box residual modelling GP-NARX (dotted line).

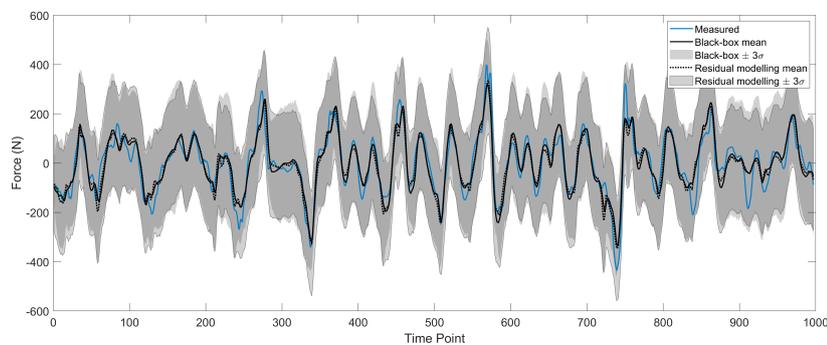


Figure 5.5: Test set prediction comparison between the MC MPO of the black-box (solid line) and grey-box residual modelling GP-NARX (dotted line).

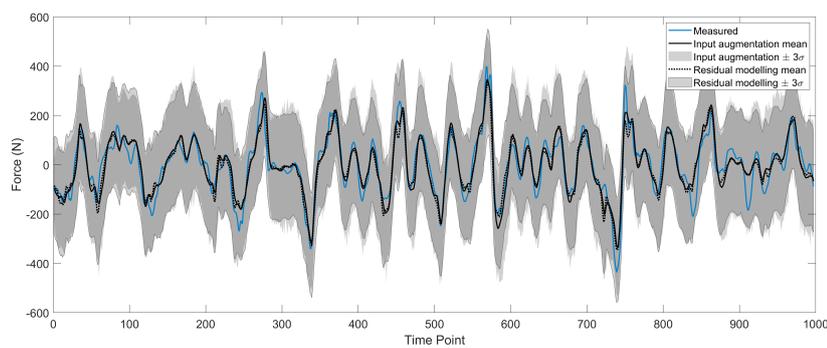


Figure 5.6: Test set prediction comparison between the MC MPO of the grey-box input augmentation (solid line) and residual modelling (dotted line) GP-NARX models.

one would expect a grey-box model to be of particular use.

To consider the extent to which a model is extrapolating, it is useful to consider how the training, validation and test sets overlap within the input space, which for the case of wave loading is formed of the velocity U and acceleration \dot{U} of the wave particles, and their lags. In the idealised case, all conditions within the test set will lie within the area/volume² covered by the training and validation sets; this guarantees that the model is interpolating at all times (assuming the training and validation set are representative of the behaviour of the system). Obtaining such a dataset is challenging in many contexts, particularly for offshore environments where conditions are highly variable and measured data for extreme events is rare. Model performance in extrapolation is, therefore, very important in such cases.

For fewer than four dimensions, one may visualise the input space for the training, validation and test sets, and consequently their overlap. Here, considering the input space in terms of the only the velocity and acceleration of the wave particles and not their lags, Figure 5.7 plots the boundaries for the three datasets used in the previous section. The boundaries of the dataset are determined to be the maximum projections from the origin of the input space in all directions encompassed by the measured data. In areas where the testing set overlaps the training and validation sets, the model will be interpolating, else it will be extrapolating. It is worth noting that even the ‘complete’ datasets used here represent only a small proportion of conditions able to be experienced by the structure (as the data were chosen where flow conditions were close to uni-directional, as discussed at the start of Section 3). The 2D boundaries of wave velocity and acceleration were found to be comparable to the boundaries of the first 2 input space PCA components which accounted for 95% of total variance and deemed representative of the input space as a whole.

For this two-dimensional case, one may use a coverage measure to assess the extent to which one is extrapolating. Here the coverage is calculated as the proportion of area within the test set boundary that lies within the boundaries of both the training and validation sets:

$$\text{Coverage} = 100 \left(\frac{A'}{A} \right) \quad (5.31)$$

where A' is the area within the test set boundary also enclosed by both the training and validation set boundaries and A is the total area within the test set boundary. Although the density of points within boundaries can vary at each coverage level, a relative measure of extrapolation is achieved, allowing for an investigation in to

²In most cases this is likely to be an n -dimensional volume, as the input space will generally have more than two dimensions

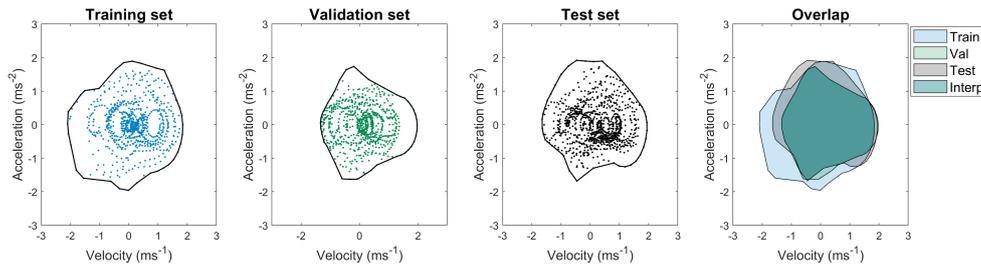


Figure 5.7: 2D visualisations of training, validation and test set boundaries and their overlapping region of interpolation. The black lines represent the boundaries enclosing the coloured points from each dataset. The blue, green and grey shaded regions represent the areas covered by each dataset whilst the shaded teal region represents the area of the test set considered interpolation.

extrapolative performance of the models. Note that outliers will have a considerable impact on the calculated coverage level, leading to an over estimation of actual coverage and should be removed from datasets where appropriate. In order to assess model performance at a range of coverage levels, the sizes of the training and validation sets were adjusted to achieve desired levels of coverage. As the number of points used increases, the boundary covered by the datasets will grow and hence cover a larger area of the test set. The rate of growth of the boundary will be highest at low quantities of training and validation points where each additional data point will have a higher chance of widening the boundary. Plots of overlapping boundaries for a range of training and validation set sizes and their respective coverage levels are shown in Figure 5.8. Results of model NMSE from 0% to 80% coverage levels in 5% coverage intervals is shown in Figure 5.9.

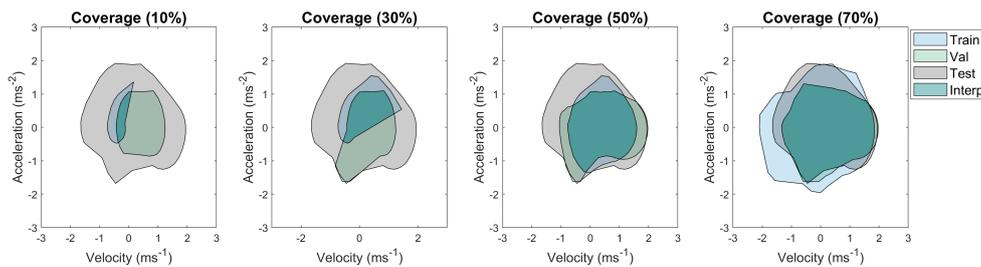


Figure 5.8: Plots of 10%, 30%, 50% and 70% test set coverage at increasing quantities of training and validation points. The blue, green and grey shaded regions represent the areas covered by each dataset whilst the shaded teal region represents the area of the test set considered interpolation.

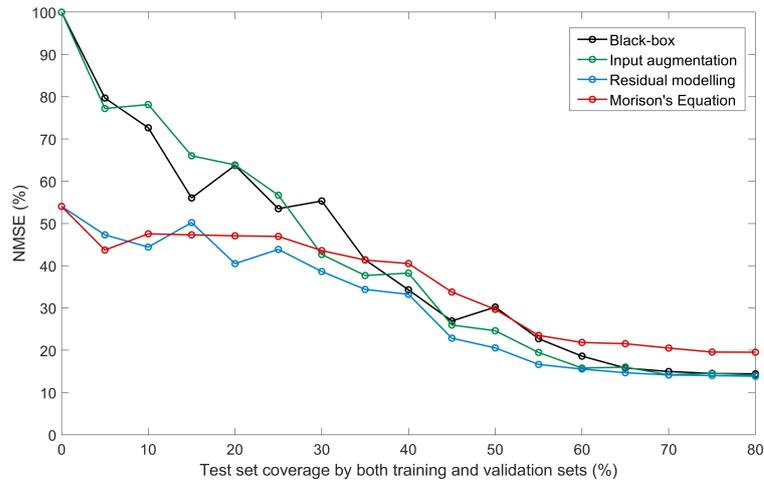


Figure 5.9: NMSE vs coverage for the MC MPO of black-box and grey-box GP-NARX models and Morison’s Equation with all models constructed using the restricted sizes of training and validation sets and prior C_d and C_m coefficients suggested by Clauss [233] and DNV-RP-C205 [234].

The largest difference in model performance is visible at the lowest levels of coverage, with the NMSE of the black-box and input augmentation models increasing steeply as the coverage approaches zero. With no supplied data and zero coverage, all models revert to their prior, resulting in an NMSE of 100% for both the black-box and input augmentation models. The poor performance of black-models at low levels of coverage is to be expected, however, the input augmentation model failed to offer improvement despite being supplied with the same Morison’s Equation prediction and prior C_d and C_m coefficients as the residual modelling approach. The physics-based component was unable to assist the input-augmentation model at low coverage due to being supplied as GP input, since the inputs of a GP may only influence a prediction within a proximity to observed data determined by the lengthscale. When far from observed data, the inclusion of Morison’s equation as an input will have minimal benefit.

The performance of residual modelling suffers significantly less when the coverage is reduced, indicating an improvement in extrapolative capabilities. In the case where no data are supplied, the black-box component of the model reverts to a zero prior, so that the predicted output of the model is now just the prediction of Morison’s equation. The usage of Morison’s Equation with prior C_d and C_m coefficients was able to achieve an NMSE of 54.03%, a significant improvement over the black-box and input augmentation models. The white-box acts as a baseline performance

for the model which may be improved if data are provided but will not override the improved black-box predictive capabilities in areas where data are available. Residual modelling combined the same white-box and black-box components as the input augmentation model but in a means that achieved superior extrapolative performance.

An alternative investigation into the effect of coverage on model performance assumed an existing white-box model could be used to assist with predictions, with C'_d and C'_m fixed independently of the supplied training data. This scenario represents the possible case of machine learning implementation within industry where a white-box model is already established and in use. To mimic an established white-box model, the coefficients of Morison's equation were established via Bayesian Linear Regression using the complete validation set, which would not ordinarily be used in white-box model creation, and kept constant throughout the variation of coverage. Results of model NMSE from 0% to 80% coverage levels in 5% coverage intervals is shown in Figure 5.10.

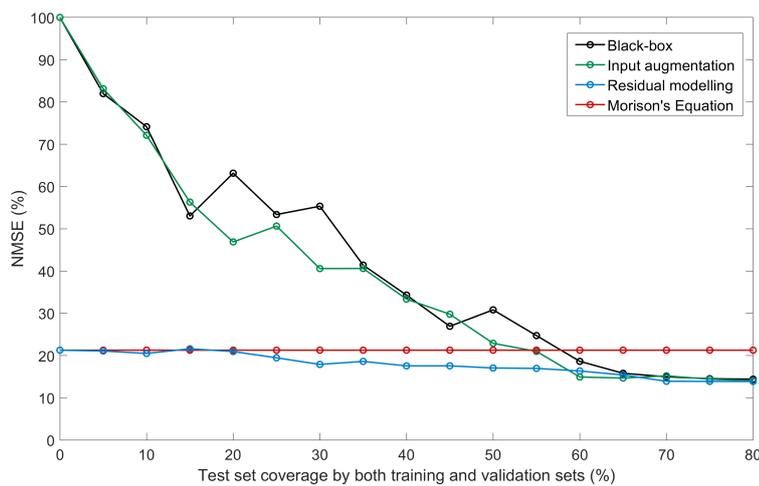


Figure 5.10: NMSE vs coverage for the MC MPO of black-box and grey-box GP-NARX models and Morison's Equation with C'_d and C'_m calculated independently of the supplied training data and kept constant throughout the variation of coverage.

By incorporating an existing white-box model rather than beginning model construction from scratch, a significant increase in residual modelling performance at low coverage levels can be seen. Where previously, the model performance here was heavily dependant on the selected prior C_d and C_m coefficients, the model is now able to revert to the prediction of the existing white-box model, thereby

reducing model NMSE to 21.26% at zero coverage. Provided that the existing model has been validated for use over the intended prediction range, this highlights the benefit of incorporating existing models within the newly created architectures. At high levels of coverage the incorporated existing model has a minimal effect on predictions, with models relying more heavily on their black-box component in areas where data are available.

5.3.4 Computation time

The trade-off between performance and computational demand is an important consideration, particularly within industrial applications. Across all models, perhaps unsurprisingly, the general trend is that those that perform best require an increased computation time. For the training and prediction of a single model this ranges from 0.9 seconds for the linear regression, to an average of 81 seconds for GPs and 2 hours 11 minutes for the extreme case of GP-NARX models.³ This difference is only exaggerated when taking in to account repeated optimisation runs to ensure stable convergence of GP and GP-NARX hyperparameters.

There are two major reasons for the considerably higher computation time of the GP-NARX models: the complexity and computational demand of the MPO NLPL cost function and the requirement of MC samples to propagate uncertainty within the output. During optimisation, for any considered approach [130], the complexity of the cost function is tied to the computational cost of the search. For the case of the MPO NLPL, the cost function was both slow to evaluate due the requirement of propagation through an independent validation set and highly sensitive to small changes in parameters, thereby creating a complex search space with high numbers of local minima. Hyperparameter sensitivity was a particular issue within the GP-NARX MPO and MC MPO due to the feedback of predictions within the model. This meant that the hyperparameters not only had an effect on predictions at the current time step but again for every instance the prediction was used as a lagged input. The high hyperparameter sensitivity meant that an increased swarm size and tighter convergence tolerance had to be used with QPSO to ensure stable optimisation which further slowed computation time. Although the larger factor within the overall run time of the GP-NARX model, the MPO NLPL cost function only affected the training time of the model, whilst the requirement of MC samples for the propagation of uncertainty affected the prediction time of the model. The computational demand within prediction is a priority for machine

³Runtimes achieved on a laptop with specification: 16GB RAM, Intel i7-9850H processor (6 core, 2.60-4.60GHz)

learning techniques as models are generally required to be trained once but make predictions repeatedly. In order to reduce prediction time, it would be possible to explore reductions within the number of MC samples used and achieve compromises between predictive stability and computation time.

An alternative means by which to reduce computation cost, with a more specific focus on training time, is through the use of sparse GPs. A subset of training points or set of pseudo-input points [266] is used to approximate the true posterior of the GP. The computational cost of training is reduced from $O(n^3)$ to $O(nm^2)$ [267], where n is the number of data points and m is the size of the subset. These methods are most useful in cases of very large datasets, with a training set of $n = 700,000$ and $m = 1000$ being effective for the estimation of flight delay times [249]. Their implementation should be considered if the dataset was expanded to cover a wider range of conditions.

5.4 Conclusions

The combining of physics-based white-box and data-based black-box modelling techniques in the form of two grey-box architectures was found to offer benefits in predictive performance over either approach used alone. The best performing grey-box model, the residual modelling GP-NARX, achieved an NMSE of 13.840% and a MSLL of -0.835, an improvement over both Morison's Equation (19.528% and -0.813) and a black-box GP-NARX (14.643% and -0.788). It was expected that grey-box models would be of specific help in assisting with extrapolation, an area in which data-based methods typically experience difficulty, and this was indeed found to be the case. Residual modelling achieved improved performance outside the range of observed training conditions, particularly in instances where a pre-established white-box was available for inclusion in to the combined model.

This chapter investigated the combining of Morison's Equation with both GP and GP-NARX regression models to predict wave loading on a structure in a real sea environment. It presented the first successful case of integrating physics and data-based methods for wave loading prediction. Two key challenges with employing the GP-NARX formulation addressed here were the selection of appropriate lag terms and the proper treatment of uncertainty propagation within the dynamic GP. A key limitation of the model, addressed in the next chapter, is a reliance on measured flow conditions close to the structure. These were used as inputs for both Morison's equation and the GP-NARX and would require the installation of flow meters.

ENHANCING GAUSSIAN PROCESS WAVE LOADING PREDICTIONS USING LINEAR WAVE THEORY

This chapter builds upon the grey-box model development of Chapter 5, in which models relied on measurements of flow conditions close to a structure to predict the wave load. These measurements are rare and can be difficult and costly to obtain. Here, aspects of linear wave theory are integrated within a GP-NARX framework to remove this reliance. Utilising physical knowledge, models are developed which rely on only incoming wave height as an input, a commonly available measured variable on many offshore structures. Models are developed using data from the monopile experiment detailed in Chapter 4, which was designed and carried out by the author. The developed PIML models provide a quantification of uncertainty alongside predictions as well as insight into the breakdown of wave load components.

6.1 Measurement of waves

The quantification of wave loading, through either measurement or prediction, presents its own set of challenges. The direct measurement of wave loads acting on offshore structures is rare, and where attempted it often requires the development and installation of bespoke systems. Even when measurements may be available, these are generally at point locations and do not provide access to a distributed

load over the structure. The prediction of wave loads across a structure, using data readily available from other sensors has the potential to provide access to wave loads where they can not be measured and reduce the cost of implementing additional measurement equipment. This chapter focuses on the utilisation of incoming wave height data, commonly available from wave radars across many offshore structures [33, 34]. Incoming wave heights are often used for the construction of directional wave spectra [268, 269], which provide an interpretable way to present a snapshot of a given sea state. Petersen [270] utilised directional spectra alongside accelerometer measurements to estimate the wave loading on the Bergsøysund bridge.

The role of a wave radar on an offshore structure was replicated in the monopile experiment, introduced in Chapter 4, through the use of wave gauges. Images of both are shown in Figure 6.1. Although both pieces of equipment provide a measurement of wave height, there are a number of differences in capability worth highlighting. Firstly a wave radar is a non-contact measurement device, able to be attached to a structure high above the waters surface, whereas a wave gauge must be partially submerged to work. In the harsh conditions of a real offshore environment, the expected lifetime of a wave gauge would be much lower than that of a wave radar. A second advantage of wave radars is that they are able to measure wave profiles over an area, therefore providing an indication of wave direction along with an instantaneous wave profile [271]. A wave gauge, however, is only capable of providing a point estimate of water depth, which may be used to estimate wave profiles over time.

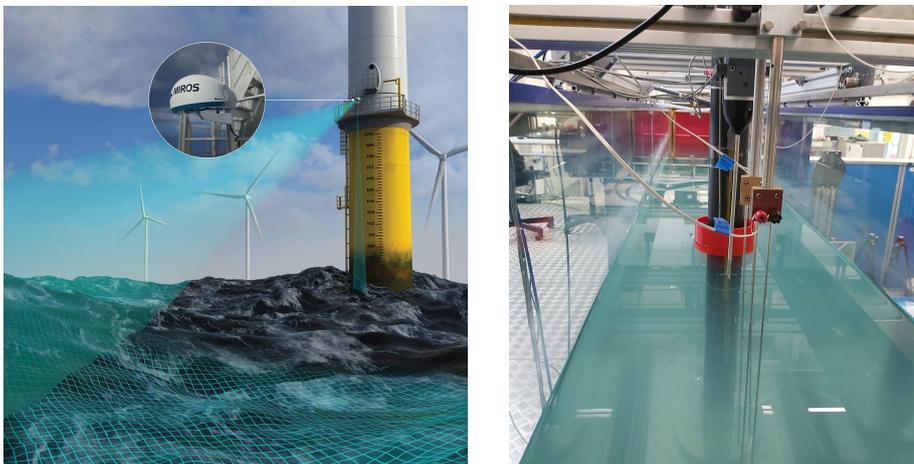


Figure 6.1: Images of a MIROS WAVEFUSION [271] wave radar installed on an offshore wind turbine (left) and the wave gauges used within the monopile experiment (right).

Due to the experiment taking place in a laboratory wave tank, the primary advantages of the wave radar over the wave gauge were not deemed to justify the significant additional expense. The conditions within the tank were much less harsh than a real offshore environment, with smaller amplitude waves, a lack of salinity and no wind. This meant that a reduced lifetime due to submerged measurement equipment was not a major factor. An ability to measure wave direction was also not of significant benefit due to only inline waves being generated within the wave tank. For these reasons, the more cost-efficient wave gauge was the preferred choice of measurement equipment.

The work in this chapter has two main objectives: to model the wave load acting on a monopile using only data of incoming wave height and to showcase how physical knowledge may be integrated within Gaussian process NARX models to improve performance and provide interpretability. Here, linear wave theory is used to improve predictive performance by providing an approximation of flow conditions as a model input whilst Morison's Equation is used to visualise the breakdown of the wave load in to drag, inertia and excluded effects. A key conclusion drawn was that although the physical knowledge included was imperfect, relying on a number of simplifying assumptions, it was still able to assist the combined final models.

6.2 Integrating physics and data

The types of physics and data-based components used, along with the means by which they are combined, has a significant effect on the model structure produced, its performance level and levels of insight provided into predictions. Here, two pieces of physics are integrated within a GP-NARX: linear wave theory is used to approximate flow conditions for use as a model input, whilst Morison's Equation is used to provide an estimate of the drag and inertia components of the wave load. A summary of included physics is shown in Table 6.1.

6.2.1 Input augmentation: Linear wave theory

The objective of models here is to predict wave load using incoming wave height as an input, however the relationship between free surface elevation and wave force is complex. This means that not only is a flexible model required, but also that the learning task is highly challenging, placing increased demand on required training data. The role of linear wave theory is to achieve an approximation of the flow

Table 6.1: Summary of physical knowledge included and its role within the combined model.

Physics included	Aim	Method	How is aim achieved?
Linear wave theory	Performance improvement	Input augmentation	Simplification of learning task
Morison's Equation	Interpretability	Mean function	Breakdown of wave load components

conditions close to the monopile, thereby reducing the complexity of the learning task, allowing an increased performance when data availability is lessened. The change in model structure from a black-box approach is summarised in Figure 6.2.

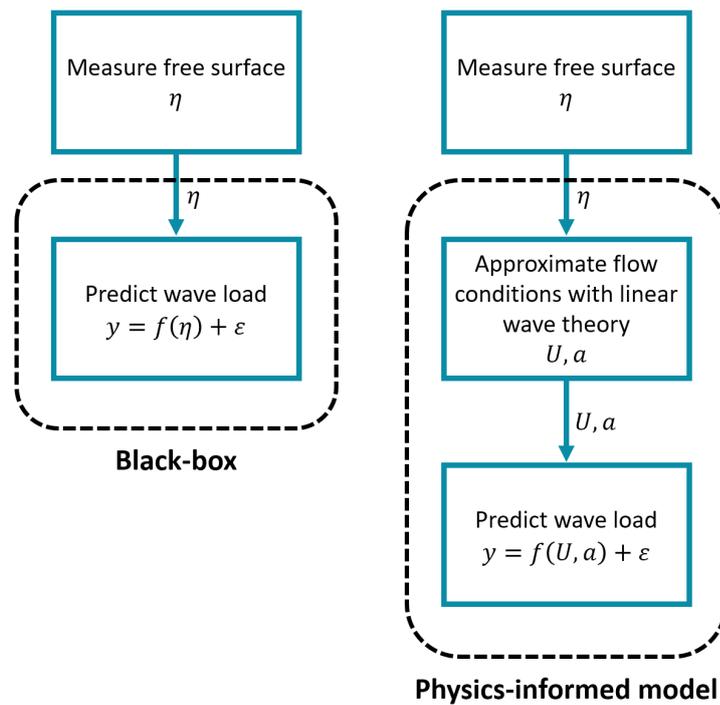


Figure 6.2: A comparison of the proposed physics-informed model structure with a black-box approach. The measured free surface η is used to approximate flow conditions (water particle velocity U and acceleration a) close to the monopile for use as an input to the model, rather than being directly used itself.

The approximation of flow conditions will be used to predict force acting on the monopile at the force collar, around the region close to the free surface ($-0.05 \leq z \leq 0.05m$). This is where flow velocities and accelerations are typically highest and where the maximum forces are likely to be experienced. The prediction of the largest forces applied to the monopile is most influential for the prediction of remaining fatigue life. Only the data from a single wave gauge will be available for use as a model input. A diagram of the relative dimensions between measurement equipment is shown in Figure 6.3.

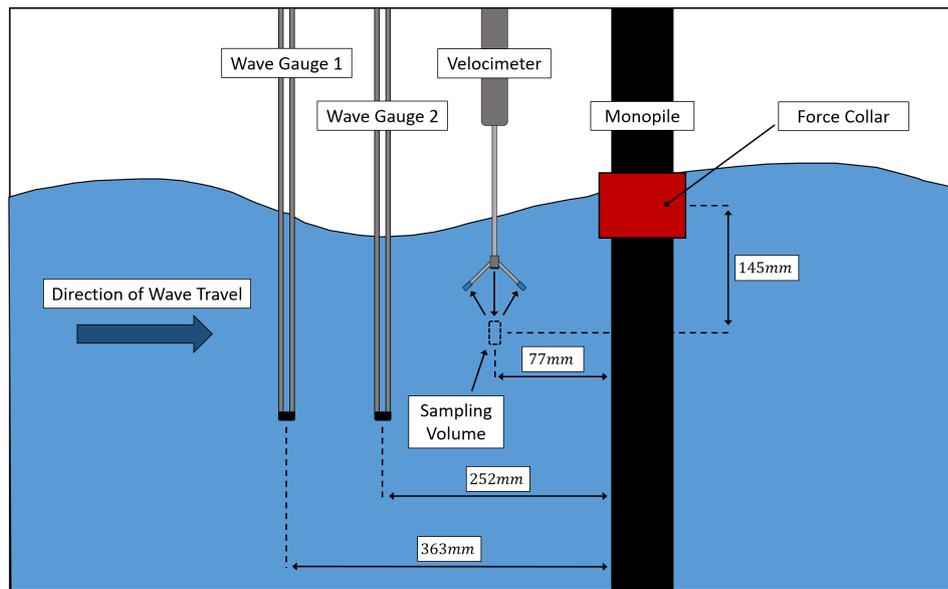


Figure 6.3: Relative positions of measurement equipment within the wave tank. Only data from wave Gauge 1 was used as input data for models. Wave Gauge 2 and the velocimeter were used for validation whilst data from the force collar was used as a training target and an unseen test target.

In order to integrate linear wave theory within the GP-NARX, the following workflow is proposed:

1. Perform a free surface reconstruction using the measured surface elevation from wave gauge 1, located 363mm from the monopile.
2. Validate the accuracy of the free surface reconstruction using the measured data from wave gauge 2, located 252mm from the monopile, none of which will be shown to the model prior to testing.

3. Construct a velocity field for the sum of linear waves obtained from the free surface reconstruction.
4. Validate the accuracy of the generated water particle velocities and accelerations generated using the measured data from the velocimeter.
5. Use the validated velocity field to achieve an estimate for the flow conditions at the force collar.
6. Use the flow conditions at the force collar as inputs within a range of models for wave loading prediction.
7. Test the performance of wave loading prediction achieved using the measured data from the force collar.

The reconstruction of the free surface follows the methods described within [272] and relies on the decomposition of the wave in to a sum of harmonic components. Here, the discrete form of the Fourier transform is employed:

$$X_k = \sum_{n=0}^{N-1} x_n \exp\left(-\frac{2\pi i}{N}kn\right) \quad (6.1)$$

This allows for the representation of a signal as a sum of N sinusoidal functions, for which the i th term may be expressed in terms of its angular frequency ω_i , amplitude A_i and phase Φ_i .

$$\omega_i = \frac{2\pi n f}{N} \quad (6.2)$$

$$A_i = \frac{1}{N} \sqrt{\Re(X_k)^2 + \Im(X_k)^2} \quad (6.3)$$

$$\Phi_i = \arctan\left(\frac{\Im(X_k)}{\Re(X_k)}\right) \quad (6.4)$$

where f is the sample rate in Hz. The i th wave number k_i is obtained through solution of the linear dispersion relation.

$$\omega_i^2 = gk_i \tanh(kd) \quad (6.5)$$

where g is acceleration due to gravity and d is depth. This allows for the expression

of free surface η as a sum of N linear waves [272].

$$\eta = \sum_{i=1}^N A_i \cos(k_i x - \omega_i t + \Phi_i) \quad (6.6)$$

For the validation of the free surface reconstruction, the surface elevation from wave gauge 1, placed 363mm from the monopile was used to determine a sum of linear waves. The accuracy of a free surface reconstruction was then tested at wave gauge 2, placed 252mm from the monopile as shown in Figure 6.4. None of the measured data from wave gauge 2 was shown to the model before performance was measured.

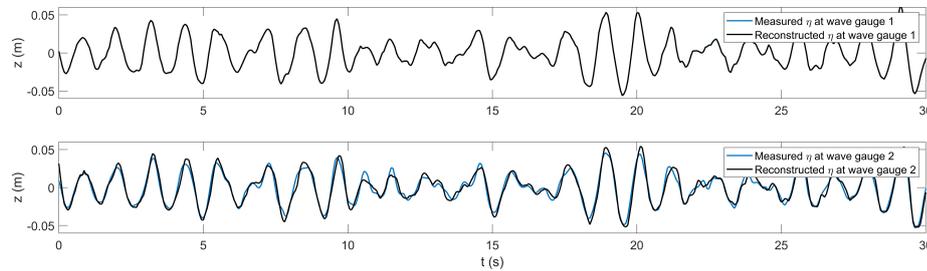


Figure 6.4: A comparison of free surface reconstruction at wave gauge 1 (top), located 363mm from the monopile, and wave gauge 2 (bottom), located 252mm from the monopile for a JONSWAP wave with $\gamma = 3.3$ and $\omega_p = 0.8Hz$.

The measured free surface elevation from wave gauge 1 was used for the reconstruction, with an average NMSE of 15.48% achieved on the free surface fit at wave gauge 2 across all JONSWAP waves.

The reconstruction of the free surface required the decomposition of the wave into a sum of linear waves, which provides a useful tool for the derivation of several properties [36]. An important one of which is the velocity potential ϕ , which when treated as a sum over N linear waves is expressed:

$$\phi = \sum_{i=1}^N \frac{A_i g \cosh(k_i(z+d))}{k_i c_i \cosh(k_i d)} \sin(k_i x - \omega_i t + \Phi_i) \quad (6.7)$$

The horizontal and vertical velocities are negative spatial derivatives of the velocity potential and their calculation for a given range of x and z will allow for the construction of a velocity field. Noting that velocities here are only meaningful

below the free surface, the horizontal velocity U_x and vertical velocity U_z are expressed:

$$U_x = -\frac{\partial\phi}{\partial x} = \begin{cases} \sum_{i=1}^N \omega_i A_i \frac{\cosh(k_i(z+d))}{\sinh(k_i d)} \cos(k_i x - \omega_i t + \Phi_i) & z \leq \eta \\ 0 & z > \eta \end{cases} \quad (6.8)$$

$$U_z = -\frac{\partial\phi}{\partial z} = \begin{cases} \sum_{i=1}^N \omega_i A_i \frac{\sinh(k_i(z+d))}{\sinh(k_i d)} \sin(k_i x - \omega_i t + \Phi_i) & z \leq \eta \\ 0 & z > \eta \end{cases} \quad (6.9)$$

from which the horizontal acceleration a_x and vertical acceleration a_z are derived as

$$a_x = -\frac{\partial^2\phi}{\partial x \partial t} = \begin{cases} \sum_{i=1}^N \omega_i^2 A_i \frac{\cosh(k_i(z+d))}{\sinh(k_i d)} \sin(k_i x - \omega_i t + \Phi_i) & z \leq \eta \\ 0 & z > \eta \end{cases} \quad (6.10)$$

$$a_z = -\frac{\partial^2\phi}{\partial z \partial t} = \begin{cases} \sum_{i=1}^N \omega_i^2 A_i \frac{\sinh(k_i(z+d))}{\sinh(k_i d)} \cos(k_i x - \omega_i t + \Phi_i) & z \leq \eta \\ 0 & z > \eta \end{cases} \quad (6.11)$$

The construction of velocity fields allows for the visualisation of an important property of flow conditions within waves, the orbital motion of water particles. The size and shape of orbits is affected by both the wave and relative depth of the water, causing a non-trivial relationship between different points within the wave field. In extreme cases, it is possible for nearby points within a wave field to have velocities in almost opposite directions. This factor contributes to the difficulty of predicting a wave load on a structure from nearby flow conditions.

By calculating the velocities and accelerations at the location of the velocimeter, where there is access to the measured variables, it was possible to test performance. It should be noted that it is not expected that results should match exactly, the true behaviour of fluid motion is far more complex than defined within linear wave theory. The aim of the velocity field reconstruction was to provide an approximation of flow conditions that would be useful as an input for wave loading prediction models. The reconstruction also has the advantage of only using the velocimeter as a validation step, with only data from a wave gauge (a cheap and widely implemented sensor) being shown to the model. A timeseries fit of the velocities and accelerations at the velocities achieved at the velocimeter location is shown in Figure 6.5.

For all wave conditions generated within the tank, an adequate approximation of the

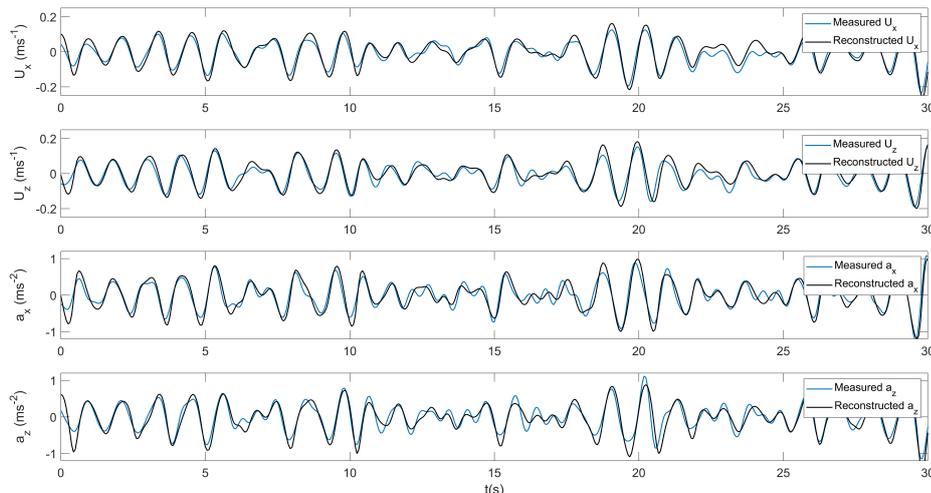


Figure 6.5: A comparison of water particle velocities and accelerations at the velocimeter from a JONSWAP wave with $\gamma = 3.3$ and $\omega_p = 0.8Hz$.

flow conditions at the velocimeter was achieved via the velocity field reconstruction. A summary of validation performance for the velocities, accelerations and free surface is shown in Table 6.2. Reconstruction accuracies of 15 – 20% were deemed acceptable, and around the order expected given the use and limitations of linear wave theory [35, 36].

Table 6.2: Summary of average NMSE achieved across all JONSWAP waves, when estimating the horizontal velocity U_x , vertical velocity U_z , horizontal acceleration a_x and vertical acceleration a_z at the location of the velocimeter and the free surface η_2 at wave gauge 2.

Variable	U_x	U_z	a_x	a_z	η_2
NMSE (%)	15.693	18.853	17.959	20.602	15.480

For the approximation of flow conditions at the location of the force collar, the region of the velocity field where $x = 363mm$ and $(-0.05 \leq z \leq 0.05m)$ was evaluated. This represented its location w.r.t wave gauge 1 where the reconstruction of linear waves was performed. A diagram of the relative positions of the wave gauges, velocimeter sampling volume and force collar surface is shown in Figure 6.6.

The reconstructed force collar velocities were averaged to provide a single represent-

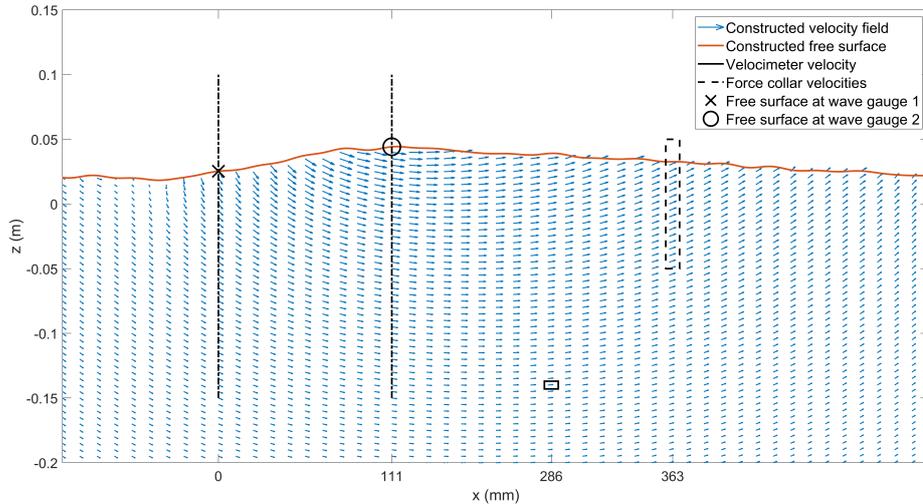


Figure 6.6: The relative positions of the wave gauges, velocimeter sampling volume and force collar surface within a reconstructed velocity field from a JONSWAP wave with $\gamma = 3.3$ and $\omega_p = 0.8Hz$.

ative velocity to calculate the total force acting on the force collar. A convergence study on the vertical grid spacing dz was performed to a velocity stable with 1% tolerance, leading to a selection of $dz = 0.2mm$.

Where the reconstruction of the flow conditions at the velocimeter could be compared with measured data, this was not the case at the force collar as it was not possible to measure flow in this location. However, it should be emphasised that the flow conditions themselves are not the quantities of interest; the wave load is. The success of the velocity field reconstruction at the force collar shall therefore be determined not by the accuracy of the reconstructed flow conditions, but whether they provide a useful input to assist a wave loading prediction model. The steps to validate the constructed sum of linear waves indicate a satisfactory representation of conditions at the velocimeter. Whether this was repeated at the force collar will be in part determined by the performance of wave loading predictions models presented in the next section.

6.2.2 Mean function: Morison's Equation

Morison's equation approximates two components of a wave load: the drag force, caused by the relative motion of fluid and a body, and the inertia force, caused by

pressure gradients within an accelerating fluid. The accuracy of Morison's Equation is, in part, determined by the extent to which a wave load is dominated by the drag and inertia components. In reality, there are many other excluded effects, contributing towards the total wave load [36, 222], including:

- Diffraction forces, caused when a wave pattern is modified when moving past a structure. These occur when the diameter of a structure is comparable in size, or larger, than the wavelength of the wave.
- Pressure forces, induced by the fluctuating pressure within a moving wave. These typically include higher-order terms and are more significant for steep waves (a large amplitude relative to wave length).
- Oscillatory forces, caused by the periodic shedding of vortices. These will occur at the Strouhal number $St = \frac{fD}{U}$, where f is the shedding frequency, D is the diameter of the structure and U is flow velocity.
- Boundary interactions, caused by the restriction of fluid flow nearby the structure. These may be caused by the close proximity of other structures (e.g. dense array of monopiles, multiple legs of a jacket structure) or presence of walls (e.g. wave tank).
- Damping forces, from the dissipation of energy when a structure is moving relative to a fluid. These are enhanced when a structure experiences sway.

Morison's Equation is known to neglect the above effects [222], which although is often a valid assumption for slender structures in deep water, is not expected to hold within the experimental setup here. The 90mm diameter monopile within a 0.5m wide wave tank results in a blockage ratio of 0.18, a value high enough to modify the passing wave pattern [222]. To capture the total wave load, a model structure is proposed of:

$$F = \underbrace{\frac{1}{2}\rho DC_d U|U|}_{\text{Drag}} + \underbrace{\frac{1}{4}\pi\rho D^2 C_m \dot{U}}_{\text{Inertia}} + \underbrace{f(x)}_{\text{Excluded effects}} \quad (6.12)$$

where $f(x)$ is a GP-NARX. Within the GP framework, this is equivalent to a Morison's Equation mean function, as in the previous chapter, with the residuals learned from data. The estimation of other effects, including diffraction, and pressure forces often involves the use of CFD and numerical methods [273, 274], which can result in high computational demands. Here, it is hoped that the flexibility of the GP-NARX may learn its structure from data.

A key challenge facing the implementation of the Morison's Equation mean function is an appropriate method of selecting C_d and C_m . The presence of other effects can impact the calculation of C_d and C_m , often leading to unrealistic coefficients to achieve a best fit to data. It is known that C_d and C_m are dependent on variables related to flow conditions, including Reynold's number Re and Keulegan-Carpenter number Kc , along with variables related to the structure, including surface roughness k and slenderness $\frac{D}{L}$. This has led to the use of standardised formulae, tables and charts from which to estimate C_d and C_m based from these variables. Here, DNV-RPC-205 Recommended Practice [234] is used to achieve appropriate estimates of the drag and inertia components of the wave load. It should be noted that these standardised coefficients aim to serve only an approximation and would not be expected to outperform fitted coefficients if the drag and inertia components dominated the wave load.

6.3 Results on the Monopile Experiment Data

This section aims to present and compare the performance of wave loading prediction models with a specific focus on the benefits of integrating physical knowledge with data. Each model was tested across a matrix of JONSWAP waves as a function of peak enhancement factor γ and peak frequency ω_p as detailed in Chapter 3. Each model was trained on 300 data points and tested on an unseen set of 1000 data points.

6.3.1 Morison's Equation

The approximation of flow conditions as described in 6.2.1 combined with the use of Morison's Equation, is able to provide a means of wave load prediction using only wave gauge data as an input with very low computational expense. Here, this will be considered a purely physics-based approach. A surface of NMSE performance over the range of wave conditions for Morison's Equation, with inputs provided through linear wave theory, is shown in Figure 6.7.

The performance of Morison's Equation ranged from 34.43-51.34% NMSE, somewhat above the $\sim 20\%$ residual [257] that might be hoped for when Morison's Equation is performing to its potential. There are two primary reasons for this observed drop in performance: the modelling process here begins with wave height rather than flow conditions, presenting an additional challenge within the task; and there is expected occurrence of other effects caused by the monopile disturbing the

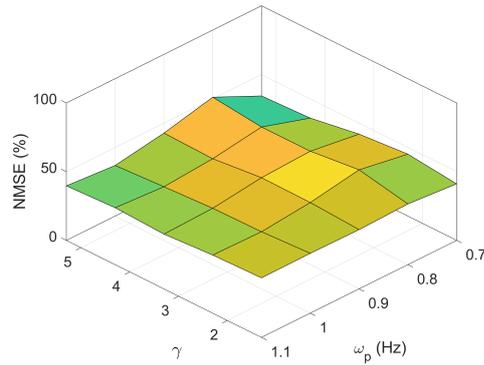


Figure 6.7: NMSE surface of Morison’s Equation using the approximated flow conditions at the force collar as inputs. Models were tested over a grid of JONSWAP waves as a function of peak enhancement factor γ and peak frequency ω_p .

passing wave profile within the narrow tank. Morison’s Equation is known to only approximate the drag and inertia components of the wave load and it is unfair to expect optimum performance from its implementation in this case.

A key performance trend observed within Figure 6.7 is a visible ridge along the 0.8Hz peak frequency wave states. The natural frequency of the structure was intentionally dropped within the operating range of the wave tank to approximately 0.8Hz , such that wave loading prediction around resonance could be observed, and this being the region of worst performance was to be expected. When excited close to a natural frequency, the motion of the structure will be largest, thereby having the greatest impact on the relative motion of the structure and fluid. Although extensions to Morison’s have focussed on accounting for moving structures [275], the standard form implemented here assumes a static structure. The aim of utilising empirical methods here was to provide a computationally efficient start point from which to develop improved model structures.

6.3.2 Autoregressive modelling

The development of linear autoregressive models provides a useful tool to bridge the gap between Morison’s Equation and GP-NARX modelling, providing a compromise between performance and computational cost. A specific motivator for their use here is their role within lag selection, for which many model optimisations are required. The completion of a similar process using GP-NARX models would have

been computationally infeasible.

AutoRegressive models with eXogenous inputs (ARX) are considered here of the form:

$$y_t = \sum_{i=0}^{l_u} \alpha_i u_{t-i} + \sum_{i=1}^{l_y} \beta_i y_{t-i} + \varepsilon \quad (6.13)$$

where the previous signal values, y_{t-i} are the wave force and the exogenous inputs u_{t-i} considered were the approximated velocity, U and acceleration, \dot{U} of the wave particles at the force collar location. This form of ARX model mimics the structure of Morison's equation, with lagged terms adding capability to cope with the upstream position of measurement equipment and representation of higher order terms. Alternatively, one may use the measured free surface elevation η as the input to the ARX model. This case is also considered here, alongside the Morison structured ARX, as a purely data-based approach as it would not rely on either linear wave theory or Morison's equation for its construction.

Lag selection: A grid search over grids

The determination of the number of selected lags for the exogeneous inputs, L_u and the model output, L_y is an important step in autoregressive model construction. The selection of lagged terms has significant effect on the structure and performance of the model. In some cases, the selection of lags may even change the name of a model, with a summary of this shown in Table 6.3. The key challenge of lag selection is computational cost; the modification of lagged terms requires the construction, fitting and evaluation of a new model. Here, this problem is amplified due to the range of generated wave states. For each generated JONSWAP wave state, in terms of γ and ω_p values, the lag selection process must be repeated.

Table 6.3: Categorisation of *linear-in-the-parameters* autoregressive Models depending on considered lag space. *Here, AR models are considered to depend only on lagged values of the target y_{t-i} .

	$L_u = 0$	$L_u > 0$
$L_y = 0$	Linear regression	Finite impulse Response (FIR) model
$L_y > 0$	AR model*	ARX model

Consistent with the approach of Chapter 5, the Bayesian Information Criterion

(BIC) [264] was selected as a model selection criterion, against which candidate lags were evaluated. It rewards model performance whilst penalising more complex models. Lag selection was done using the Morison structured ARX model as a grid search up to a maximum of 20 lags, beyond which models would be become increasingly likely to overfit. A heatmap of ΔBIC is shown for the JONSWAP wave file $\gamma = 3.3$, $\omega_p = 0.8$ Hz in Figure 6.8. The ΔBIC is the difference in BIC between the optimum model and any other result and is useful as a method of comparison between models. The optimum model will have $\Delta BIC = 0$, with a high BIC being an indicator of a poor model selection.

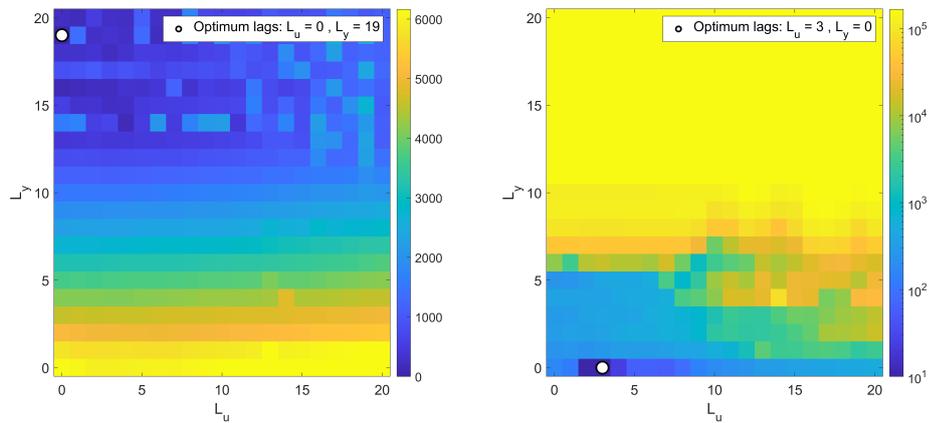


Figure 6.8: ΔBIC heatmap for OSA and MPO predictions of the Morison ARX model for a JONSWAP wave with $\gamma = 3.3$ and $\omega_p = 0.8$ Hz. For an OSA prediction, inputs to the model include previously measured values of the target, whereas for the MPO, they are previous predictions of the target.

Since the selection of lags defines the structure of the model and a single model was required to work across a range of conditions, it was necessary to select a single combination of optimum lags. The selection of multiple optimum lags for each set of generated wave conditions would require the construction of multiple models. To achieve a single lag selection, the lag selection process was first repeated for each γ and ω_p dependent JONSWAP wave. An overview of heatmaps is shown in Figure 6.9. The results were then averaged across the set of conditions to provide a single indicator of how a lag selection performed across a range of conditions. The averaged ΔBIC heatmap across the generated set of JONSWAP waves is shown in Figure 6.10.

The optimum lags selected were $L_u = 4$ and $L_y = 0$, indicating that the lagged flow conditions provided a more useful model input than lagged predictions of the

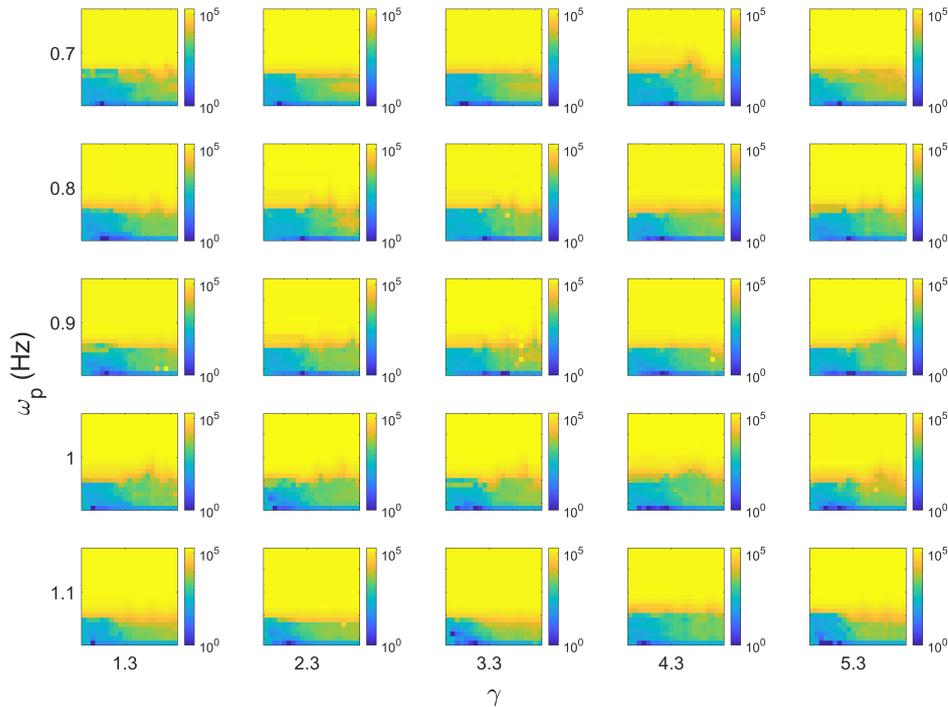


Figure 6.9: ΔBIC heatmaps from the Morison ARX model for the generated set of JONSWAP waves as a function of γ (columns) and ω_p (rows). Each subplot represents a grid search of lags L_u and L_y for a particular generated JONSWAP wave. A consistent pattern of ΔBIC suggests the selection of a single set of L_u and L_y would be appropriate.

wave force. This is expected to be due to a compounding of errors due to the prediction feedback. Even errors that begin small, have the potential to grow with each feedback cycle of an autoregressive model. Autoregressive model structures dependent on exogenous inputs only are sometimes referred to as Finite Impulse Response (FIR) models (See Table 6.3).

Visualisation of the cross correlation can provide a helpful tool in interpreting the selected lags. A high correlation between variables, particularly at specific lagged time steps, is an indicator of a useful model input. The cross correlation of the measured wave force F with itself, horizontal velocity U_x , horizontal acceleration a_x and predicted wave force \hat{F} is shown in Figure 6.11.

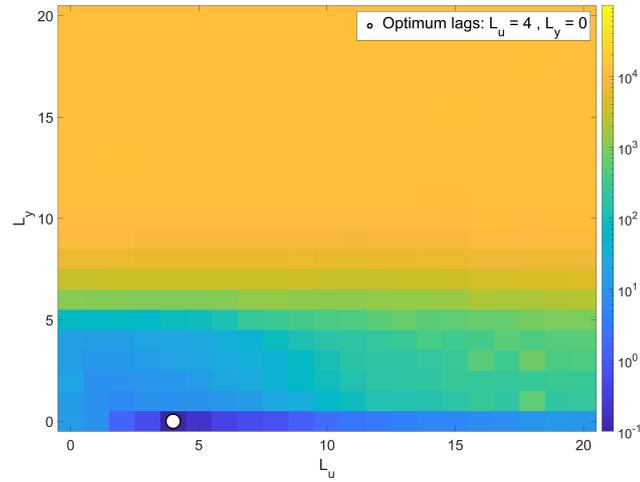


Figure 6.10: Averaged ΔBIC heatmap from the Morison ARX model across the generated set of JONSWAP waves as a function of γ and ω_p . An optimum lag selection of $L_u = 4$ and $L_y = 0$ was achieved.

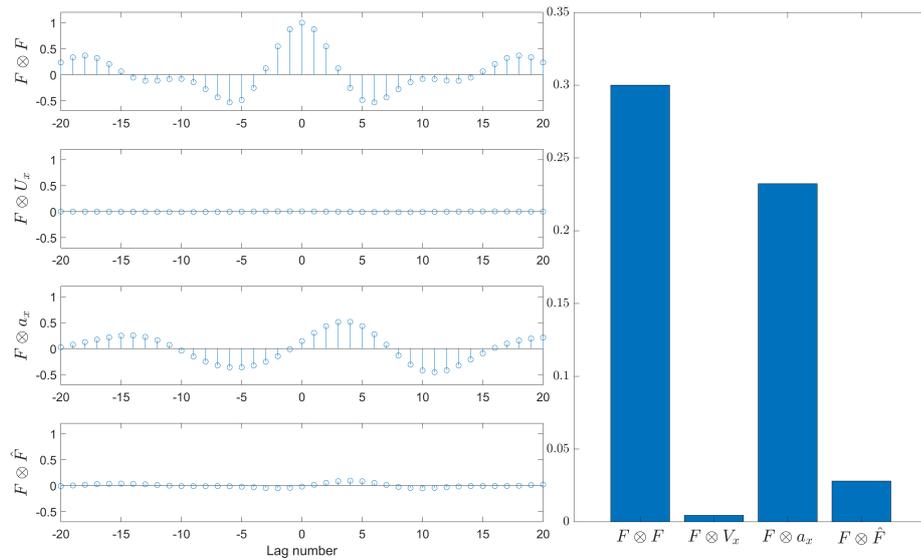


Figure 6.11: The cross correlation of the measured wave force F with itself, horizontal velocity U_x , horizontal acceleration a_x and predicted wave force \hat{F} (left) and their average magnitude of cross correlation (right) within a 20 lag range.

A few key observations may be drawn from the cross correlations of Figure 6.11. Firstly, there is a large difference in average correlation magnitude between the measured and predicted wave loads (F and \hat{F}), highlighting a key difference between OSA and MPO model structures. For OSA models, where access to previous measurements of the target variable are available, the selection of these as lagged model inputs is generally much more likely due to their high correlation with the current target. For the MPO model constructed here, the lower correlation of the lagged predicted wave loads with the current measured target explains why no output lags were deemed preferable. Another important observation is that the magnitude of correlation for horizontal acceleration a_x was much higher than that of horizontal velocity U_x . This is an indicator of the flow regime and aligns with what would be expected when the wave interaction is inertia dominated (Keulegan-Carpenter number < 1 and Reynolds number > 4000). Here, the acceleration based C_m term within Morison's equation has most influence on predictions.

Autoregressive model performance

Here, two autoregressive model structures are tested, both of which will use lags $L_u = 4$ and $L_y = 0$ found during lag selection for fairness of comparison. For any given JONSWAP wave state, each model was trained on the same 300 data points and tested on an unseen set of 1000 data points. The model structures are summarised as:

- The Morison FIR model, using the approximated water particle velocity and acceleration as inputs:

$$y_t = \sum_{i=0}^{l_u} \alpha_i^{(1)} U_{t-i} |U_{t-i}| + \sum_{i=0}^{l_u} \alpha_i^{(2)} \dot{U}_{t-i} + \varepsilon \quad (6.14)$$

- The wave gauge FIR model, using the measured free surface from the wave gauge directly as an input:

$$y_t = \sum_{i=0}^{l_u} \alpha_i \eta_{t-i} + \varepsilon \quad (6.15)$$

A comparison of the performance of wave gauge input FIR model and the Morison's Equation FIR model using the approximated flow conditions as inputs is shown in Figure 6.12.

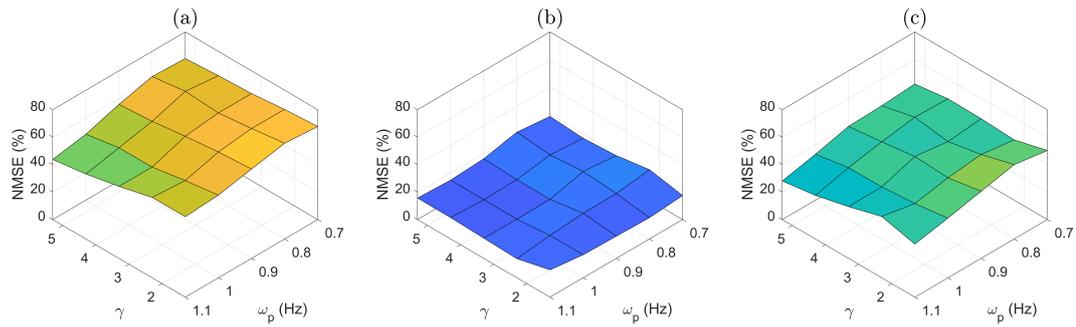


Figure 6.12: (a) NMSE surface of a FIR model using wave gauge data as an input. (b) NMSE surface of a FIR form of Morison's Equation as a function of approximated flow conditions. (c) The improvements in performance of the Morison FIR over the wave gauge FIR model. Models were tested over a grid of JONSWAP waves as a function of peak enhancement factor γ and peak frequency ω_p .

The approximation of flow conditions at the force collar for use as a model input rather than using the wave height directly was able to improve performance over the full range of wave states by between 28.1% and 54.2% NMSE. Although a flexible model, the FIR struggled to adequately capture the relationship between the wave height and wave load and by approximating the flow conditions, the complexity of the learning task was reduced, thereby increasing performance. The consistency of results was also aided, observed within the flatter NMSE surface across the various wave conditions for the Morison FIR model. The greatest areas of improvement were seen in the low peak frequency waves, for which the amplitude of waves was larger. Larger waves produce larger particle velocities and accelerations, and therefore larger wave loads, for which accurate prediction is more impactful for fatigue life estimation.

The autoregressive model structure helped the Morison FIR model outperform the standard form of Morison's Equation used within Figure 6.7. This was due to a number of factors: the Morison FIR had a higher number of model parameters (8 vs 2), allowing for increased flexibility and an ability to model more complex relationships; the lagged timesteps may have been useful in helping capture the time delay between the monopile and upstream measurement equipment; and the multiple time lags may have been used to approximate higher order terms, being helpful in capturing phenomena typically excluded by Morison's Equation e.g. vortex shedding [228].

6.3.3 GP-NARX

A GP-NARX¹ provides the advantages of autoregressive modelling previously discussed within a GPR framework. The ability of a GP to model relationships without a restriction to a fixed functional form is especially helpful when a process is only partially understood. Here, the SE kernel, suitable for modelling smooth continuous functions, is used to capture phenomena present within wave loading. A purely data-based approach is first considered in which the wave gauge data are passed directly to the GP-NARX as an input, with the relationship between them learned. This is then compared with a physics-informed case of approximating the flow conditions at the force collar using linear wave theory for use as a GP-NARX input. A comparison of the performance for these cases across the range of wave states is shown in Figure 6.13. The same training (300 points) and testing (1000 unseen points) were used for each JONSWAP wave state as previous models.

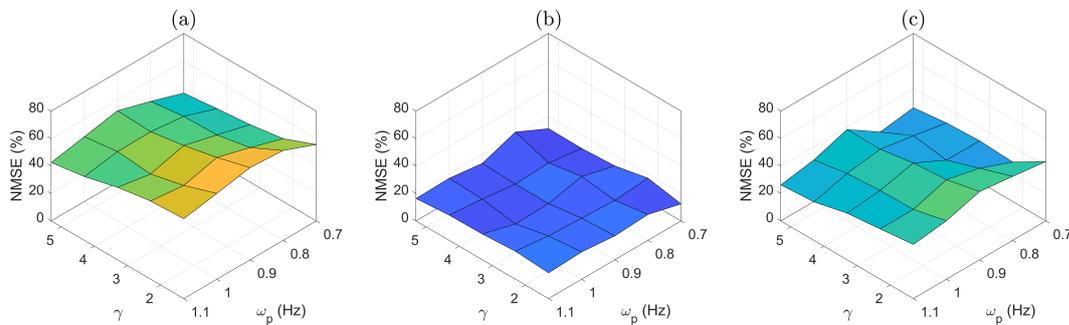


Figure 6.13: (a) NMSE surface of a GP-NARX with a SE kernel using wave gauge data as an input. (b) NMSE surface of a GP-NARX with a SE kernel as a function of approximated flow conditions. (c) The improvements in performance achieved as a result of using the approximated flow conditions rather than the wave gauge data as an input to the GP-NARX. Models were tested over a grid of JONSWAP waves as a function of peak enhancement factor γ and peak frequency ω_p .

The best model in terms of mean square error is the GP-NARX using physically-informed inputs with an NMSE of 15.93%, a very respectable performance, considering only incoming wave height data was required to be measured. As well as highlighting the results in terms of mean squared error, the discussion here aims to explain motivation for the use of machine learning and physical knowledge

¹Since the GP-NARX was constructed with $L_y = 0$, it might be more accurately described as a GPX, however the term GP-NARX will be used here for consistency with existing literature.

in combination. Beginning with data-based benefits, both implemented cases of the GP-NARX were able to outperform their equivalent linear FIR models with 51.37% vs 59.17% NMSE for the wave gauge input, and 15.93% vs 17.05% for the reconstructed flow conditions. The increased flexibility of the GP-NARX, much of which is provided through the SE kernel, was better able to capture the complex relationship between variables. The bigger benefits, both in terms of absolute and proportional improvements, were seen when wave gauge data was used as an input. Here, the learning task was more challenging, favouring more flexible models. The performance of the wave gauge FIR model was also worse, meaning that there was more scope for absolute improvement.

Even though a GP-NARX is a flexible and effective data-based technique within a variety of applications, the inclusion of the reconstructed flow conditions as a model input was still able to improve prediction quality. By reducing the complexity of the learning task, the use of linear wave theory was able to assist the machine learning model, even though it only provides an approximation of flow conditions. This highlights perhaps the major finding of this work, that physical knowledge integrated within machine learning models does not have to be exact to be helpful. The limitations of linear wave theory are well understood [35, 36], and are unlikely to fully hold even within laboratory environments. This is observed within the $\sim 15 - 20\%$ errors in the velocities, accelerations and free surface within Table 6.2. However, whether or not the approximated flow conditions were close to their true value, they were still able to assist the GP-NARX with wave loading prediction, which was their primary goal.

6.3.4 Interpretability through inclusion of prior knowledge

Where the approximation of flow conditions was able to offer improvements in performance, the inclusion of empirical methods aims to provide interpretability within results. This is done by providing a breakdown of the wave load in to drag, inertia and other effects. Due to the obstruction caused by the monopile within the relatively narrow wave tank (blockage ratio of 0.18), the wave profile cannot be assumed to be undisturbed and effects excluded by Morison's Equation should not be ignored [222].

When used alone, Morison's Equation and suitable values for C_d and C_m may adequately estimate the drag and inertia components of a wave load within appropriate flow regimes. However, in the presence of other effects, the fitting of C_d and C_m to data becomes challenging, with the correct estimation of drag loading often biased. Fitted coefficients in the presence of other effects may result in moderate

performance, but without physically meaningful values. The 34.43-51.34% NMSE achieved by the fitted Morison's Equation in Figure 6.7 often produced exceedingly high C_d coefficients, with $C_d > 10$ in some cases, thereby losing physical interpretability within results. The primary aim of implementing Morison's Equation here is not to maximise the performance of Morison's Equation itself, but to provide interpretability within the results of the final model.

Here, the guidance of DNV-RPC-205 Recommended Practice [234] is used to define C_d and C_m , with the aim being to approximate the drag and inertia loading without being skewed by the presence of other effects. A GP-NARX with a SE kernel is then used to capture the residual to sum within the final model. The performance of Morison's Equation accounting for only drag and inertia loading, along with GP-NARX residual models also accounting for excluded effects is shown in Figure 6.14.

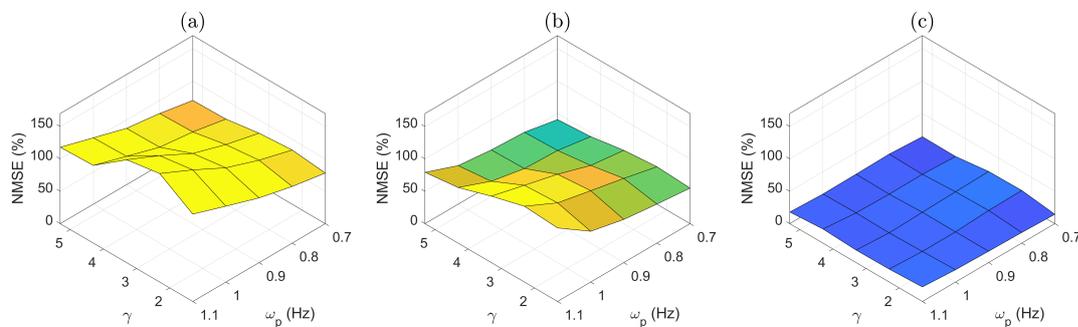


Figure 6.14: (a) NMSE surface of Morison's Equation when used with the C_d and C_m coefficients suggested by DNV-RPC-205 [234]. (b) NMSE surface for a GP-NARX residual model using wave gauge data as an input. (c) NMSE surface for a GP-NARX residual model using approximated flow conditions as an input. Models were tested over a grid of JONSWAP waves as a function of peak enhancement factor γ and peak frequency ω_p .

The poor performance of Morison's Equation used with the C_d and C_m suggested by DNV-RPC-205 [234] highlights the importance of understanding the limitations of empirical models. Used in this way, Morison's Equation only approximates the drag and inertia components of the wave load. Attempts to capture excluded effects with the GP-NARX led to improvements in performance, with greater performance achieved when the approximated flow conditions were used as inputs rather than the wave gauge data directly. It is expected, like with other models presented, that the approximation of flow conditions helped to simplify the learning task and help the GP-NARX better learn structure from smaller amounts of training data.

The breakdown of drag, inertia and excluded effects for a JONSWAP wave with $\gamma = 3.3$ and $\omega_p = 0.8Hz$ is shown in Figure 6.15. The most basic information gained from the plot is the relative magnitudes of the force components; both the inertia load and the load due to excluded effects were significantly higher than the drag load. Although here, their effects are combined, the quantification of pressure forces, diffraction forces and other excluded effects are particularly helpful for a number of specific offshore structure calculations. These include the calculation of deck heights and required air gaps, and the calculation of wave loading on nearby structures [222]. The identification of specific features within force components, including conditions in which peaks occur, can also be helpful in understanding wave loads. A trend specific to this experiment is the non-zero mean of the drag load, due to the positioning of the force collar at the water line. When the wave height at the monopile is a maximum, and velocities approach their peak, the wave load acts across the entire force collar, producing a large positive peak in drag loading. However, in the opposite case, when a wave trough passes the monopile and velocities reach their maximum negative values, the wave load acts below where the force collar is placed.

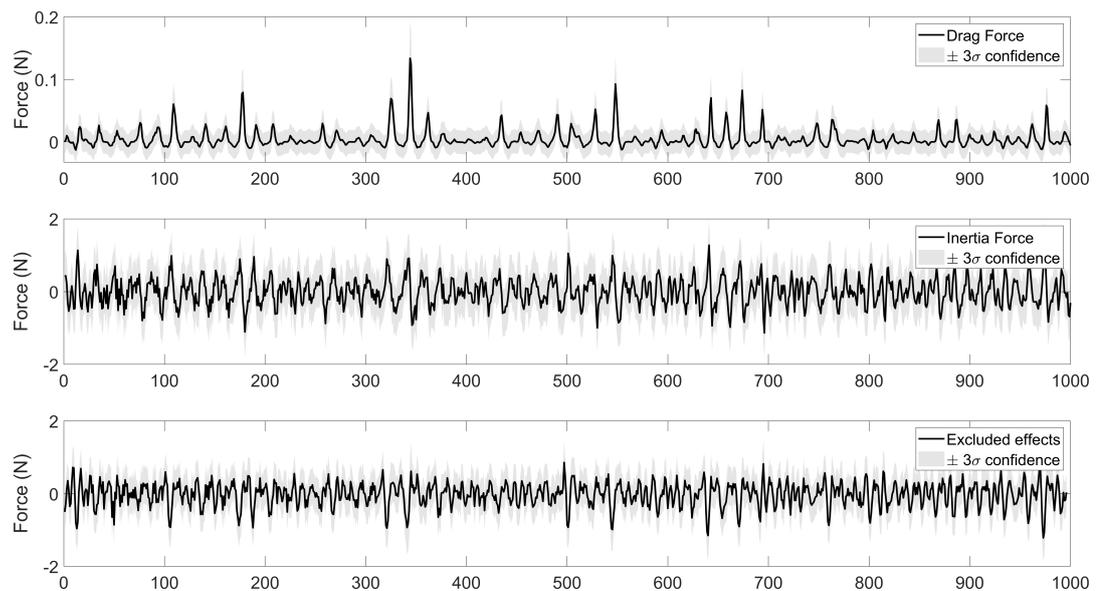


Figure 6.15: A breakdown of the relative drag, inertia and combined other effects within the wave load from a JONSWAP wave with $\gamma = 3.3$ and $\omega_p = 0.8Hz$. The prior C_d and C_m coefficients suggested by [234] were used to estimate drag and inertia forces, with excluded effects captured by a GP-NARX with a SE kernel.

6.4 Conclusions

Novel wave loading prediction models were developed that utilised only incoming wave height as a model input. They were implemented on an experimental dataset of a monopile structure within a wave tank across a variety of representative ocean state spectra. The data from a single wave gauge, performing an equivalent role to wave radars installed on offshore structures, was used to predict a wave load measured by a force collar.

The use of a GP-NARX, Morison's Equation and linear wave theory in combination within a physics-informed model was able to offer increased performance over any of its components used alone. The model was also able to offer a degree of interpretability, by providing a breakdown of the wave load in to drag, inertia and excluded effects. Model results were able to be visualised and interpreted easily, along with quantification of uncertainty.

Although the results of this work are promising, achieving wave height to wave force prediction errors of $\sim 15\%$, it is important to highlight that the models were developed using experimental data. There are many benefits of working within laboratory environments, particularly for research, including more tightly controlled conditions, increased repeatability of test states and an ability to install extensive sensor networks. However, it is important to note the difference between a wave tank and a real offshore environment; weather, salinity, wind and the nature of generated waves will all make the construction of models more challenging. Further testing and development of models, utilising data from real structures will be the natural progression of this work, with the validity of linear wave theory an expected important area for investigation. The progression to Stokes waves [276] or other similar extensions might prove a necessary inclusion.

PARTIAL KNOWLEDGE INCLUSION THROUGH KERNEL DESIGN

The models developed so far in the thesis have incorporated physical knowledge into Gaussian process mean functions and within the inputs to models. This chapter focuses on the incorporation of prior physical knowledge within the kernels of Gaussian processes, as introduced in section 2.2.4 of the literature review in Chapter 2. Modelling scenarios where one might have partial knowledge of a system or process are investigated through the combination of physics-informed kernels and more flexible ‘data-based’ kernels. Case studies of kernel structures are presented with application to wave loading prediction and a vibrating cantilever beam. The incorporation of physical knowledge in the kernel is rare within engineering generally, with very few examples currently in the field of SHM. Detailing the design process for the construction of physical kernels and how this might reflect a users physical insight should be a useful resource.

7.1 Kernel design

The covariance function (kernel) of a Gaussian process controls the behaviour of generated functions and completely defines a GP along with the mean function. Frequently used kernels such as the squared exponential and Matérn kernels correlate points closer together within the input space more highly. This is useful for a general model and can be highly flexible, however it does not embed any prior

physical knowledge about the system. It can be beneficial, in terms of performance, insight and trust within a model, to incorporate physical knowledge within the design of kernels e.g. the oscillatory behaviour of a spring mass damper system [149]. The design of the kernel may influence the prediction of the model to obey expected behaviours and constrictions.

Unlike mean functions, which may be any arbitrary function [139], kernels must satisfy several constraints:

1. The kernel must be symmetric.

$$k(X, X') = k(X', X) \quad (7.1)$$

2. The kernel must be positive semi-definite.

$$\sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j k(X_i, X_j) \geq 0 \quad \text{for all } n \in \mathbb{N}, X \in D, \alpha \in \mathbb{R} \quad (7.2)$$

3. The kernel must be measurable. This ensures the kernel represents a valid inner product within some feature space [277].

$$k(X, X') \in \mathcal{M}(X \times X', \mathbb{R}) \quad (7.3)$$

where $\mathcal{M}(X \times X', \mathbb{R})$ is the set of measurable functions acting on the input spaces X and X' .

4. The kernel must be non-infinite.

$$k(X, X') < \infty \quad \text{for all } X, X' \in X \times X' \quad (7.4)$$

The second constraint can be challenging to prove and poses the largest obstacle within the design of kernels. A common method to ensure this constraint is met for kernel designs is through the use of Bochner's theorem [278].

A positive definite, continuous, stationary function $k(\tau)$ may be expressed as the Fourier transform of a finite positive measure μ .

$$k(\tau) = \int_{-\infty}^{\infty} e^{-i\omega\tau} d\mu(\omega) \quad (7.5)$$

Provided a Fourier transform may be taken, the definition of finite positive measures within the frequency domain is a useful means of constructing kernels. It is often easier than attempting to prove positive semi-definiteness directly. The spectral mixture kernel [279] utilises Bochner’s theorem in its design. A symmetric sum of Gaussians is constructed and its Fourier transform taken to construct a valid kernel. The means and widths of the Gaussians determine the properties of the kernel, allowing for a high flexibility within kernel modelling capabilities.

There are a variety of means by which kernels may be constructed, with a summary provided in Figure 7.1. In general, these may be separated in to two categories: kernels that are constructed from scratch, and kernels that are constructed utilising other kernels in combination or as a starting point. A useful property of kernels is that linear operators between kernels produces a valid kernel as an output. This allows existing valid kernels to be combined to produce new kernels.

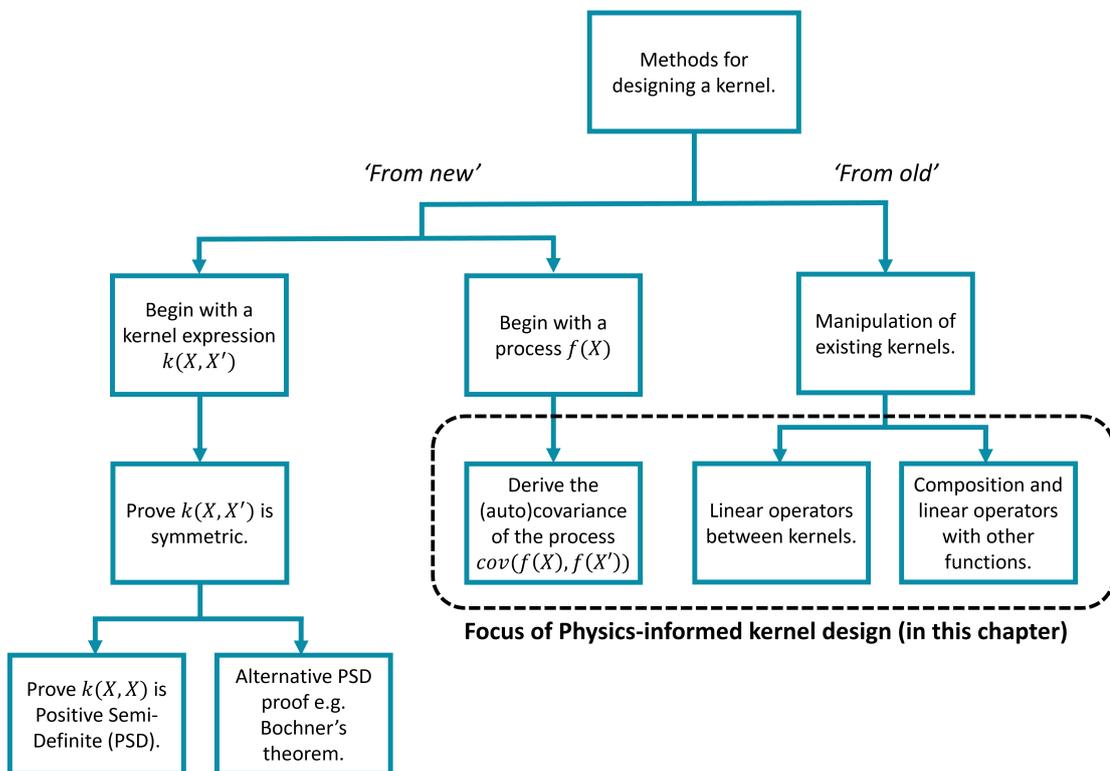


Figure 7.1: A summary of potential means by which to design a kernel for a Single Output Gaussian Process (SOGP). ‘From new’ refers to methods that begin from scratch, whilst ‘From old’ refers to methods that utilise existing kernels.

Combining kernels can be an effective means of accounting for partial understanding

of a process. A physically derived kernel, $K_{Phy}(X, X')$ may be used to capture some aspect of a process that is well understood, with a more flexible kernel, $K_{Data}(X, X')$ used to capture the remainder. For example, an additive kernel structure may aim to capture phenomena not represented within a physics based model:

$$K(X, X') = \underbrace{K_{Phy}(X, X')}_{\text{Understood process}} + \underbrace{K_{Data}(X, X')}_{\text{Excluded phenomena}} + \underbrace{\sigma_n^2 \delta_{ij}}_{\text{Noise}} \quad (7.6)$$

This model structure is presented in this chapter with application to wave loading prediction. It is revisited, along with a comparison to mean function based approaches, in Chapter 8. A multiplicative kernel structure may aim to incorporate partial understanding of a particular process:

$$K(X, X') = \underbrace{\overbrace{K_{Phy}(X, X') K_{Data}(X, X')}^{\text{Combined process}}}_{\substack{\text{Understood} \\ \text{aspect}}} + \underbrace{\sigma_n^2 \delta_{ij}}_{\text{Noise}} \quad (7.7)$$

This chapter will introduce a physics-informed kernel of this form applied to a vibrating cantilever beam, showing how kernels may be derived for a physical process and combined with flexible kernels. This is used to reflect instances where knowledge about a particular aspect of a process is missing. The chosen example is relevant to a wide variety of engineering applications and the kernel structure maps nicely onto how one would typically structure a physical model of the system.

7.2 Morison's Equation in kernel form

Within a GP, the kernel defines the family of functions from which predictive samples may be drawn. Through the design and selection of kernels, one may control and restrict the behaviour of predictions generated from a GP and enforce desirable or physically derived constraints. It is often the case that one may wish to design a kernel to capture the behaviours of some physical process $f(x)$. If a functional form for $f(x)$ is known, this is possible through calculation of the (auto)covariance of the process:

$$\text{cov}(f(X_i), f(X_j)) = \mathbb{E}[f(X_i)f(X_j)] - \mathbb{E}[f(X_i)]\mathbb{E}[f(X_j)] \quad (7.8)$$

Earlier work in this thesis investigated the use of Morison's equation for the prediction of wave loads on offshore structures [32, 100]. Morison's equation relies on a number of simplifying assumptions and does not account for phenomena such as vortex shedding and turbulence [45, 256]. In Chapter 5, it was used within the mean function, with the learning of the excluded processes the role of the data-based GP-NARX. This improved predictive performance, particularly in instances of reduced training data coverage. Here, a kernel is derived to represent Morison's Equation.

A GP with a Morison's kernel will have equivalent modelling capabilities to a Bayesian Linear Regression of Morison's equation, except that the model training time is now of order $O(n^3)$ rather than $O(n)$. Used alone, an implementation of this model would be needlessly inefficient; the advantage of kernel representation here comes from the ability to combine with other kernels. In many engineering applications, there is some aspect of a process that is not well understood and is not accounted for within a physics-based model. A useful property of kernels for varying the inclusion of prior knowledge within a GP is the ability to be combined, through addition, multiplication and composition, with other kernels. The use of physically derived kernels, K_{Phy} , in combination with flexible, more generally applicable, kernels, such as a Squared Exponential, K_{Data} , allows for the creation of model structures where K_{Phy} aims to encode some aspect of prior knowledge and K_{Data} captures unknowns.

The derivation of the Morison kernel begins with the expression for Morison's Equation first presented in Chapter 3. For a stationary, rigid, slender, cylinder of diameter D positioned within waves of velocity U and acceleration \dot{U} , the force per unit axial length F by Morison's Equation is expressed:

$$F = \underbrace{\frac{1}{2}\rho DC_d U|U|}_{C'_d} + \underbrace{\frac{1}{4}\pi\rho D^2 C_m \dot{U}}_{C'_m} \quad (7.9)$$

where ρ is the fluid density, C_d is the drag coefficient and C_m is the inertia coefficient. The dimension specific terms may be grouped to form two constants C'_d and C'_m relating to the drag and inertia forces of the wave. This allows the structuring of

Morison's Equation as a Bayesian Linear regression.

$$X = [U|U|, \dot{U}] \quad (7.10)$$

$$\beta = [C'_d, C'_m]^T \quad (7.11)$$

$$p(F|X, \beta, \sigma_n^2) \sim \mathcal{N}(X\beta, \sigma_n^2\mathbb{I}) \quad (7.12)$$

The covariance of a process $f(X) = X\beta$ with prior $\beta \sim \mathcal{N}(0, \sigma_\beta^2\mathbb{I})$ between two input vectors X_i and X_j is derived

$$\begin{aligned} \text{cov}(f(X_i), f(X_j)) &= \mathbb{E}[f(X_i)f(X_j)] - \mathbb{E}[f(X_i)]\mathbb{E}[f(X_j)] \\ &= \mathbb{E}[(X_i\beta)(X_j\beta)^T] - \mathbb{E}[X_i\beta]\mathbb{E}[X_j\beta] \quad (7.13) \\ &= X_i\mathbb{E}[\beta\beta^T]X_j^T \\ &= X_i\Sigma_\beta X_j^T \end{aligned}$$

The expression for covariance can then be used as a kernel within a physically-informed GP of the form

$$f(X) \sim \mathcal{GP}(0, K_{Phy}(X, X')) \quad (7.14)$$

where $K_{Phy}(X, X') = X\Sigma_\beta X'^T + \sigma_n^2\delta_{X, X'}$ and $X = [U|U|, \dot{U}]$. If one considers the target to be the sum of $f(x)$ and some function $g(x)$ to account for unmodeled phenomena, then it follows that the covariance structure is

$$\begin{aligned} K(X, X') &= K_{Phy}(X_1, X'_1) + K_{Data}(X_2, X'_2) + \sigma_n^2\delta_{X, X'} \\ &= \underbrace{X_1\Sigma_\beta X_1'^T}_{\text{Morison's Equation}} \\ &\quad + \underbrace{\sigma_f^2 \exp\left(-\frac{1}{2}(X_2 - X'_2)\Lambda^{-1}(X_2 - X'_2)^T\right)}_{\text{Excluded phenomena}} \quad (7.15) \\ &\quad + \underbrace{\sigma_n^2\delta_{X, X'}}_{\text{Noise}} \end{aligned}$$

where a Squared Exponential (SE) has been used as the covariance of the unknown

behaviour under a zero mean assumption. Here $X_1 = [U|U], \dot{U}]$, $X_2 = [U_t, \dot{U}_t, U_{t-1}, \dot{U}_{t-1}, \hat{y}_{t-1}, \hat{y}_{t-2}, \hat{y}_{t-3}]$, σ_f^2 is the signal variance, Λ is the matrix of length scales such that $diag(\Lambda) = [l_1^2, l_2^2, \dots, l_D^2]$ for a D dimensional input and σ_n^2 is the noise variance. The selected maximum lags, $l_u = 1$ and $l_y = 3$, have been carried forward from the work in Chapter 5.

7.2.1 Model Predictive Performance

Models were implemented on the Christchurch Bay Tower (CBT) [230] dataset, detailed in Chapters 4 and 5. All models were trained using the first 500 datapoints of the training and validation sets highlighted in Figure 4.2. For the GP-NARX models, this involved the optimisation of MPO performance on the validation set (see Algorithm 1 of Chapter 5). To summarise the model structures being tested:

- Morison's Equation, structured as a Bayesian Linear Regression.

$$X_1\beta \quad (7.16)$$

where $X_1 = [U|U], \dot{U}]$ is an input space of flow conditions and $\beta = [C'_d, C'_m]^T$ are the combined drag and inertia coefficients.

- A black-box, zero-mean GP-NARX with a Squared Exponential kernel.

$$\mathcal{GP}(0, K_{SE}(X_2, X'_2)) \quad (7.17)$$

where $X_2 = [U_t, \dot{U}_t, U_{t-1}, \dot{U}_{t-1}, \hat{y}_{t-1}, \hat{y}_{t-2}, \hat{y}_{t-3}]$ is the lagged input space of flow conditions and previous predictions of the wave load.

- A physics-informed GP-NARX, with Morison's Equation in the mean function.

$$\mathcal{GP}(X_1\beta, K_{SE}(X_2, X'_2)) \quad (7.18)$$

- A physics-informed GP-NARX, using the kernel defined in Equation 7.15.

$$\mathcal{GP}(0, K_{Lin}(X_1, X'_1) + K_{SE}(X_2, X'_2)) \quad (7.19)$$

The performance of models was measured using two metrics: the Normalised Mean Square Error (NMSE) and the Mean Standardised Log Loss (MSLL). The MSLL is a probabilistic measure, with superior models having more negative scores. A baseline of zero is equivalent to setting the predictive mean and variance for all test

set points as the mean and variance of the training set. A comparison of results is shown in Table 7.1.

Table 7.1: Model performance comparison.

Model	Structure	NMSE (%)	MSLL
Morison’s Eq.	$X_1\beta$	19.528	−0.813
	$X_1\beta$ (<i>Normalised</i>)	15.458	−0.857
GP-NARX	$\mathcal{GP}(0, K_{SE})$	19.383	0.450
	$\mathcal{GP}(X_1\beta, K_{SE})$	14.797	−0.598
MC MPO	$\mathcal{GP}(0, K_{Lin} + K_{SE})$	14.913	−0.939

A promising trend observed within Table 7.1 is that models using physics and data in combination outperformed both Morison’s Equation and black-box data-based models. The residual modelling GP-NARX and combined kernel GP-NARX were the two best performing models, with similar NMSE scores and an improved MSLL on the combined kernel model. A potential reason for the worse MSLL of the residual model is that attempting to fit a GP to a residual rather than the complete signal greatly reduces the signal to noise ratio. This can make it challenging to pick out remaining structure within the signal, favouring long lengthscale, high noise variance model fits.

Inline with findings within the literature [46, 223–225], Morison’s Equation was found to perform very well considering its simplicity and outperformed the black-box GP-NARX. This was in part helped by the dataset, which was from a region of primarily unidirectional flow, a key assumption of Morison’s Equation. The training, validation and test sets were all from drag-dominated flow regimes, meaning that parameters learned within the construction of the model were likely to be appropriate for the test conditions. When using Morison’s Equation, it is important to consider the likely flow regime in which it may be implemented and adjust C_d and C_m accordingly.

Although achieving a moderate NMSE score, the black-box GP-NARX performed very poorly with regards to MSLL. This highlighted two issues: the importance of considering multiple performance metrics and unexpected behaviours that may arise from the use of black-box models, particularly highly flexible ones such as a GP-NARX. The cause of a poor MSLL here was the explosion of uncertainty intervals due to feedback of samples within the MC MPO prediction, a common phenomena within autoregressive model structures.

7.2.2 Posterior contribution breakdown

Along with benefits in overall performance, an advantage of combining physics with data-driven techniques is the insight that can be provided by looking at the contribution of each kernel component. Whilst the separation of a mean function and residual model is already clear within the combined model structure, the individual kernel posterior contributions within a combined kernel require decomposition. Following [156] it is possible to derive the conditional predictive distribution for the contribution of a kernel K_i within a combined additive kernel of the form $K = \sum_{i=1}^{i=n} K_i$.

$$p(f_i^* | f^*, X^*, f, X, \theta) \sim \mathcal{N}(K_i^{*T} (\sum_{i=1}^{i=n} K_i) y, K_i^{**} - K_i^{*T} (\sum_{i=1}^{i=n} K_i) K_i^*) \quad (7.20)$$

where f_i^* is the prediction contribution of kernel K_i within the combined prediction $f^* = \sum_{i=1}^{i=n} f_i^*$. The breakdown of the mean function and residual GP contributions is shown in Figure 7.2, with the breakdown of kernel component contributions of the combined kernel model shown in Figure 7.3.

In both cases, Morison's Equation is able to capture the majority of structure within the wave force via either the mean function or linear kernel component. Here, this provides interpretability in to the flow conditions, highlighting the presence of a drag dominated flow regime and primarily unidirectional flow. For alternative flow conditions, particularly those breaking the assumptions of Morison's Equation, the relative contributions of model components would likely be very different.

Due to the good fit of the mean function, the signal to noise ratio of the residual is very poor and the residual fit GP struggles to pick out the remaining structure. The middle plot of Figure 7.2 shows a long lengthscale fit with a large estimated noise. This contributes towards the larger variance and poorer MSL of the combined residual model.

The consideration of individual component contributions within a model is of particular importance during extrapolation. Far from observed data, the data-based component of the model (or stationary kernel) will revert to its zero prior, with the overall model therefore dependant on the quality of either the mean function or physics-derived (nonstationary) kernel component. Independent of the overall model performance for a given testing dataset, ensuring that model

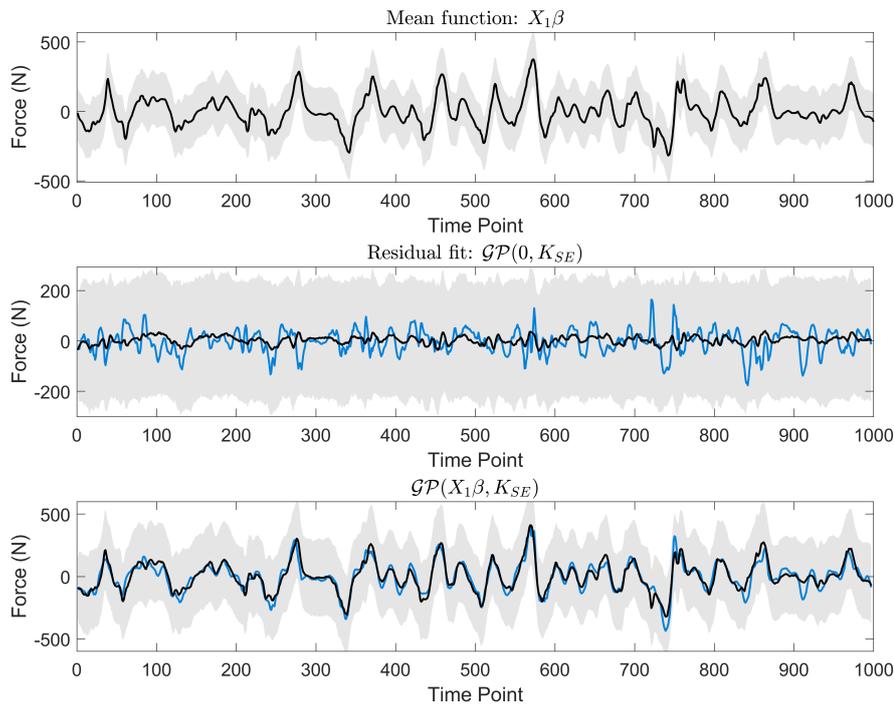


Figure 7.2: Residual modelling contribution breakdown of the linear mean function (top), residual GP-NARX with SE kernel (middle) and combined model (Bottom).

components are capturing their intended processes is therefore important. The ability of kernel components to be learned simultaneously will help to reduce the ‘biasing’ phenomena of mean functions. Here, the GP with the derived-kernels has captured each component very well, indicating that it is likely to perform well in extrapolation.

7.2.3 Conclusions

The use of Morison's Equation and a GP-NARX in combination was found to increase predictive performance over either technique used independently. Residual modelling was able to improve NMSE at the expense of MSL, whilst the combined kernel was able to offer advantages in both NMSE and MSL.

Inline with the literature, Morison's Equation was found to achieve satisfactory

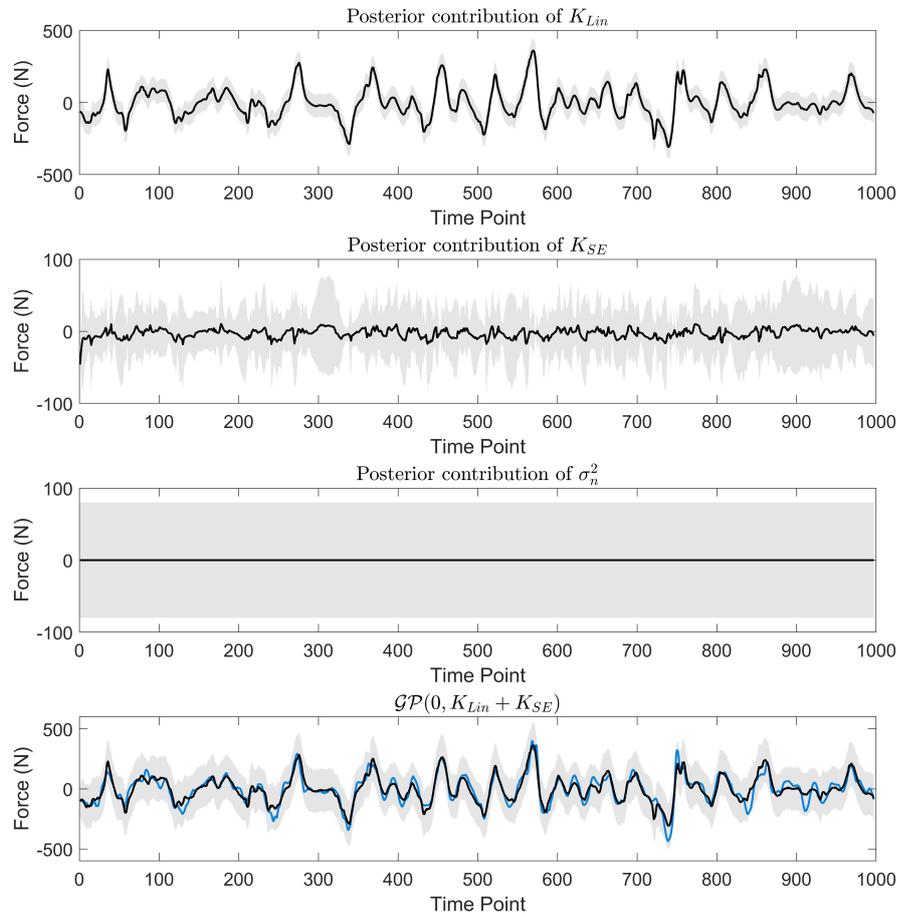


Figure 7.3: Combined kernel model contribution breakdown of the linear (Morison) kernel (top), SE GP-NARX (second), noise variance (third) and combined model (bottom).

performance when used in an appropriate flow regime for little computational cost. It provided a sensible start point for the development of combined physics and data-based models.

The decomposition of combined models was found to offer physical insight into the role of each model component. Here, due to the flow conditions, the majority of structure was captured via the physics-based component. Further work investigating how model component contributions vary over a range of flow conditions is planned to further explore how data-based learning may assist prediction in conditions

where the performance of Morison's Equation would typically suffer.

This case study investigated a case of an additive partial knowledge structure presented in Equation 7.6, applied to the task of wave loading prediction. The second case study of the chapter will now present a more broadly applicable engineering example of a cantilever beam in free vibration. The example also takes the multiplicative structure present within Equation 7.7, with the aim of presenting a range of potential kernel structures to the reader.

7.3 Physical kernels for the free vibration of a cantilever beam

To demonstrate the proposed method of handling partial knowledge within the GPR framework, a simulated case study of a cantilever beam is employed. This case study is useful for a number of reasons; a beam assumption is applicable to a wide range of engineering applications from the modelling of offshore structures to aeroplane wings, and importantly, it provides us with an opportunity to explore the embedding of various levels of physical insight. Where a structure of interest is well modelled by a cantilever beam, there are full analytical solutions available. If a structure is cantilever but not of uniform cross section or homogeneous material, knowledge of the mode shapes may no longer be assumed (this is the example pursued here). Finally, one may face a situation where only confident about boundary conditions, in which case employing constraints within the GPR may be appropriate, for example [197, 201]. The work in this section of the chapter presents the first use of the derived SDOF kernel [149] within a larger MDOF kernel structure. It highlights how along with incorporating physics within individual kernel components, the structure of how multiple kernels are combined can itself mimic a physical model or process.

The dataset used within the study was from a simulated cantilever beam as shown in Figure 7.4, where x is distance along the beam from the fixed end.

The beam was modelled as a steel cylinder, 1m in length, with a radius of 0.01m. The modal amplitudes A_i were assumed to decay logarithmically; this was to impart more energy into the lower modes of the beam [244, 280]. The parameters of the simulation are given in Table 7.2.

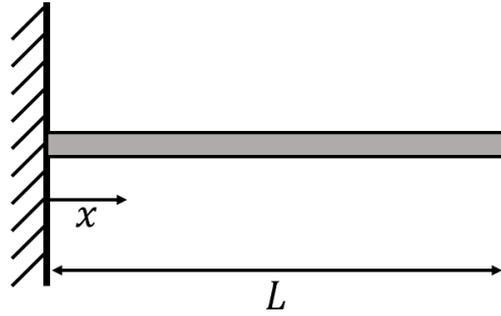
Figure 7.4: Cantilever beam of length L .

Table 7.2: Parameters of the cantilever beam simulation.

Quantity	Value
Length, L	1 m
Radius, r	0.01 m
Young's Modulus, E	190 GPa
Density, ρ	8050 kgm ⁻³
Damping ratio, ζ	0.005
Simulation frequency, f_s	8192
Number of modes, N	4
Modal amplitudes, A_i	$10^{\frac{0-(i-1)(N-1)}{N-1}}$ for $i = 1, \dots, N$

The free vibration response of a beam can be found through separation of variables with a solution of the form $Y(x, t) = W(x)T(t)$. The response of a beam in free vibration may be calculated through the superposition of normal modes [244].

$$Y(x, t) = \sum_{i=1}^{\infty} W_i(x)T_i(t) = \sum_{i=1}^{\infty} W_i(x)e^{-\zeta\omega_n^i t}(A_i \cos(\omega_n^i t) + B_i \sin(\omega_n^i t)) \quad (7.21)$$

where A_i and B_i are determined from the initial conditions of the beam. For many sets of known boundary constraints it is possible to derive an exact analytical expression for the mode shapes $W_i(x)$. Following Blevins [243], for a fixed-free beam of uniform cross section in free vibration, this gives

$$W_i(x) = \cos(\beta_i x) - \cosh(\beta_i x) - \frac{\cos(\beta_i L) + \cosh(\beta_i L)}{\sin(\beta_i L) + \sinh(\beta_i L)} (\sin(\beta_i x) - \sinh(\beta_i x)) \quad (7.22)$$

where L is the length of the beam and β_i is a constant specific to the boundary conditions and mode. The analytical solutions of Blevins [243] were used to create a simulated dataset of a fixed-free beam in free vibration. The first four modes of the beam were investigated, with plots of the mode shapes, $W_i(x)$, and oscillatory behaviour, $T_i(t)$, shown in Figure 7.5.

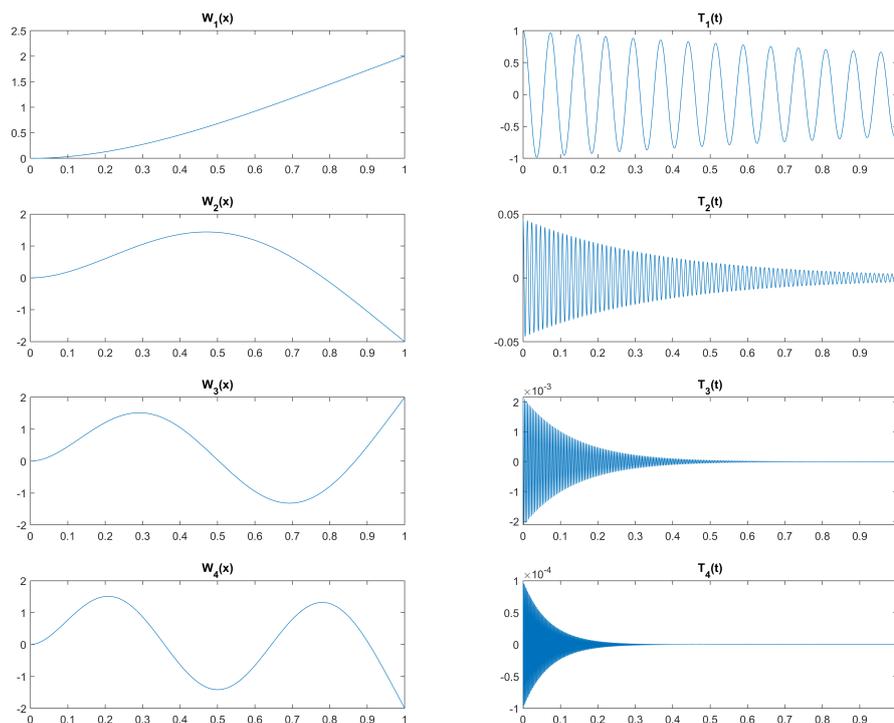


Figure 7.5: $W_i(x)$ and $T_i(t)$ plots of the first four modes of a simulated cantilever beam in free vibration.

It is worth noting that the main focus of this case study is not to minimise prediction error on a test set, as for the cantilever beam simulated this could be achieved through an analytical solution. The aim here is to investigate how to vary the level of physical knowledge incorporated within a machine learning model and discuss the effects this will have. At one extreme would lie “complete knowledge”, i.e. the analytical solutions of Blevins [243] and Rao [244], whilst a purely data-based approach would represent zero prior knowledge.

7.3.1 Relaxing assumptions through kernel design

The construction of physics-based (white-box) models requires the approximation of phenomena and a reliance on assumptions and simplifications remaining representative. In the case of a cantilever beam these may include a perfectly fixed beam, a uniform beam cross section or homogenous material properties. When assumptions hold, physics-based models often provide efficient, accurate solutions to a range of engineering problems. However, the usage of physics based models when their assumptions breakdown can lead to significant increases in error. For a cantilever beam, examples of this might range from a non-uniform cross section of a turbine blade or aircraft wing, marine growth along the length of an offshore monopile or a soil foundation being approximated as fixed. The relaxing or removal of assumptions within a physics-based model will therefore allow for the widening of conditions in which the model may be used. Examples of restricted modelling cases are summarised in Table 7.3.

Table 7.3: Summary of assumptions and modelling restrictions within the analytical expressions of Blevins [243] and Rao [244] for a fixed-free beam in free vibration.

Modelling assumption	Restrictive modelling case
Uniform cross section.	Non-uniform cross section of a turbine blade or aircraft wing.
Perfectly fixed base.	Deflection of structures within soil or sea bed foundations [281].
Homogeneous material properties.	Composite materials with directional material properties.
Equal mass distribution along the length of the beam.	Marine growth adding mass along the length of offshore monopiles [282].

Many key restrictions of the physics-based model of the cantilever beam come from the analytical expression for the mode shape, $W_i(x)$, which is specific to the boundary conditions of the beam. This work suggests a means of removing these restrictions, accounting for partial knowledge by combining different covariance functions, each to capture different components of the combined target function. For a beam where the mode shapes are assumed to be unknown, the displacement will be a product of oscillatory behaviour and the unknown modes, $Y(x, t) = W(x)T(t)$ where $W(x)$ is unknown. The covariance of Y in this case is equal to

$cov(Y, Y') = K_W(x, x')K_T(t, t')$. Here the SE kernel is used to model the covariance, K_W , of the unknown mode shapes:

$$K_{SE}(\tau) = \sigma_f^2 \exp\left(-\frac{1}{2}(\tau)^T \Lambda^{-1}(\tau)\right) + \sigma_n^2 \delta_{ij} \quad (7.23)$$

where σ_f^2 is the signal variance, σ_n^2 is the noise variance, Λ is the matrix of length scales such that $diag(\Lambda) = [l_1^2, l_2^2, \dots, l_D^2]$ for a D dimensional input and $\tau = x - x'$ is the distance between a pair of input points x and x' . To capture the covariance of the oscillatory behaviour, K_T , the SDOF kernel [149] is employed:

$$K_{SDOF}(\tau) = \frac{\sigma^2}{4m^2\zeta\omega_n^3} e^{-\zeta\omega_n|\tau|} \left(\cos(\omega_d\tau) + \frac{\zeta\omega_n}{\omega_d} \sin(\omega_d|\tau|) \right) \quad (7.24)$$

where the hyperparameters of the kernel now relate to physical properties of a SDOF oscillator: m is the mass, $\zeta = c/2\sqrt{km}$ is the damping ratio, $\omega_n = \sqrt{k/m}$ is the natural frequency and $\omega_d = \omega_n\sqrt{1-\zeta^2}$ is the damped natural frequency. Draws from this kernel are constrained to obey the behaviour of a decaying SDOF oscillator, a useful property to encode.

For the prediction of the full beam response, the SE and SDOF kernels are multiplied together to predict the response contribution of a single mode i . These kernel products may then be summed for a given number of modes N to predict the combined response of the beam. This will lead to a combined kernel structure of

$$K(\tau_x, \tau_t) = \sum_{i=1}^N \underbrace{K_{SE}^i(\tau_x)}_{\text{Mode shape}} \underbrace{K_{SDOF}^i(\tau_t)}_{\text{Oscillatory behaviour}} + \underbrace{\sigma_n^2 \delta_{ij}}_{\text{Noise}} \quad (7.25)$$

where $\tau_x = x - x'$ and $\tau_t = t - t'$ represent distances between points in the spatial and temporal inputs respectively. The combined model has the flexibility to recover $W_i(x)$ and $T_i(t)$ with an important distinction; the analytical form of the cantilever mode shape has not been fixed and the corresponding assumptions used to construct it have been relaxed.

7.3.2 Recovery of mode shapes

The recovery of the mode shapes $W_i(x)$ and temporal functions $T_i(t)$ can be useful to estimate stress in difficult-to-reach areas of a structure [283]. Although the simulated beam will obey the analytical expressions for $W_i(x)$, this constraint was removed from the model structure via the introduction of the SE kernel. As such, the capability of the model to recover the mode shapes is a useful property to measure. The use of a simulation is useful for this case as it allows the exact mode shapes used to construct the response to be compared with the predictions and performance levels to be measured. Eight evenly spaced data points were selected along the length of the beam as training points, with the response simulated at $8192Hz$. Model performance was measured on an unseen test set of 100 points along the beams length simulated for one second.

To extract the mode shapes from the combined beam response prediction, the posterior contributions for each mode i within the kernel sum will require decomposition. This was achieved through the use of Gaussian conditionals, following [156], similarly to the decomposition of Morison kernel components presented in section 7.2.2. The predictive performance for each mode was measured using two metrics: the Normalised Mean Square Error (NMSE) and the Mean Standardised Log Loss (MSLL). The MSLL is a probabilistic measure, with superior models having more negative scores. A breakdown of predictive performance is shown in Table 7.4 and plots of the recovered $W_{1:4}(x)$ and $T_{1:4}(t)$ are shown in Figure 7.6.

Table 7.4: Modal performance breakdown of beam response prediction.

Model target	Function	NMSE (%)	MSLL
First mode	$W_1(x)$	0.002	-5.472
	$T_1(t)$	0.411	-3.088
Second mode	$W_2(x)$	0.225	-3.520
	$T_2(t)$	0.285	-2.677
Third mode	$W_3(x)$	9.949	-1.118
	$T_3(t)$	0.240	-2.055
Fourth mode	$W_4(x)$	1.221	-2.522
	$T_4(t)$	0.541	-1.969

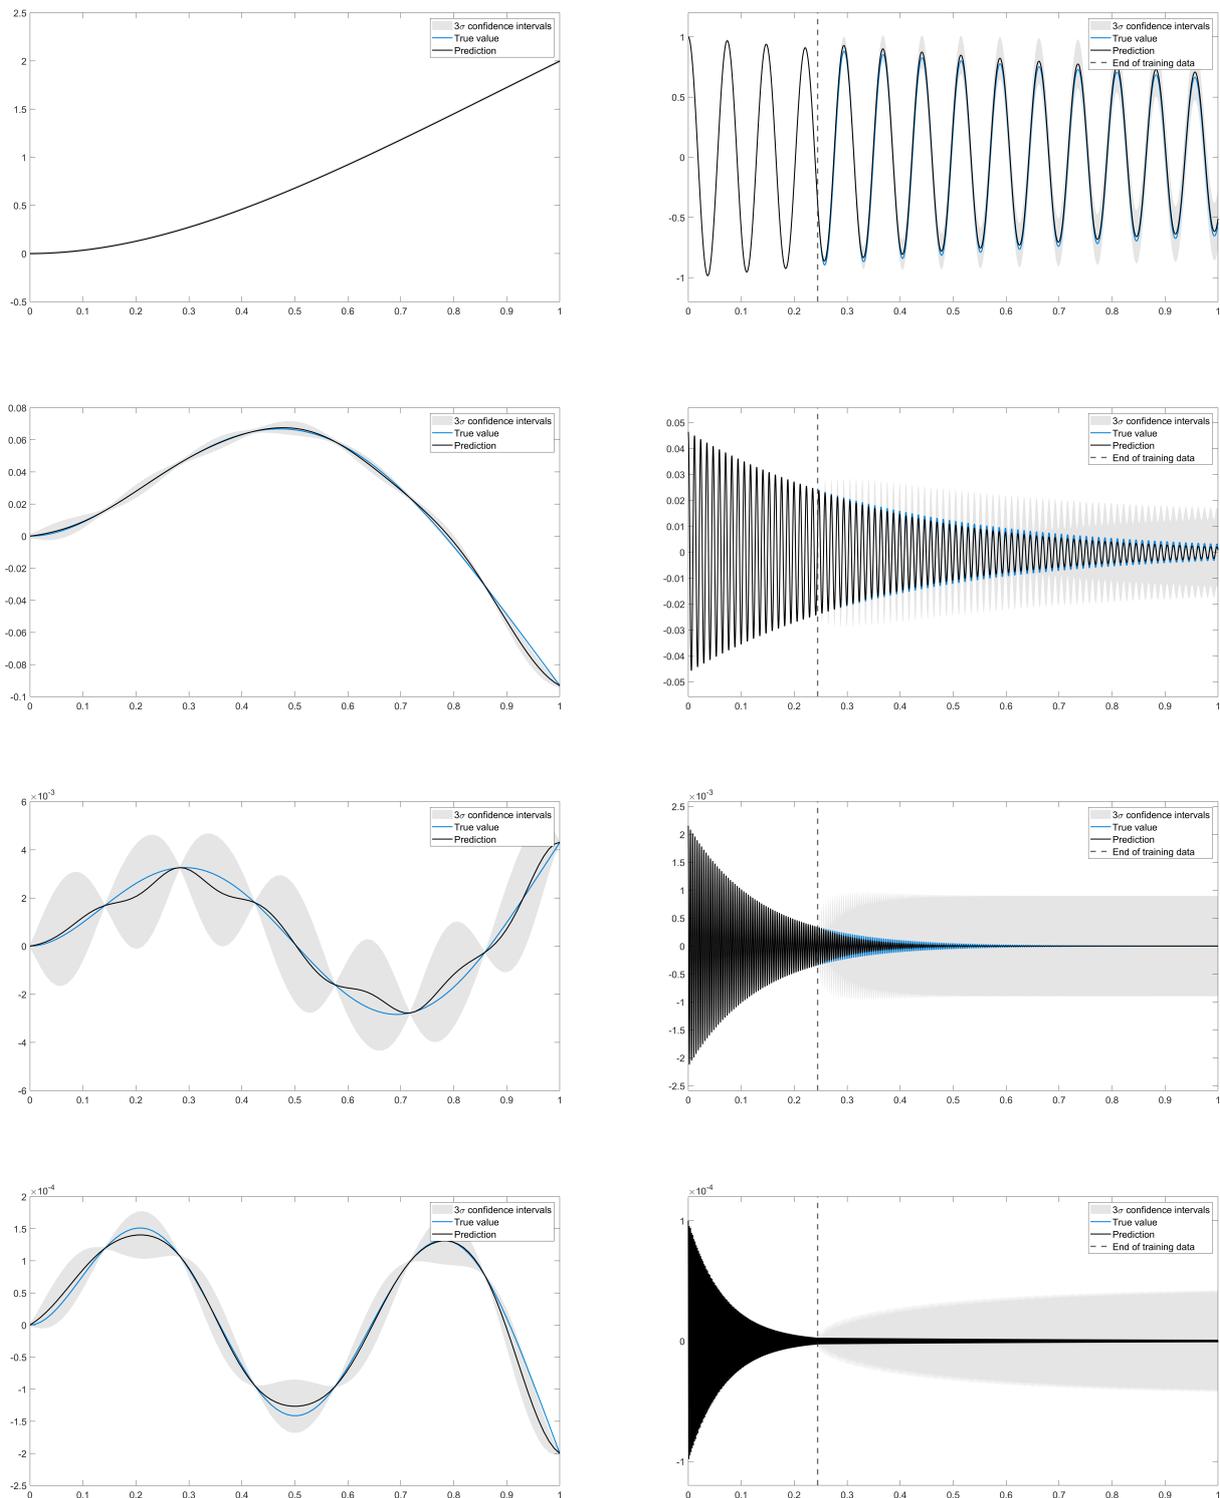


Figure 7.6: Breakdown of recovered $W_{1:4}(x)$ (left column) and $T_{1:4}(t)$ (right column) from a cantilever beam in free vibration by a combined kernel of the form $K(\tau_x, \tau_t) = \sum_{i=1}^4 K_{SE}^i(\tau_x)K_{SDOF}^i(\tau_t) + \sigma_n^2\delta_{ij}$. All models observed eight evenly spaced points along the beams length at time points $[1:2:2000]$ as training data.

A key trend observed within Table 7.4 is that performance generally worsens for the higher modes of the beam. This is to be expected for a number of reasons; firstly, the mode shapes are more complex for the higher modes, meaning that the same eight spatial points measured along the beams length must be used to represent additional, sharper changes in direction. The magnitude of vibration of the higher modes was also significantly lower, making identifying a modes' contribution from the full response more challenging. A useful property of the SDOF kernel for identifying modal contributions was having physically interpretable hyperparameters, specifically the natural frequency ω_n . The first four natural frequencies of the beam were able to be identified and the corresponding hyperparameters fixed within the kernel. This reduced the number of required hyperparameters to optimise whilst encouraging the model to learn the response contributions at specific frequencies. Physically interpretable hyperparameters provide an avenue through which knowledge may flow in and out of a system. If parameters are learnt instead of fixed, optimisation can provide a means of retrieving this information from the system.

Within the plots of $W_{1:4}(x)$, the ability of the model to recover the mode shapes is observed. A useful property of the GP here is the quantification of uncertainty within predictions. For the first mode, where an excellent fit is achieved, a narrow confidence interval indicates the model is certain in its prediction. For $W_3(x)$, behaviour indicative of an underestimation of the lengthscale within the SE kernel is observed. The influence of observed points decays very quickly with distance and the model attempts to revert to its zero prior between the training points along the length of the beam. The confidence intervals of the GP expand quickly within these areas to reflect this however, preventing an overconfident, incorrect prediction.

One of the major advantages of including physical knowledge within a model is improving the ability to extrapolate. This can be seen in the $T_{1:4}(t)$ plots within Figure 7.6, where the model continues to predict beyond the end of observed training data (Dotted line). The behaviour of a decaying SDOF oscillator is encoded within the design of the SDOF kernel and has shown to be useful when conditioning models on a reduced quantity of data [149]. An important consideration when using prior knowledge to extrapolate however is any assumptions present within the construction of the model.

7.4 The balance between physics and data

Relating back the proposed model structure to Figure 1.1 and the early discussions of Chapter 1, it can be useful to consider how this changes the models relative position on the scale of prior knowledge inclusion. By relaxing the assumptions made within the physics-based model, the reliance on data has increased. If unsure about whether assumptions of a physics-based model may hold, e.g. the examples presented within Table 7.3, the inclusion of them has the potential to cause overly confident, erroneous model predictions. Were these assumptions maintained, for example via the inclusion of a mean function, a higher reliance on the physics-based component would be achieved, thereby moving the model's position on the scale.

The balance between physics and data within a model is an important consideration within the construction of physics-informed machine learning models. Figure 7.7 highlights the concept of an optimum level of prior knowledge inclusion for a specific modelling scenario. A modelling scenario here is categorised by its relative levels of available knowledge and data. Ideally, one would wish for both high levels of data and prior knowledge, for example, a heavily sensed structure in well defined laboratory conditions; however in real life industrial applications this would be rarely achievable. In circumstances when neither knowledge or data are abundant, the predictions of models may often be unreliable or erroneous.

The concept of an “optimum” level of prior knowledge within a model should be considered when developing models along with the consequences of incorrect placement on the scale. The over inclusion of prior knowledge within a model can lead to a higher reliance on the physics-based component of the model which may not accurately represent the modelled system or phenomena. This can lead to reduced model flexibility and overly confident models. An under inclusion of prior knowledge may waste a key resource, useful for increasing model interpretability and capabilities in extrapolation. Models with reduced prior knowledge inclusion will rely more heavily on available data, increasing the demand for expensive data collection.

The work done so far has highlighted potential model structures for the relaxing of assumptions imposed on the vibration of a cantilever beam, however the initial testing of the model has focussed on a simulated dataset, representing an idealised case. The aim for future work is to implement and test the methods on a measured dataset, for which the assumptions made in the analytical formulas of Blevins [243] and Rao [244] would be less valid.

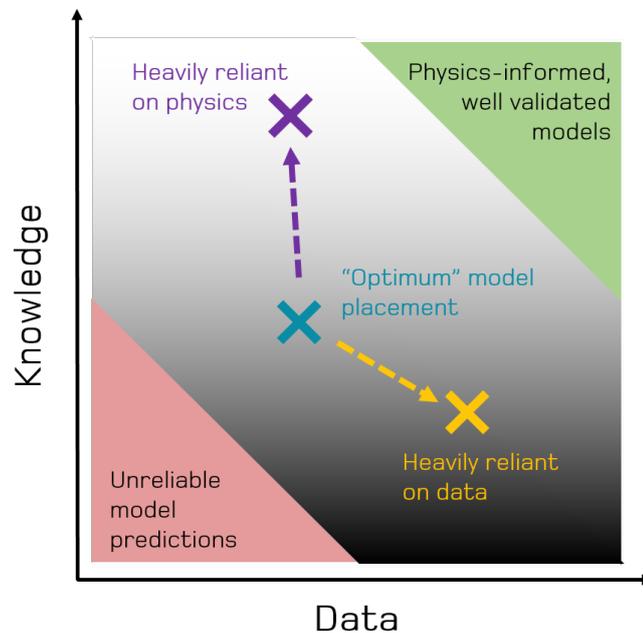


Figure 7.7: The concept of an optimum level of prior knowledge inclusion within a physics-informed machine learning model for a given modelling scenario. Examples of incorrectly estimating a models placement on the scale, leading to an heavy reliance on physics or data are shown.

7.5 Conclusions

The idea of an optimum level of prior knowledge to be included within a model changing depending on the modelling scenario was highlighted within the context of a cantilever beam in free vibration. The benefits of achieving this optimum level such as improved extrapolation, interpretability and reduced reliance on data collection were discussed. Consequences of being incorrectly placed on the scale of prior knowledge inclusion were also presented. These included reduced flexibility and potential reliance on unrealistic assumptions for a physics-heavy model and poor extrapolation, interpretability and increased demand for data collection for a data-heavy model.

The model developed in the second half of this chapter explored a scenario where only partial prior knowledge was available and it was assumed the assumptions used to construct physics based models could not be guaranteed to hold. These assumptions were relaxed through the use of kernel design within a GP. The mode

shapes of the beam, for which the analytical solutions were removed from the model as to relieve specific boundary conditions, were able to be recovered for the first four modes of the beam.

This chapter as a whole outlined the design process for how one might incorporate physical knowledge within the kernels of a Gaussian process and highlighted important links between how the nature of prior knowledge possessed by a user changes the optimum kernel structure. Two case studies were presented in the chapter: a kernel representation of Morison's Equation, the first example of physically-informed kernel for wave loading prediction, and an MDOF kernel structure capturing the free vibration response of a cantilever beam. The beam kernel presented the first utilisation of the derived SDOF kernel within a MDOF scenario and the first to extract the mode shapes of a vibrating system utilising a posterior kernel contribution breakdown.

EXPLORING THE EFFECTS OF PRIOR KNOWLEDGE INCLUSION ON GAUSSIAN PROCESS MODEL STRUCTURE

Deciding on the most appropriate means of integrating prior knowledge within a GP can be challenging, with the capabilities of models not always clear before they are implemented. An important step to aid with making this decision is understanding how the means of incorporating prior knowledge within a GP effects the properties, structure and performance of a constructed model. This chapter aims to directly compare differences in models with modifications to their mean and covariance functions, through the use of a case study utilising a consistent form of prior knowledge. An example is considered where a linear function represents some known aspect of a system's behaviour, with the remainder captured by a flexible kernel; this allows the maths for all examples to be worked through completely, whilst highlighting differences more clearly. A toy dataset is then used to compare the performance of models in circumstances of limited training data, an important area of focus for grey-box models. The work of this chapter is early in its field and provides an important platform from which to construct the decision framework presented later in Chapter 9.

8.1 Prior knowledge through linear models

To allow for the direct comparison of model structures, it is important that the form of prior knowledge integrated within models remains consistent. Here, linear prior knowledge is used to represent an approximation of a known behaviour across all models. The use of linear models is common across many applications within engineering, some examples of which are shown in Table 8.1. The applicability of linear prior knowledge to a range of fields, along with an ability to more easily show differences in mathematical derivations made it an effective case study example.

Table 8.1: Examples of linear models within engineering applications.

Model	$f(X)$	Description
Hooke's Law	$F = xk$	Force F required to deflect a material by some distance x .
Charles' Law	$V = Tk$	Volume V and temperature T relationship of a gas at constant pressure.
Archard Wear Equation	$Q = \frac{W}{H}K$	The volume of wear Q produced by a load W sliding a distance L over a material of hardness H .
Morison's Equation	$F = U U C'_d + \dot{U}C'_m$	Wave force F acting on a slender cylinder given wave velocity U and acceleration \dot{U} .

A model is deemed linear (or *linear-in-the-parameters*) if the output $f(X)$ may be represented as some combination of the input space X multiplied by a set of constants β :

$$f(X) = \beta_0 + X_1\beta_1 + X_2\beta_2 + \dots + X_n\beta_n \quad (8.1)$$

expressed in matrix form as

$$f(X) = X\beta \quad (8.2)$$

where X is the design (input) matrix and β is the vector of parameters. It is important to note that X is not restricted to only linear combinations of variables; for example, $X = [\sin(X), XY, Z^2]$ would still produce a model that was *linear-in-the-parameters*.

It is possible to represent any *linear-in-the-parameters* model as a GP, via relation to Bayesian linear regression [156]. This represents the most simple case of GP construction and is useful to highlight parallels between GPs, regression and linear models. Assuming a process $f(X) = X\beta$ with prior $\beta \sim \mathcal{N}(0, B)$ between two input vectors X_i and X_j , the covariance is derived

$$\begin{aligned} \text{cov}(f(X_i), f(X_j)) &= \mathbb{E}[f(X_i)f(X_j)] - \mathbb{E}[f(X_i)]\mathbb{E}[f(X_j)] \\ &= \mathbb{E}[(X_i\beta)(X_j\beta)^T] - \mathbb{E}[X_i\beta]\mathbb{E}[X_j\beta] \xrightarrow{0} \\ &= X_i\mathbb{E}[\beta\beta^T]X_j^T \\ &= X_iBX_j^T \end{aligned} \quad (8.3)$$

with expectation

$$E[f(X)] = 0 \quad (8.4)$$

leading to Gaussian process of

$$f(X) \sim \mathcal{GP}(0, X_iBX_j^T) \quad (8.5)$$

This is equivalent to a zero-mean GP with a linear kernel. If used in this form, the model will perform equivalently to a Bayesian linear regression, but with a higher computational cost. The advantage of this representation comes from an ability to be easily integrated within other GP models, as will now be presented.

For all upcoming model examples, a case will be considered where linear prior knowledge represents only partial knowledge of a system's complete behaviour. A process will be assumed of the form:

$$y = \underbrace{f(X)}_{\text{Prior knowledge}} + \underbrace{g(X)}_{\text{Unknown process}} + \underbrace{\epsilon}_{\text{Noise}} \quad (8.6)$$

where $f(X) = X\beta$ represents prior knowledge of a process in linear form, $g(X) \sim \mathcal{GP}(0, K)$ represents purely data-based capture of unknown phenomena and $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$ is Gaussian noise.

8.2 Modification of GP model structures

The modification of either the mean or covariance function presents two clear avenues via which one might wish to alter the modelling capabilities of a GP. The modification of the mean function is generally more common within the literature [8, 100, 107, 142–147], however, as this chapter investigates, is not always the optimum method through which to incorporate prior knowledge. The capabilities of a model, and therefore the most appropriate modification to structure, should be tuned to best fit a given modelling task.

This chapter explores differences between four model structures. Though more detail is provided in the upcoming sections, these models are summarised as:

- Deterministic mean function: Certain knowledge of how the mean of a process behaves.
- Stochastic mean function: Knowledge of how the mean of a process behaves with an associated degree of uncertainty.
- Physics-informed kernel with zero mean prior: Knowledge of the nature of functions present with a process without knowledge of the mean behaviour.
- Physics-informed kernel with known mean: Knowledge of both the nature of functions present within a process and its mean behaviour.

The following sections derive the model structures applied to fit Equation 8.6. This allows results to be directly compared.

8.2.1 Deterministic mean function

A deterministic model assumes a fixed output for any given input. Within an engineering context, this is a strict assumption, even for highly controlled environments and well understood phenomena. A point estimate, rather than a predictive distribution (or degree of uncertainty), implies that the user has certain knowledge of how the mean of a process behaves. For circumstances such as simulations, certain knowledge of outputs is possible, and deterministic functions can offer a computationally efficient way to enforce trends within the mean of a process.

Although $m(X)$ may be any arbitrary function, for the purposes of consistency and comparison to other models, a linear mean function is used here of the form

$m(X) = X\beta$, where X is a set of inputs and the regression weights have no associated uncertainty $\beta \sim \mathcal{N}(b, 0)$. For a process of the form in Equation 8.6, a deterministic mean leads to a structure of:

$$y = \underbrace{X\beta}_{\text{Deterministic mean}} + \underbrace{g(X)}_{\text{Unknown process}} \quad (8.7)$$

where the unknown process, $g(X)$ is captured by a GP with a flexible kernel. To derive the mean and covariance functions of the complete model, the expectation of the deterministic mean is first calculated:

$$\mathbb{E}[m(X)] = X\mathbb{E}[\beta] = Xb \quad (8.8)$$

followed by its covariance

$$\begin{aligned} \text{cov}(m(X_i), m(X_j)) &= \mathbb{E}[m(X_i)m(X_j)] - \mathbb{E}[m(X_i)]\mathbb{E}[m(X_j)] \\ &= \mathbb{E}[(X_i\beta)(X_j\beta)^T] - (Xb)(Xb)^T \\ &= X_i b b^T X_j^T - X_i b b^T X_j^T \\ &= 0 \end{aligned} \quad (8.9)$$

This leads to a GP of

$$y = \mathcal{GP}(Xb, K(X_i, X_j) + \sigma_n^2 \mathbb{I}) \quad (8.10)$$

where $K(X_i, X_j)$ is a kernel of choice. Note that since there is no effect from the deterministic mean on the covariance of the GP, the predictive variance remains unchanged. However, the predictive mean now becomes

$$E[y^*] = X_* b + K(X, X_*)(K(X, X) + \sigma_n^2 \mathbb{I})^{-1}(y - Xb) \quad (8.11)$$

This may be interpreted as the sum of the mean function and a zero mean GP implemented on the residuals $(y - Xb)$. Relating back to the interpretations of Equation 3.8 in Chapter 3, the predictions of a GP will follow the mean function when predicting far from observed training data. This makes mean functions an effective way to assist with extrapolation.

8.2.2 Stochastic mean function

Stochasticity acknowledges the presence of uncertainty and is necessary to account for variability within real world applications. A stochastic mean function will produce an added contribution $\sigma_{m(X)}^2$ within the covariance function. The added contribution will be dependent on the mean function used and accounts for the uncertainty present within the output of the stochastic model. For example, one might wish to use the output of a Finite Element (FE) model as a mean function, with a variance output of the model (if provided) as the covariance contribution. In other circumstances, such as for a linear mean function, it is possible to calculate the covariance contribution to be included. The inclusion of a stochastic mean function in a GP leads to a model of the form:

$$y = \mathcal{GP}(m(X), K(X_i, X_j) + \sigma_{m(X)}^2 + \sigma_n^2 \mathbb{I}) \quad (8.12)$$

To align with the earlier deterministic case, a stochastic mean function is assumed here of the form $m(X) = X\beta$ with $\beta \sim \mathcal{N}(b, B)$. Though the expectation remains the same, $\mathbb{E}[m(X)] = Xb$, it is now important to account for the covariance contribution, which is derived as

$$\begin{aligned} \sigma_{m(X)}^2 = \text{cov}(m(X_i), m(X_j)) &= \mathbb{E}[m(X_i)m(X_j)] - \mathbb{E}[m(X_i)]\mathbb{E}[m(X_j)] \\ &= \mathbb{E}[(X_i\beta)(X_j\beta)^T] - X_i b b^T X_j^T \\ &= X_i \mathbb{E}[\beta\beta^T] X_j^T - X_i b b^T X_j^T \\ &= X_i [B + b b^T] X_j^T - X_i b b^T X_j^T \\ &= X_i B X_j^T \end{aligned} \quad (8.13)$$

An important consideration when implementing a stochastic prior mean is how the components of the model are optimised; this will impact both the final model structure and performance. A **residual model** first fits a mean function to a dataset, with the residual then passed to a zero mean GP. For a stochastic linear mean this would be expressed:

$$y = \underbrace{\mathcal{GP}(Xb, XBX^T)}_{\text{Stochastic mean fitted first}} + \underbrace{\mathcal{GP}(0, K(X_i, X_j) + \sigma_n^2 \mathbb{I})}_{\text{Residual GP}} \quad (8.14)$$

Residual models provide a flexible and effective means of incorporating stochastic mean functions with little modification to existing GP architecture. The residual GP

in Equation 8.14 is zero mean, with a chosen flexible kernel, a widely implemented choice of model structure. Although the linear case is shown in Equation 8.14, there is no restriction on types of mean functions that may be used, making residual models suitable for building upon existing physics based approaches (e.g. simulation outputs or empirical laws). The major drawback of residual models occurs when the presence of unknown processes, aimed to be modelled by a residual GP, makes the initial fit of the mean function challenging. The biasing of model parameters and mis capture of an intended function component can be common problems, with this phenomenon highlighted later through a toy example.

An alternate implementation of a stochastic prior mean is to jointly optimise the whole model, learning the stochastic mean alongside the remainder of the GP in a single step. For a stochastic linear mean, this is expressed:

$$y^* = \mathcal{GP}(Xb, K(X_i, X_j) + XBX^T + \sigma_n^2\mathbb{I}) \quad (8.15)$$

Implemented in this way, this is equivalent to the inclusion of a linear kernel with a linear prior mean [139]. The simultaneous learning of model components can be of particular help when the learning of a linear component is disturbed by the presence of other processes.

8.2.3 Physics-informed kernel with zero mean prior

The additive inclusion of covariance functions with zero mean priors will not affect the mean function of a GP. In general, a physics-informed kernel with a zero mean prior is useful to include when one knows something about the form of a process (e.g. linear, periodic, smooth) but not specifically about magnitudes of function parameters or mean process behaviour. Alternatively, a process may be known to have a zero mean, a common example of which is white noise, $f(X) = \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma_n^2\mathbb{I})$, with covariance calculated:

$$\begin{aligned} K(X_i, X_j) &= cov(f(X_i), f(X_j)) = \mathbb{E}[f(X_i)f(X_j)] - \mathbb{E}[f(X_i)]\mathbb{E}[f(X_j)] \\ &= \mathbb{E}[\epsilon\epsilon^T] - \underbrace{\mathbb{E}[\epsilon]\mathbb{E}[\epsilon]}_0 \\ &= \sigma_n^2\mathbb{I} \end{aligned} \quad (8.16)$$

An alternative view of incorporating noise within a GP is the addition of the ‘white noise kernel’. The structure of the kernel ($\sigma_n^2\mathbb{I}$) implies values are only correlated with themselves (zero off the diagonal), with the magnitude of the

variance controlled by the noise variance σ_n^2 .

To represent any of the engineering knowledge present within Table 8.1, a kernel may be derived for a linear process $f(X) = X\beta$ with prior $\beta \sim \mathcal{N}(0, B)$:

$$\begin{aligned}
 K(X_i, X_j) &= \text{cov}(f(X_i), f(X_j)) = \mathbb{E}[f(X_i)f(X_j)] - \mathbb{E}[f(X_i)]\mathbb{E}[f(X_j)] \\
 &= \mathbb{E}[(X_i\beta)(X_j\beta)^T] - \mathbb{E}[X_i\beta]\mathbb{E}[X_j\beta] \xrightarrow{0} \\
 &= X_i\mathbb{E}[\beta\beta^T]X_j^T \\
 &= X_iBX_j^T
 \end{aligned} \tag{8.17}$$

Note how this is equivalent to both the linear kernel (8.3) and the stochastic linear mean function covariance contribution in (8.13). When working with similar processes, similar terms will often appear and whether these are determined a ‘kernel’, a ‘covariance contribution’ or another name may depend on the frame of reference.

8.2.4 Physics-informed kernel with known mean

Along with an ability to encode information about the form of a process, the introduction of non-zero priors within a covariance function allows information about the magnitude of function components, specifics about functional behaviours (e.g. long or short term trends, quickly or slowly varying processes) and mean process behaviour to be incorporated within a model.

As earlier stated, the additive inclusion of a linear covariance function with a linear prior is equivalent to a jointly optimised stochastic linear mean function. The advantage of this approach over residual modelling is that since all function components are now learned simultaneously, the likelihood of model components learning their intended part of a process is increased. If an initial model component is first fit to the data, the biasing of parameters may occur due the presence of processes not represented within the physical model.

Continuing with a model structure consistent with earlier examples, the introduction of a covariance component with a linear mean can be shown through treatment of the whole process:

$$y = f(X) + g(X) + \epsilon \tag{8.18}$$

where $f(X) = X\beta$ with prior $\beta \sim \mathcal{N}(b, B)$ and $g(X) \sim \mathcal{GP}(0, K)$. The covariance

of the combined process is derived as

$$\begin{aligned}
\text{cov}(y, y) &= \text{cov}(f(X_i), f(X_j)) + \text{cov}(g(X_i), g(X_j)) + \text{cov}(\epsilon, \epsilon) \\
&= \mathbb{E}[f(X_i)f(X_j)] - \mathbb{E}[f(X_i)]\mathbb{E}[f(X_j)] + K(X_i, X_j) + \sigma_n^2\mathbb{I} \\
&= \mathbb{E}[(X_i\beta)(X_j\beta)^T] - X_i b b^T X_j^T + K(X_i, X_j) + \sigma_n^2\mathbb{I} \\
&= X_i \mathbb{E}[\beta\beta^T] X_j^T - X_i b b^T X_j^T + K(X_i, X_j) + \sigma_n^2\mathbb{I} \\
&= X_i [B + b b^T] X_j^T - X_i b b^T X_j^T + K(X_i, X_j) + \sigma_n^2\mathbb{I} \\
&= X_i B X_j^T + K(X_i, X_j) + \sigma_n^2\mathbb{I}
\end{aligned} \tag{8.19}$$

with expectation

$$\mathbb{E}[f(X)] = X\mathbb{E}[\beta] = Xb \tag{8.20}$$

leading to Gaussian process of

$$f(X) \sim \mathcal{GP}(Xb, K_y + XBX^T) \tag{8.21}$$

where $K_y = K(X_i, X_j) + \sigma_n^2\mathbb{I}$. Taking the expression for predictive mean and variance

$$\mathbb{E}[f(X_*)] = m(X_*) + K_*^T (K + \sigma_n^2\mathbb{I})^{-1} (y - m(X)) \tag{8.22}$$

$$\mathbb{V}[f(X_*)] = K_{**} - K_*^T (K + \sigma_n^2\mathbb{I})^{-1} K_* \tag{8.23}$$

substituting the expressions for mean and covariance will give

$$\mathbb{E}[f(X_*)] = X_* b + (K_* + X B X_*^T)^T (K_y + X B X^T)^{-1} (y - X b) \tag{8.24}$$

$$\begin{aligned}
\mathbb{V}[f(X_*)] &= (K_{**} + X_* B X_*^T) \\
&\quad - (K_* + X B X_*^T)^T (K_y + X B X^T)^{-1} (K_* + X B X_*^T)
\end{aligned} \tag{8.25}$$

these may be rearranged beginning with the predictive mean as follows

$$\mathbb{E}[f(X_*)] = X_* b + (K_* + X B X_*^T)^T (K_y + X B X^T)^{-1} (y - X b)$$

which may be expanded using the Woodbury inversion [284]

$$\begin{aligned}
\mathbb{E}[f(X_*)] &= X_* b + (K_*^T + X_* B X^T) \\
&\quad (K_y^{-1} - K_y^{-1} X (B^{-1} + X^T K_y^{-1} X)^{-1} X^T K_y^{-1}) (y - X b)
\end{aligned}$$

leading to the form

$$\mathbb{E}[f(X_*)] = X_*\hat{\beta} + K_*^T K_y^{-1}(y - X\hat{\beta}) \quad (8.26)$$

where $\hat{\beta} = (B^{-1} + X^T K_y^{-1} X)^{-1}(B^{-1}b + X^T K_y^{-1} y)$. This form of the predictive mean is useful in terms of interpretability as it maintains the same structure as Equation 8.22, with an effective mean function of $m(X) = X\hat{\beta}$. This allows the prior b and posterior $\hat{\beta}$ linear model parameters to be compared. Another benefit occurs when one may be interested in the case of uninformative (or very weak) priors. As $B \rightarrow \infty$, the term $XB X^T$ present in the original expression for predictive mean would lead to numerical errors. However, as $B^{-1} \rightarrow 0$, the now present term B^{-1} may cancel. As both expressions may work for smaller B values, it is sensible to stick with a single implementation [102].

Alternatively, the predictive mean may be expressed

$$\mathbb{E}[f(X_*)] = K_*^T K_y^{-1} y + (X_* - K_*^T K_y^{-1} X)\hat{\beta} = K_*^T K_y^{-1} y + R^T \hat{\beta} \quad (8.27)$$

where $R = X_*^T - X^T K_y^{-1} K_*$. This form of the expression is useful to consider for implementation reasons as the first term is equivalent to a standard predictive mean. This makes modification to existing code easier to implement and allows for comparison of model contributions.

For the the predictive variance

$$\mathbb{V}[f(X_*)] = (K_{**} + X_* B X_*^T) - (K_* + X B X_*^T)^T (K_y + X B X^T)^{-1} (K_* + X B X_*^T)$$

expanding using the Woodbury inversion [284]

$$\begin{aligned} \mathbb{V}[f(X_*)] = & (K_{**} + X_* B X_*^T) \\ & - (K_*^T + X_* B X^T) (K_y^{-1} - K_y^{-1} X (B^{-1} + X^T K_y^{-1} X)^{-1} X^T K_y^{-1}) \\ & (K_* + X B X_*^T) \end{aligned}$$

leading to the form

$$\mathbb{V}[f(X_*)] = (K_{**} - K_*^T K_y^{-1} K_*) + R^T (B^{-1} + X^T K_y^{-1} X)^{-1} R \quad (8.28)$$

where $R = X_*^T - X^T K_y^{-1} K_*$. Similarly to the predictive mean expressions, the form of predictive variance is more suitable for implementation in cases of uninformative priors. Occurrences of B are no longer present and have been replaced

by occurrences of B^{-1} . Advantages in implementation are also achieved due to the first term being equivalent to the expression for standard predictive variance.

8.2.5 Summary of methods

It is often the case that when modifying either the mean or covariance function of a GP, changes to the other may be produced. Figure 8.1 highlights how beginning from differing perspectives may produce similar effects within final model structures.

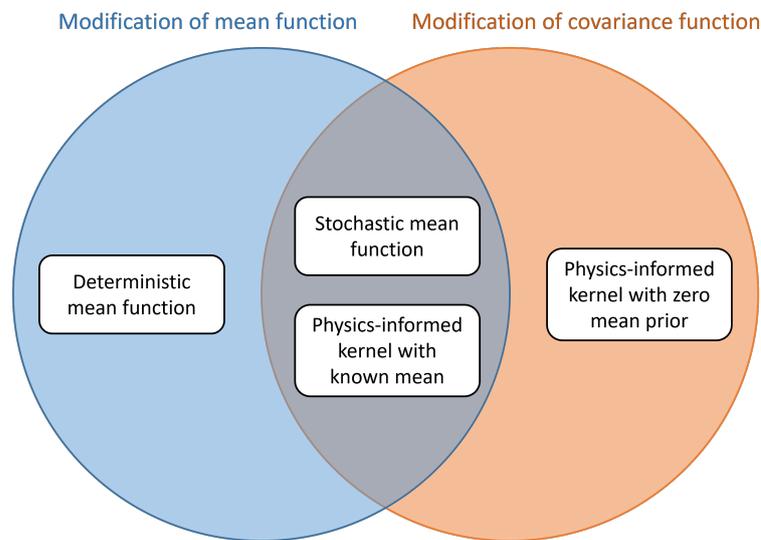


Figure 8.1: Overlap in consequences of modification of the mean and covariance functions of a GP.

Deterministic mean functions will modify the mean function of a Gaussian process but will have no impact on the covariance function. This is because unlike **stochastic mean functions**, there is no degree of uncertainty associated to the mean function. They are most suitable for use when one has certain knowledge about an underlying trend, for example, in cases of simulated data.

Utilising **physics-informed kernels with zero mean priors** will change the structure of a GPs covariance function, but does not include prior knowledge within the mean function. They are useful when one might have knowledge about the nature of functions present within a process, but not about the magnitude of underlying trends.

Stochastic mean functions and **physics-informed kernels with known means** incorporate prior knowledge within both the mean and covariance functions of a GP. For **stochastic mean functions**, an additional term is included within the covariance function, accounting for uncertainty within the representation of the mean function. For **physics-informed kernels**, a prior mean provides an avenue through which to incorporate knowledge of physical parameter magnitudes or underlying process behaviours.

8.3 Toy case study: Learning from limited data

One of the scenarios where prior knowledge can be most helpful is when observed data are ‘limited’. As introduced in Chapter 1, this does not necessarily mean that only a small amount of data has been collected, but that the collected data may cover only a partial area of potential conditions. It is very possible within a monitoring campaign to measure many hundreds of Gigabytes (or even Terabytes) of data and still account for a small fraction of the potential conditions observable. The ability of models to perform outside the ranges of their observed training data is therefore very important and is investigated here through the use of a toy data example.

An important first point to highlight is that a function need not grow very complex before it becomes easy to misrepresent the function by observing a limited range of data. Consider the fit of a linear model, having been shown various training ranges of a function of the form $y = \alpha x + \epsilon$, shown in Figure 8.2.

In the case of the linear function, the reduction in training data shown to the model does not significantly impact performance. A similar fit is achieved whether the model has observed the majority of the dataset or a smaller subset. This would be a favourable phenomenon to occur within an engineering dataset, and would allow for high predictive performance outside observed data ranges. However, such a modelling task would be unlikely to occur within complex systems and changing real world conditions.

Consider, now the introduction of a second process within the target function such that it is now of the form $y = \alpha x + \beta \sin(\omega x) + \epsilon$. The fit of a linear model, having been shown varying ranges of training data is repeated and shown in Figure 8.3.

Here, the introduction of the additional process within the target function made the capture of the underlying linear component more difficult. Although the underlying linear component was captured well when the model observed a large data range,

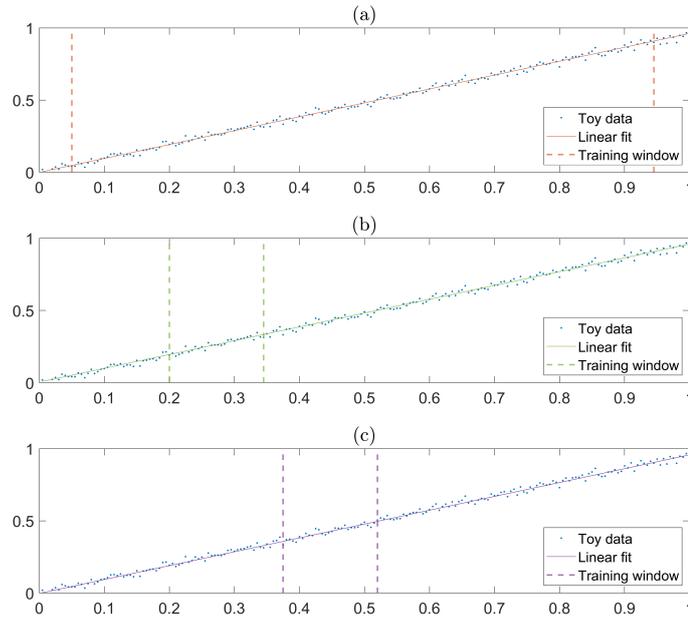


Figure 8.2: The fit of a linear model to a function of the form $y = \alpha x + \epsilon$. A large training range is shown to the model in (a), with smaller training ranges shown in (b) and (c). The fit of the model remains reasonable for all training windows.

in cases where the model observed a limited range, it was possible to both under and over estimate the true gradient. This toy example aims to highlight how misrepresentation of a process can occur when only a subset of conditions are observed. The models introduced here will highlight how manipulation of the kernel and mean function can combat this. Remaining with a toy function of the form $y = \alpha x + \beta \sin(\omega x) + \epsilon$, models are shown the first 25% of the dataset during training and predict across the full range.

8.3.1 Black-box SE kernel fit

The dataset is first modelled using a zero mean GP with Squared Exponential (SE) kernel, a model of the form

$$f(X) \sim \mathcal{GP}(0, K_{SE}(X_i, X_j) + \sigma_n^2 \mathbb{I}) \quad (8.29)$$

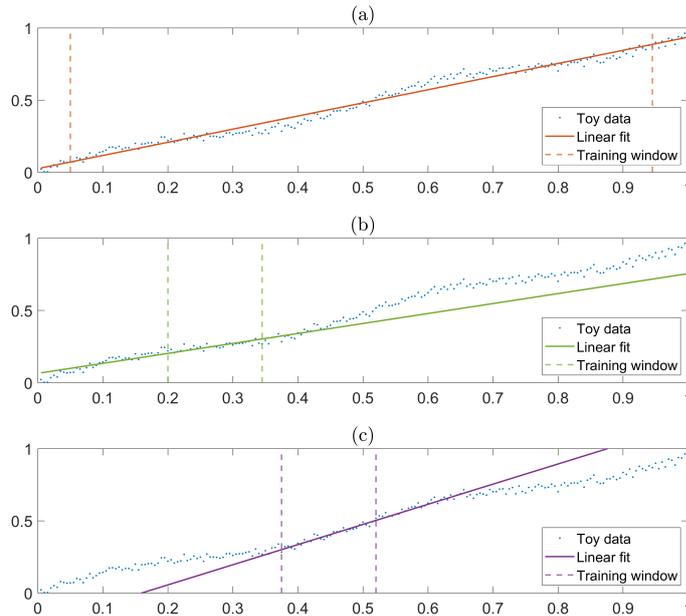


Figure 8.3: The fit of a linear model to a function of the form $y = \alpha x + \beta \sin(\omega x) + \epsilon$. The model is able to capture the underlying linear component of the target function when a large training range is shown to the model (a). However when limited training ranges are shown, a linear fit may underestimate (b) or overestimate (c) the gradient of the underlying linear component within the target function.

This is what one might consider an ‘out of the box’ or ‘default’ modelling approach. No particular considerations have been made about the process being modelled. A SE kernel can be suitable for any process that is smoothly varying and continuous and so often provides a start point from which to build more complex models. It may be used as an initial model to study a dataset or integrated within final model designs. As seen from the training window in Figure 8.4, SE kernels generally perform well in interpolation, with an excellent fit to previously observed data. However, when the training data runs out, performance quickly drops and the model refers to its prior. As one would expect, the SE kernel alone is unable to extrapolate effectively.

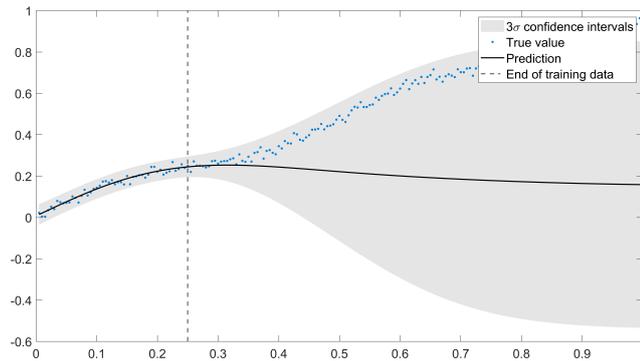


Figure 8.4: Prediction of a zero mean GP with a SE kernel on a toy dataset of the form $y = \alpha x + \beta \sin(\omega x) + \epsilon$. The model is shown the first 25% of the dataset during training.

8.3.2 Stochastic mean function

A residual model is one way of incorporating a stochastic mean function within a GP. A model is first fit to the target dataset, here a Bayesian linear regression, with the residual passed to a zero mean GP. This leads to a combined model of the form

$$f(X) \sim \mathcal{GP}(Xb, XBX^T) + \mathcal{GP}(0, K_{SE}(X_i, X_j) + \sigma_n^2 \mathbb{I}) \quad (8.30)$$

A linear model is first fit to the data to provide an estimate for a linear mean function, with the result of this shown in Figure 8.5. However, the presence of additional processes, represented here as $\beta \sin(\omega x)$, may sometimes disrupt the learning of a linear function component. This is a particular problem in circumstance of limited data where the chance of a non-representative data region (i.e. a region too steep or shallow) is more likely.

The GP is able to achieve a reasonable fit to the residual within bounds of previously observed data. However, when far from observed data, the residual GP reverts to its prior. The GP fit to the residual is shown in Figure 8.6.

When the residual GP is combined with the linear model, the overall model performance is improved within the bounds of observed data. Far from observed data, the combined model reverts back to the originally fitted linear model, which acts as a stochastic mean function. This has the potential to be a highly useful property if an appropriate or representative mean function is able to be fitted. It allows for improvements in performance where data are available, and may otherwise

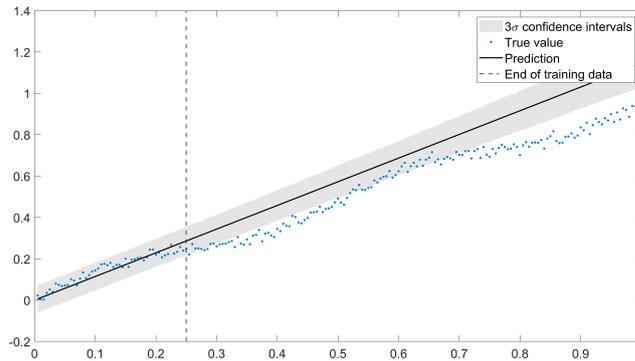


Figure 8.5: Bayesian linear regression on a toy dataset of the form $y = \alpha x + \beta \sin(\omega x) + \epsilon$. The model is shown the first 25% of the dataset during training. The model overestimates the gradient of the linear function component (α) due to the presence of the sinusoidal term ($\beta \sin(\omega x)$).

revert to a baseline performance. Where this type of model may underperform is in circumstances highlighted in Figure 8.7. Here the underlying linear model was poorly fit and the model therefore performed poorly in extrapolation. Far from observed data, residual models are highly dependent on their initially fit model (mean function) and care should be taken to ensure this model is performing as expected.

8.3.3 Physics-informed kernel with zero mean prior

If knowledge of the functional form of a process is known, this structure may be represented within the covariance function (kernel) of a Gaussian process. The kernel is responsible for controlling the family of functions from which predictions may be drawn and design choices may be made to enforce desired behaviours. For the example of a linear function component, using the earlier derived linear kernel in Equation 8.3, it may be included within the covariance function of the Gaussian process such that it is now of the form:

$$f(X) \sim \mathcal{GP}(0, K_{SE}(X_i, X_j) + XBX^T + \sigma_n^2 \mathbb{I}) \quad (8.31)$$

One advantage of kernel representation of processes is the potential for single step learning. This can help to avoid the biasing/disruption caused by trying to learn a function component as part of a combined process. An aim of kernel design is to

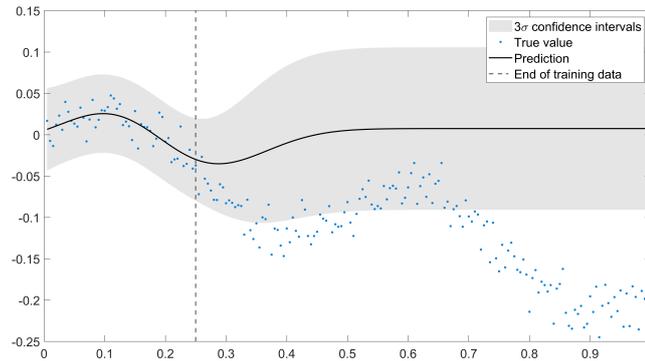


Figure 8.6: Prediction of a zero mean GP with a SE kernel on the residual of a Bayesian linear regression on a toy dataset of the form $y = \alpha x + \beta \sin(\omega x) + \epsilon$. The model is shown the first 25% of the dataset during training.

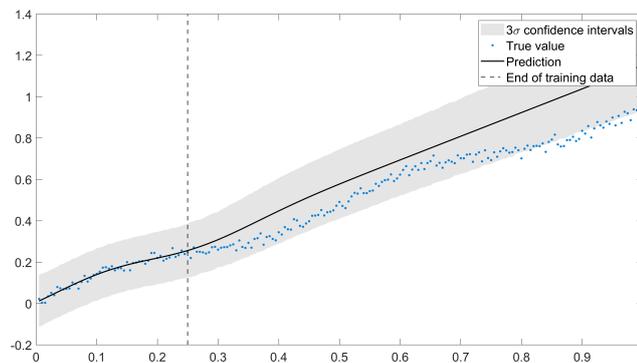


Figure 8.7: Prediction of a residual model on a toy dataset of the form $y = \alpha x + \beta \sin(\omega x) + \epsilon$. The model is comprised of a Bayesian linear regression and zero mean GP fit on the residuals; it is shown the first 25% of the dataset during training.

learn all processes simultaneously, such that each kernel component is able to best capture its intended function component. This phenomena can be seen within the improved linear component capture and extrapolation within Figure 8.8.

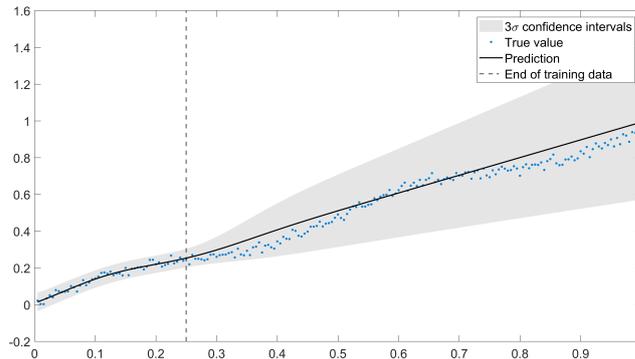


Figure 8.8: Prediction of a GP with a SE plus a zero mean linear kernel with structure $f(X) \sim \mathcal{GP}(0, K_{SE}(X_i, X_j) + XBX^T + \sigma_n^2\mathbb{I})$ on a toy function of the form $y = \alpha x + \beta \sin(\omega x) + \epsilon$. The model is shown the first 25% of the dataset during training.

8.3.4 Physics-informed kernel with known mean

The inclusion of covariance functions with non-zero prior means maintains the advantages of single step learning provided by representation of function components within the kernel. In addition, if specific prior knowledge is known about magnitudes of prior coefficients, this may now be included. Here the GP will be of the form

$$f(X) \sim \mathcal{GP}(Xb, K_{SE}(X_i, X_j) + XBX^T + \sigma_n^2\mathbb{I}) \quad (8.32)$$

The linear function component and SE kernel will be learnt at the same time following the theory from section 8.2.4. The prediction of the model is shown in Figure 8.9

The implementation of priors through the covariance function can also provide a degree of interpretability via access to the posterior distributions of specific parameters. The prior and posterior distributions of the linear parameter β is shown in Figure 8.10. This is of particular help where such parameters may have physical meaning e.g. coefficients within empirical laws.

Any advantages provided through the inclusion of prior means alongside covariance functions is highly dependent on the accuracy and strength of priors and their appropriate selection is therefore an important modelling decision. It is very possible to worsen model performance through inclusion of priors, particularly in cases of overly confident inaccurate parameters.

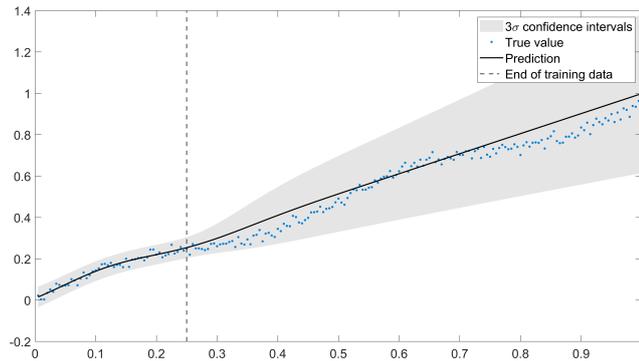


Figure 8.9: Prediction of a GP with a SE plus a non-zero mean linear kernel with structure $f(X) \sim \mathcal{GP}(Xb, K_{SE}(X_i, X_j) + XBX^T + \sigma_n^2\mathbb{I})$ on a toy function of the form $y = \alpha x + \beta \sin(\omega x) + \epsilon$. The model is shown the first 25% of the dataset during training.

8.3.5 Performance summary

The performance of models was measured using two metrics: the Normalised Mean Square Error (NMSE) and the Mean Standardised Log Loss (MSLL). The MSLL is a probabilistic measure, with superior models having more negative scores. A baseline of zero is equivalent to setting the predictive mean and variance for all test set points as the mean and variance of the training set. A comparison of results is shown in Table 8.2.

Table 8.2: Performance summary of models on toy dataset.

Model	Structure	NMSE	MSLL
Zero mean SE kernel	$\mathcal{GP}(0, K_{SE} + \sigma_n^2\mathbb{I})$	223.247	-14.727
Linear prior mean, residual model	$\mathcal{GP}(Xb, XBX^T) + \mathcal{GP}(0, K_{SE} + \sigma_n^2\mathbb{I})$	20.413	-15.458
Zero prior mean, linear + SE kernel	$\mathcal{GP}(0, K_{SE} + XBX^T + \sigma_n^2\mathbb{I})$	2.920	-16.948
Linear prior mean, linear + SE kernel	$\mathcal{GP}(Xb, K_{SE} + XBX^T + \sigma_n^2\mathbb{I})$	3.032	-16.954

The zero mean GP with a SE kernel, what might be considered the 'out of the box' approach, is suitable for a wide range of modelling tasks and capable of capturing a

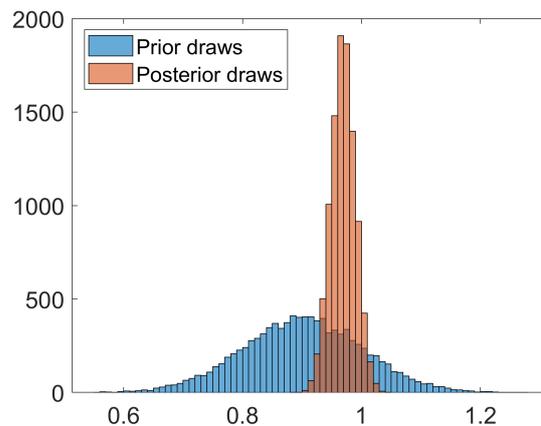


Figure 8.10: Prior and posterior draws of the linear model coefficients for a joint optimisation GP with a linear mean function and SE covariance $y = \alpha x + \beta \sin(\omega x) + \epsilon$. Training data was normalised using its calculated mean and standard deviation.

wide variety of functional forms. However, it performed poorly during extrapolation on the toy data set. This behaviour was to be expected as a stationary kernel, such as a SE, will revert to its prior far from data. In general, the implementation of flexible modelling approaches are only suitable when working within the limits of observed data.

A key observation within these results is the performance improvement of the kernel based methods over the residual model. This is primarily due to the mis-learning of the linear component via the initial fitting of the linear function. In an attempt to account the the sinusoidal component, an over estimation of the linear slope was achieved. Although a simplistic case is highlighted here, this scenario presents a common problem posed with residual model structures. If they are adopted in cases in which a model known to exclude parts of a process, it is important to first check, via the use of an additional validation step, that acceptable fit of the mean function may be achieved.

The performance of the kernel based methods for this example were found to be very similar, being within 0.112% NMSE and 0.006 MSL of each other. The difference in performance here is due to the placing of a prior over the parameter α . In this case, the true value for the toy dataset was known to be $\alpha = 1$, whilst a prior of $\alpha \sim \mathcal{N}(0.9, 0.1)$ was used as a linear prior mean. The fact that the linear prior mean in this case performed so similarly to a zero prior highlights the importance of both accurate and appropriately confident priors. Though within

10% of the true underlying value, the relatively high confidence placed over the prior limited the potential benefit to the model, with a slight underestimate of the true underlying linear gradient achieved.

8.4 Conclusions

This chapter explored how changing the way in which prior knowledge is incorporated within a Gaussian process, can change the capabilities and properties of the constructed model. To allow for direct comparison of results, a consistent example of prior knowledge was used in the form of a linear function $f(X) = X\beta$. This was then incorporated via both the mean function and kernel of a GP. As shown in Table 8.1, a linear form of prior knowledge is relevant to a wide range of engineering tasks, however, future efforts hope to expand the forms of prior knowledge considered. Constraints, non-linear equations and simulated prior knowledge would all be valuable to investigate in a similar manner.

The derivation of model expressions highlighted why differences in model capability occur due to the introduction of additional terms within expression for the predictive mean and variance. The role and utility of how these terms can be exploited by a user were discussed. For example, an additional term included within the predictive variance allows a stochastic mean function to account for uncertainty with the expression of a mean function, which a deterministic mean function does not possess.

The examples of models presented in this chapter were all of the additive partial knowledge structure presented in Equation 8.6, allowing for more direct comparison of the expressions produced. Although an additive structure is arguably the intuitive first case to investigate, as discussed in Chapter 7, it is not the only potential way to represent physical knowledge within a model. The investigation of other cases, such as a multiplicative partial knowledge structure, are a natural continuation of this work. Efforts of Jones [285], alongside the author are currently ongoing to investigate and compare alternative structures.

The classification of model structures utilising an equivalent form of prior knowledge was a useful step to aid the work of the next chapter of the thesis. In the upcoming Chapter 9, a proposed framework to assist a user with PIML model construction is developed. Based on the nature of prior physical knowledge possessed by a prospective modeler, the framework aims to suggest an appropriate model architecture. For this reason, it was very important to distinguish, understand and

clearly define architectures with similar capabilities. The first section of the chapter applies a similar logic to the type of physical knowledge possessed by the user, aiming to classify and describe common categories of prior physical knowledge.

A FRAMEWORK FOR PRIOR KNOWLEDGE INCLUSION WITHIN GAUSSIAN PROCESSES

Effective incorporation of prior knowledge within machine learners brings an array of potential benefits, including interpretability into model results, reduced demand for expensive data collection and increased operator confidence. However, deciding how best to integrate any given piece of prior knowledge within a model can be challenging and will have an effect on final model structure and performance. This chapter investigates important relationships between the form in which knowledge presents itself (e.g. equations, simulations, bounds), the source from which it comes (e.g. domain experts), its cost of accrual and how best to incorporate it within a model. A proposed framework is developed to assist with deciding how best to integrate prior knowledge within Gaussian Process Regression (GPR) models. The aim is to help classify a given piece of prior knowledge and suggest an appropriate model architecture. To the author's knowledge, this is the first attempt to develop a tool to assist a user with the development of PIML models. This chapter is built upon work from throughout the thesis, relying upon the understanding of GP model structures presented in Chapters 3 and 8, and utilising the PIML model development of Chapters 5, 6 and 7 as case studies with which to test the developed framework.

9.1 Categorising prior knowledge

The most appropriate means of integration of prior knowledge within any machine learner is highly dependent on the knowledge being integrated. Here, therefore, some steps are taken to understand and categorise common types of prior knowledge a prospective modeler might utilise. Factors including where the knowledge comes from, here defined as the knowledge source, along with its structure (equations, simulations, logical relations etc.), here defined as the knowledge form, are considered.

9.1.1 Types of knowledge source

The source of prior knowledge can have a significant influence on its value, ease of implementation within a model and potential improvements in performance [10]. Here, the following definitions of knowledge sources are proposed:

Universal knowledge: Known by an average person, or able to be very quickly explained. Explanations should be immediately obvious, not requiring a discussion e.g. “The sky is blue”.

Basic reasoning: Something that could be easily inferred from looking at a scenario. Explaining may require some level of justification e.g. “The trolley will roll down the hill”.

Scientific knowledge: Based on scientific theory but need not require any specialised knowledge from a specific field. Approximate level of science that might be taught to teenage school children e.g. $F = ma$.

Domain knowledge: Technical knowledge relating to a specific field. It should be known and understood by the majority of people working within the field e.g. $\omega_n = \sqrt{\frac{k}{m}}$, the natural frequency ω_n of a system, given its stiffness k and mass m .

Expert knowledge: Knowledge requiring years of experience, or high levels of expertise within a specific application e.g. Soil interaction effects at a particular wind farm.

With each subsequent knowledge source presented, the availability of knowledge reduces, becoming more difficult to acquire and with generally increasing costs. This concept is highlighted in Figure 9.1. An important consideration relating to the source of prior knowledge is the trade-off between the cost of knowledge (in terms of money, time and expertise) and the potential benefit that may be provided to a model. This is maximised when easily accessible universal knowledge or basic reasoning is able to provide a large benefit to model performance. The reverse case, in which expensive or challenging to acquire expert knowledge provides little benefit to a model should be avoided. In the case where knowledge is already available e.g. existing results of a simulation, the utilisation should be maximised, so long as results are validated.

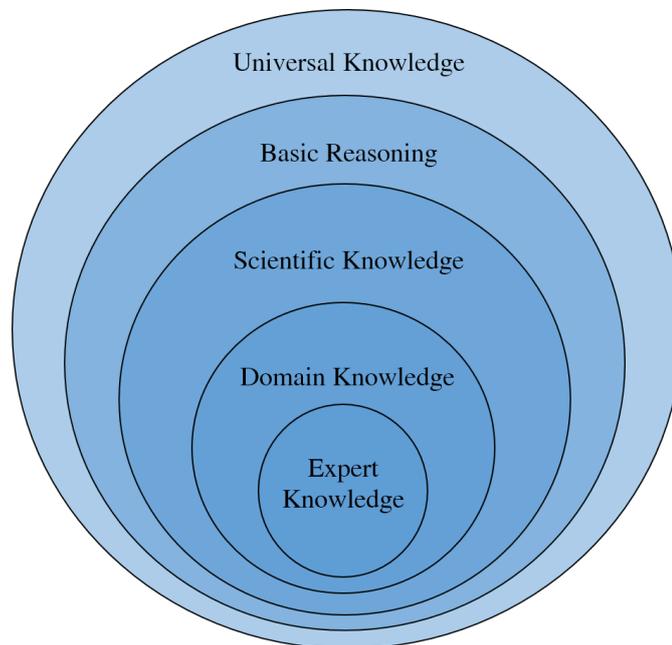


Figure 9.1: Diagram of knowledge sources. Rings represent knowledge contained within a particular source and are intentionally overlapping; for example, an expert would also be expected to have access to scientific and universal knowledge.

The author believes that the classification of knowledge sources is a useful step before integration of prior knowledge within machine learning; however, their classification can be subjective. It is helpful to consider how the process might be applied to an engineering example. A mass spring damper system is shown in Figure 9.2, with a summary of knowledge source examples in Table 9.1.

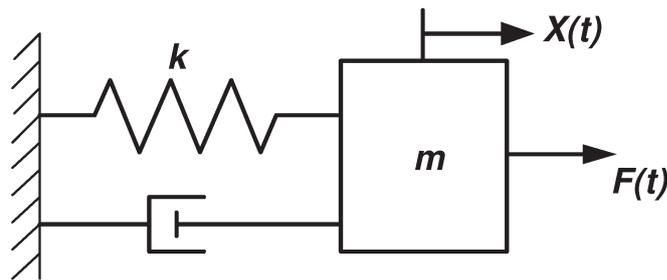


Figure 9.2: Mass spring damper system.

Table 9.1: Knowledge source examples for a mass spring damper system.

Knowledge type	Knowledge	Description
Universal knowledge	$m \geq 0$	The block exists and has some mass.
Basic reasoning	$x \geq x_{wall}$	The block may not move left of the wall boundary.
Scientific knowledge	$F = ma$	Newton's second law.
Domain knowledge	$m\ddot{x} + c\dot{x} + kx = F$	Equation of motion for a linear mass spring damper system.
Expert knowledge	$k = f(T)$	Expected stiffness of a specific spring as a function of temperature.

9.1.2 Knowledge forms

It is argued here that, the form which prior knowledge takes and the structure of how it may be expressed is the primary factor in how best it should be integrated within a machine learning model. The relationship between knowledge form and model structure is studied by von Rueden *et al.* [10]. There are many types of knowledge forms, the classification of which in to broader categories is presented here:

Variable dependence

When constructing models, an assessment of variable dependence is an important step when selecting inputs; this involves distinguishing between causal and associative relationships. If inappropriate inputs are selected for a model, very little may be done with model structure to remedy this [66]. Although statistical tests are helpful for the detection of dependence, determining whether the relationship is causal is challenging without the use of prior knowledge. Strong correlations between variables, indicating the presence of a linear relationship, are very possible without any direct causal link [76]. Associative relationships between model inputs and the target have the potential to lead to highly confident erroneous model predictions.

Bounding of variables

Without knowledge of the specific behaviour of a variable, it is often the case the extreme limits might be known either as an upper bound, a lower bound or both. The bounding of input variables might occur through knowledge of expected operational conditions e.g. temperature ranges. The bounding of outputs can be useful for preventing models from giving physically impossible or unrealistic predictions. The bounding of parameters (or hyperparameters) could be used to limit search spaces, reducing the computational burden during optimisation.

Functions of variables

Where relationships between variables can be expressed as a function, a series of powerful methods for prior knowledge inclusion become available. Where an approximate functional form of a process is known, e.g. through an empirical law, a fitted mean function can provide means through which to improve extrapolative capabilities. Access to the functional form of parts of a process allows for the mimicking of behaviours within the covariance function design. Where desired inputs for a model are unavailable or can not be measured, the approximation of these variables in function form has the potential to simplify learning tasks, particularly if the variable is known to strongly influence a model target.

Simulation outputs

The use of simulations is an important part of many steps within engineering applications including design, manufacturing and development of maintenance strategies. Used alongside experimental and physical studies, simulations are an excellent way to acquire helpful data to support engineering decisions. Simulations of a target variable can be utilised as a mean function, with a stochastic mean function possible where uncertainty estimates of variables are also provided. Similarly to

functional estimates of model inputs, simulations can provide useful access to additional model variables where measured data are unavailable. In all cases where a simulation output might be utilised, additional care should be taken that the outputs of the simulation are validated and to be trusted by the user.

Dimension reduction

The use of dimension reduction via prior knowledge is an effective means of simplifying the learning task presented to a data-based model. Dimension reduction here is used to refer to any means by which the operational space is transformed from a higher-dimensional space to a lower-dimensional one whilst aiming to retain important information present within the data. Examples of dimension reduction include the use of rotational and reflectional symmetries, representation of 3D problems as a 2D slice and 1D ‘inline’ representation of problems. Although the act of dimension reduction itself can simplify problems, the decision of where and when such simplifications are valid can be a challenging task.

Knowledge source examples

To highlight how various knowledge forms may present themselves within an engineering context, it is useful to present examples from a specific application. Much of the work within this thesis focusses on the wave loading of offshore structures, an important subtask in the estimation of their useful remaining life. Table 9.2 presents relevant prior knowledge of varying forms.

Table 9.2: Examples of varying knowledge forms within the context of wave loading of offshore structures.

Knowledge form	Expression	Description
Variable dependence	$p(F \cap X) = p(F)p(X)$	Wave load independent of lightning flash.
Bounding of variables	$6^\circ C \leq T \leq 18^\circ$	Average North Sea temperature range [286].
Functions of variables	$F = U U C'_d + \dot{U}C'_m$	Morison’s Equation.
Simulation outputs	$f(X) = F_{FE}$	Results from an external Finite Element model.
Dimension reduction	$U(x) \sim U(x, y, z)$	Inline representation of problem.

The form in which prior knowledge is expressed is often linked to the source from which it is provided, for example, knowledge of appropriate dimension reduction for a problem would not be provided through universal knowledge. It can be useful to consider this relationship when constructing physics informed models, with the aim of maximising the benefit provided to a model whilst minimising the resources spent (time, money and expertise) in acquiring knowledge. A key advantage of physics-informed models is that prior knowledge may be used to reduce a models reliance on data collection, thereby reducing monitoring efforts and costs. If a large amount of resources are spent in acquiring the knowledge to do so, the net benefits are diminished. The form of knowledge, closely linked to how most appropriately it should be integrated to a model, the benefits in performance this achieves, along with the knowledge source and expense of knowledge accrual should all be considered during model construction. The approximate relationship between knowledge sources and forms is presented in Figure 9.3. There is, of course, a large degree of subjectivity and overlap about what constitutes a boundary between knowledge sources; however the important takeaway is the relative positions between knowledge forms.

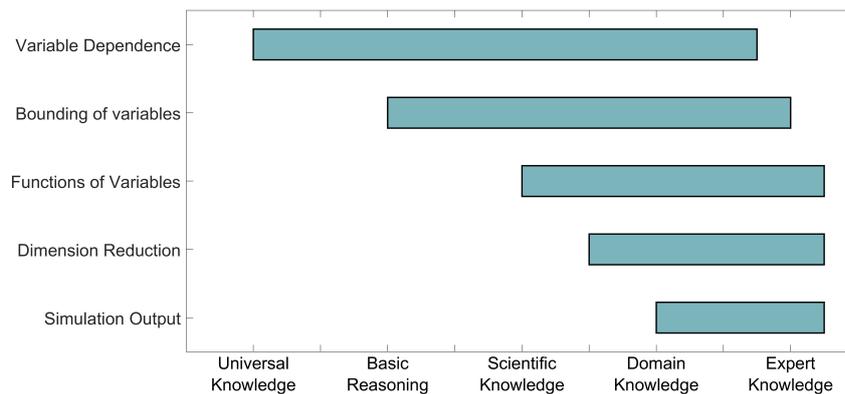


Figure 9.3: A comparison of the relationship between knowledge source and knowledge form.

A first key relationship between knowledge source and knowledge form drawn from Figure 9.3, is the difficulty of acquiring simulation data. Typically, the construction of any form of simulation, (CFD, FEA or otherwise) requires a high degree of domain expertise, placing a relatively high resource cost on this knowledge form. A consequence of this is that before any effort is spent creating simulations specifically to assist a machine learner, care should be taken to ensure alternative options could not provide similar benefit to the model.

The scope of knowledge sources for both variable dependence and the bounding

of variables was the widest, covering almost the full range of knowledge sources. The difficulty of both tasks and level of expertise required varies strongly with application, however examples from across the full range can be very valuable with PIML models. For instance, bounding variables within physical laws, or limitations can ensure predictions are physically meaningful [193].

9.2 A framework for prior knowledge inclusion

Deciding how best to integrate any given piece of prior knowledge within a machine learning model can be challenging and will have an effect on final model structure and performance. This section provides a framework, in the form of the flowchart in Figure 9.4, to assist with deciding how best to integrate prior knowledge within Gaussian process models. The aim is to break the decision down in to a series of more easily answered questions that may help classify a given piece of prior knowledge and suggest an appropriate means of integration within a model. An expected user of the framework should be comfortable with implementing Gaussian process models but need not be experienced with physics-informed modelling. For users who may have enough experience to develop their own judgments of how best to integrate a piece of prior knowledge, the framework should act as a sense check that may be useful alongside personal opinion.

Alongside the goal of suggesting an appropriate model architecture, there were several soft constraints placed on its design. Though these may have to be disregarded, if future versions necessitate it, they were deemed useful for the usability of the framework. These design constraints included:

- The flowchart should be able to fit on a single page. This would allow for use without flicking between pages, and an ability to be displayed in formats such as a poster.
- The boxes within the flowchart should be a consistent size. This would allow for an aligned layout and a clearer presentation.
- Avoid the use of mathematical expressions within flowchart boxes. Where possible, text descriptions of models and decisions are preferable will avoid complications with notation and clarity to a variety of readers.
- Maintain a readable text size. Attempts to shrink the flowchart to a single page should not be so extreme that the text present becomes difficult to read.

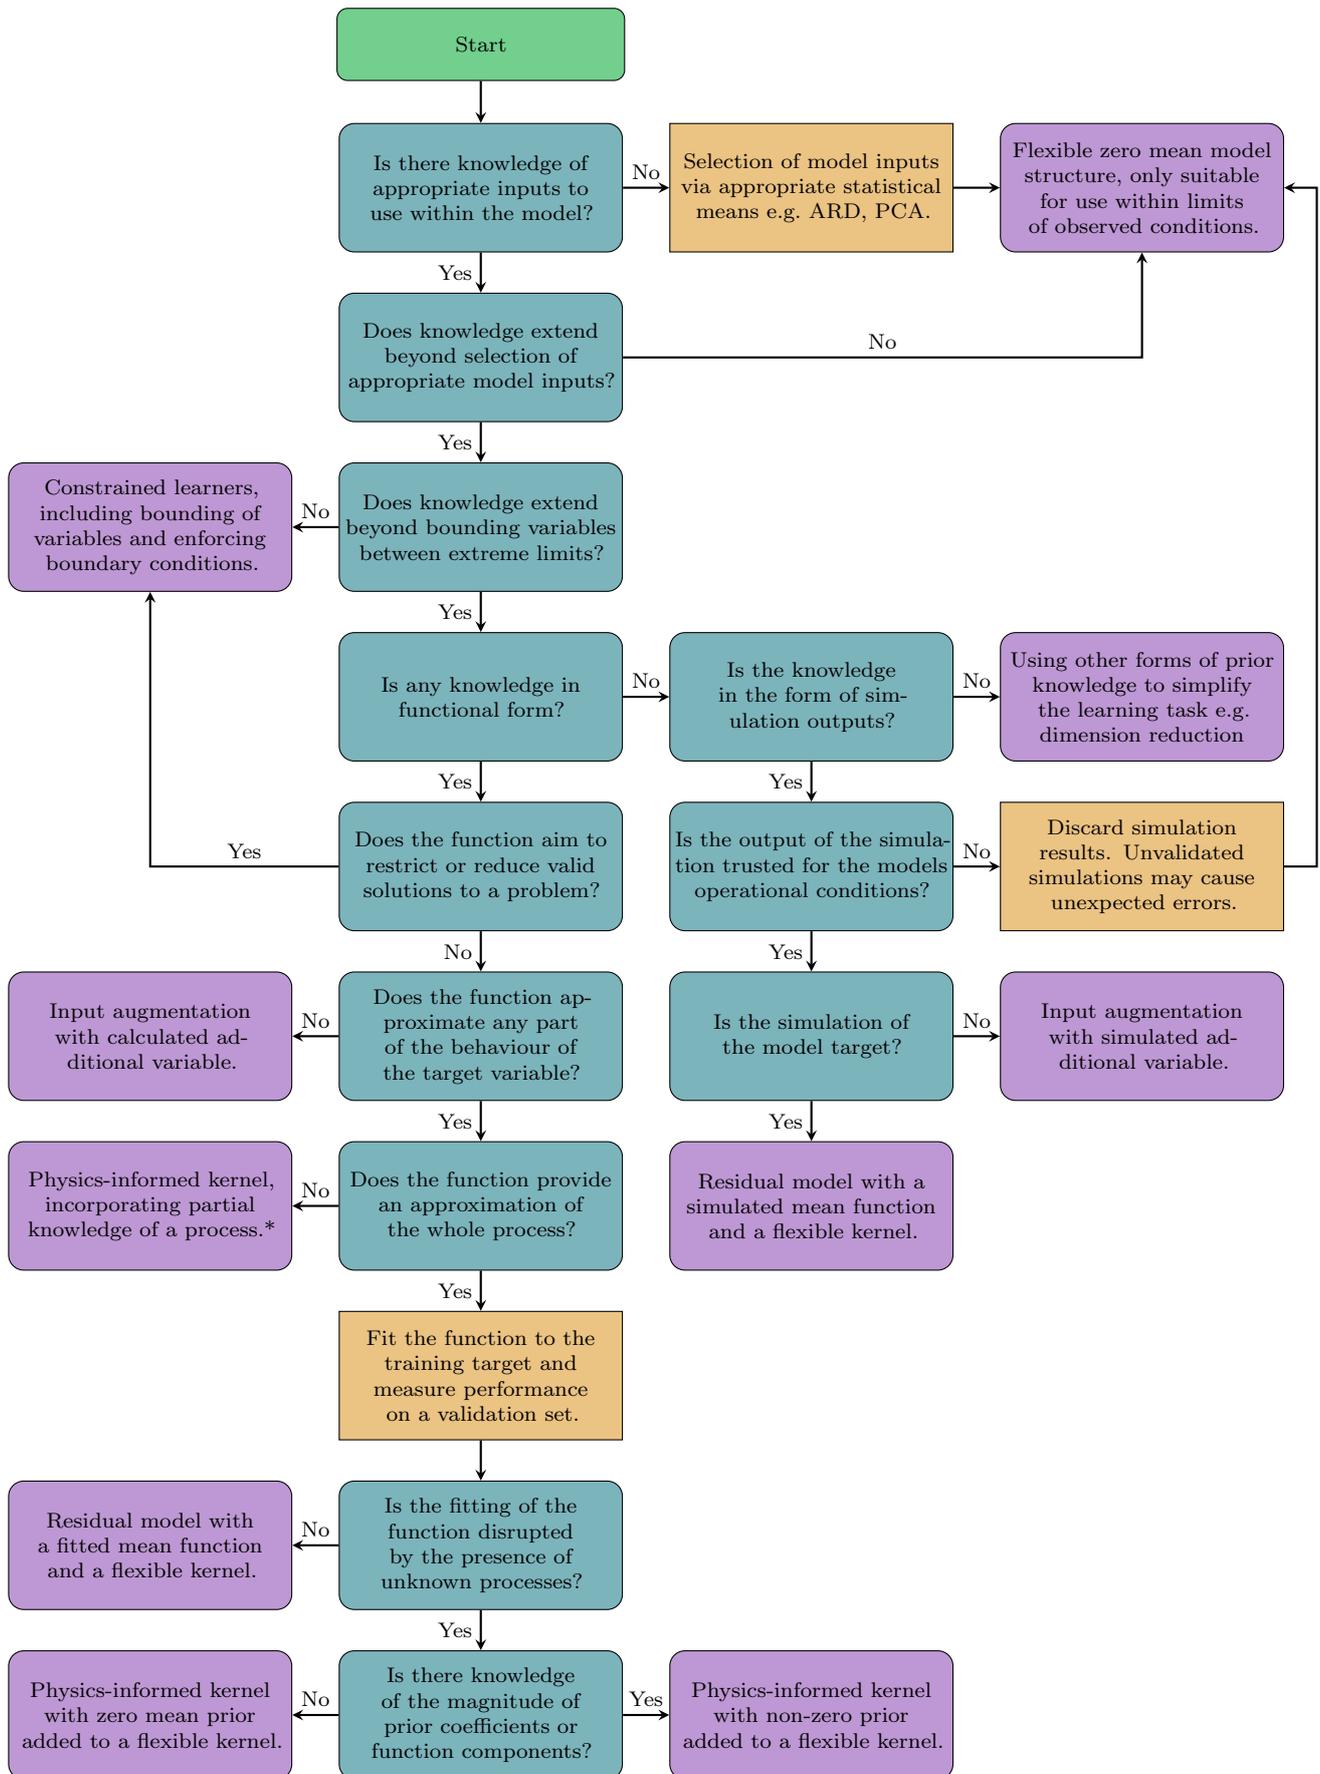


Figure 9.4: A flowchart for the inclusion of prior knowledge within Gaussian process design. Teal nodes represent decisions, orange nodes represent actions and purple nodes represent outcomes. *The design of kernels to account for partial knowledge of a process presents its own widely varying research area, which is detailed further in Chapter 7.

Application of the framework

In order for the framework to be useful, it is important that the outcomes and suggestions it produces should be sensible and that navigation of the flowchart itself is user friendly. This necessitates adherence to the following aims:

- The language used within the flowchart should be clear, concise and easily understood by an expected user.
- The outcomes produced by navigating the flowchart should align as closely as possible with effective implementation of prior knowledge found within this thesis and other literature.
- It should not be possible to end up in a situation in which a question asked within the flowchart does not apply or make sense, given the users modelling scenario.

To test the reliability of the framework, examples of prior knowledge used throughout the thesis are used as case studies. The hope is that successful examples of prior knowledge integration would be suggested as model structures via navigation of the flowchart. In each upcoming example, questions from the flowchart were answered based on only the prior knowledge with no knowledge of the PIML model structure assumed.

9.2.1 Example A: Morison's Equation

The first application example focusses on Morison's Equation, which formed the physics-based component of the first PIML model within the thesis in Chapter 5. For the modelling of wave loads on slender members, which many offshore structures are comprised of, Morison's Equation has been the most widely used such method since its introduction in 1950 [32]. Here, the prior knowledge (Morison's Equation) takes the form:

Prior knowledge example A

$$f(X) = \frac{1}{2}\rho DC_d U|U| + \frac{1}{4}\pi\rho D^2 C_m \dot{U} \quad (9.1)$$

where ρ is the fluid density, D is cylinder diameter, U is water particle velocity, \dot{U} is water particle acceleration, C_d is the drag coefficient and C_m is the inertia coefficient [36, 222].

Navigating the flowchart in Figure 9.5 for this example, takes the following route:

Q1. *Is there knowledge of appropriate inputs to use within the model?*

A1. Yes. The water particle velocity and acceleration were known to be appropriate inputs for wave loading prediction models.

Q2. *Does knowledge extend beyond selection of appropriate model inputs?*

A2. Yes. Morison's Equation is utilised here.

Q3. *Does knowledge extend beyond bounding variables between extreme limits?*

A3. Yes. Morison's equation approximates the wave load.

Q4. *Is any knowledge in functional form?*

A4. Yes. Morison's Equation is a function.

Q5. *Does the function aim to restrict or reduce valid solutions to a problem?*

A5. No. Morison's Equation does not constrain what values for wave load are or aren't valid.

Q6. *Does the function approximate any part of the behaviour of the target variable?*

A6. Yes. Morison's equation approximates the wave load, the target of the model.

Q7. *Does the function provide an approximation of the whole process?*

A7. Yes. Morison's equation approximates the whole wave load.

Q8. *Is the fitting of the function disrupted by the presence of unknown processes?*

A8. No. When Morison's Equation was fit to the data, an NMSE of $< 20\%$ was achieved on an unseen test set. This is inline with expected performance and implied Morison's equation was working as well as could be expected [257]. This suggests that the presence of processes such as vortex shedding, not captured by Morison's Equation did not impede or disrupt its fitting.

- **Suggested Model: Residual model with a fitted mean function and a flexible kernel.**

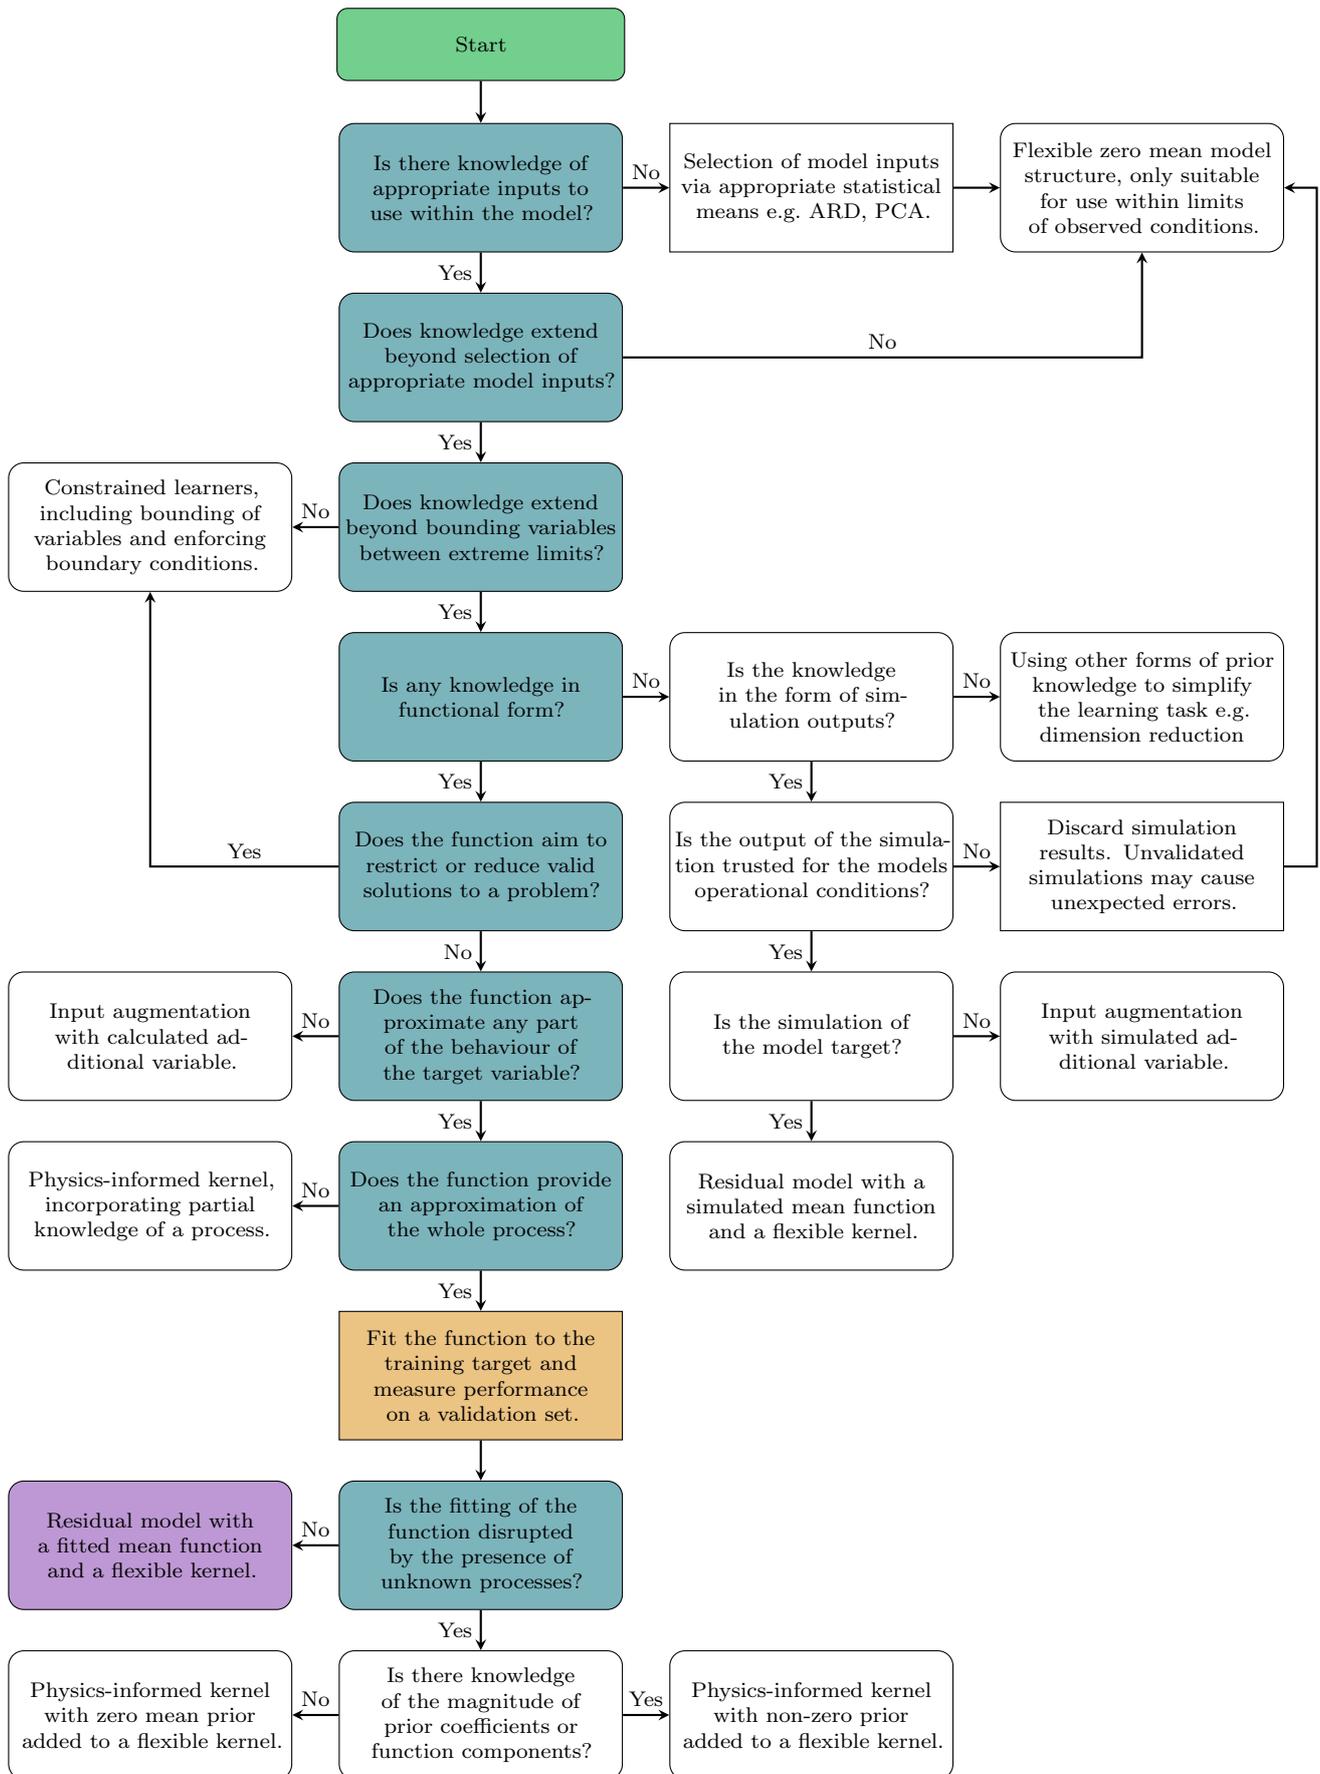


Figure 9.5: Application of the flowchart for the incorporation of Morison's Equation within a Gaussian process. This work is presented in Chapter 5. Teal nodes represent decisions, orange nodes represent actions and purple nodes represent outcomes.

In this case, the suggested model of the flowchart aligned with the best performing PIML model structure within Chapter 5, with the implemented structure outperforming a black-box model. Particular improvements were seen where access to measured data was limited and models relied on their capability to extrapolate. The chapter also investigated a case where Morison's Equation was used for input augmentation, with a residual model structure found to be preferable, further supporting the alignment of the flowchart decision.

In terms of flowchart navigation, the most challenging section was Q8 '*Is the fitting of the function disrupted by the presence of unknown processes?*'. The answer to this question is arguably the most subjective, with the answer determined by whether a satisfactory fit of Morison's Equation was achieved. Here, due to the model obtaining an NMSE of $< 20\%$ on an unseen validation set, Morison's Equation was deemed to be working as expected [257], and not disrupted by the presence of unknown processes. If the answer to this question was instead deemed to be 'yes', with an alternate route of the flowchart taken, the suggestion of the flowchart would have been a physics-informed kernel added to a flexible kernel. This model structure was implemented for Morison's Equation in the work within Chapter 7 and was also found to be an effective choice of PIML structure. The additive kernel structure and residual model were compared within Chapter 7 and found to perform similarly. The residual model achieved a slightly improved NMSE (14.797% vs 14.913%) whilst the additive kernel structure achieved an improved MSL (-0.939 vs -0.598). As models become more similar, in terms of both capability and performance, their separation within the framework becomes more challenging. An advantage of this however, is that though questions may become more subjective later in the framework, the consequences of decisions appears to reduce. In this case, due to comparable performance and model capabilities, either model structure would be deemed appropriate in the author's opinion.

9.2.2 Example B: Linear wave theory

The second application example for the flowchart looks at aspects of linear wave theory, which were integrated within PIML models in Chapter 6. Multiple useful quantity expressions were used in this chapter, including the free surface η , water particle velocities U_x, U_z and water particle accelerations a_x, a_z [35, 36]. A common theme however is that they all relied upon the linear solution for velocity potential. Though the multiple expressions used may be referred to in Chapter 6, only the expression for velocity potential is given here:

Prior knowledge example B

$$\phi = \sum_{i=1}^N \frac{A_i g \cosh(k_i(z+d))}{k_i c_i \cosh(k_i d)} \sin(k_i x - \omega_i t + \Phi_i)$$

where A is wave amplitude, $k = \frac{2\pi}{\lambda}$ is the wave number, $c = \frac{\omega}{k}$ is the wave speed and Φ is the phase. The horizontal and vertical water particle velocities are negative spatial derivatives of the velocity potential and their calculation for a given range of x and z allowed for the construction of a velocity field. Deriving w.r.t t provided access to the water particle accelerations a_x and a_z .

Navigating the flowchart in Figure 9.6 for this example, takes the following route:

Q1. *Is there knowledge of appropriate inputs to use within the model?*

A1. Yes. The models in this chapter were designed to use incoming wave height as an input, a commonly measured variable on many offshore structures.

Q2. *Does knowledge extend beyond selection of appropriate model inputs?*

A2. Yes. Multiple aspects of linear wave theory is utilised here.

Q3. *Does knowledge extend beyond bounding variables between extreme limits?*

A3. Yes. Linear wave theory is able to approximate many important wave quantities, including, velocity, free surface and pressure.

Q4. *Is any knowledge in functional form?*

A4. Yes. Linear wave theory has multiple functions.

Q5. *Does the function aim to restrict or reduce valid solutions to a problem?*

A5. No. Though assumptions must be relied upon when utilising linear wave theory (e.g. no temperature gradients, small amplitude waves), the equations themselves do not constrain what values for each respective quantity are or aren't valid.

Q6. *Does the function approximate any part of the behaviour of the target variable?*

A6. No. Linear wave theory doesn't approximate any component of the wave load.

- **Suggested Model: Input augmentation with calculated additional variable.**

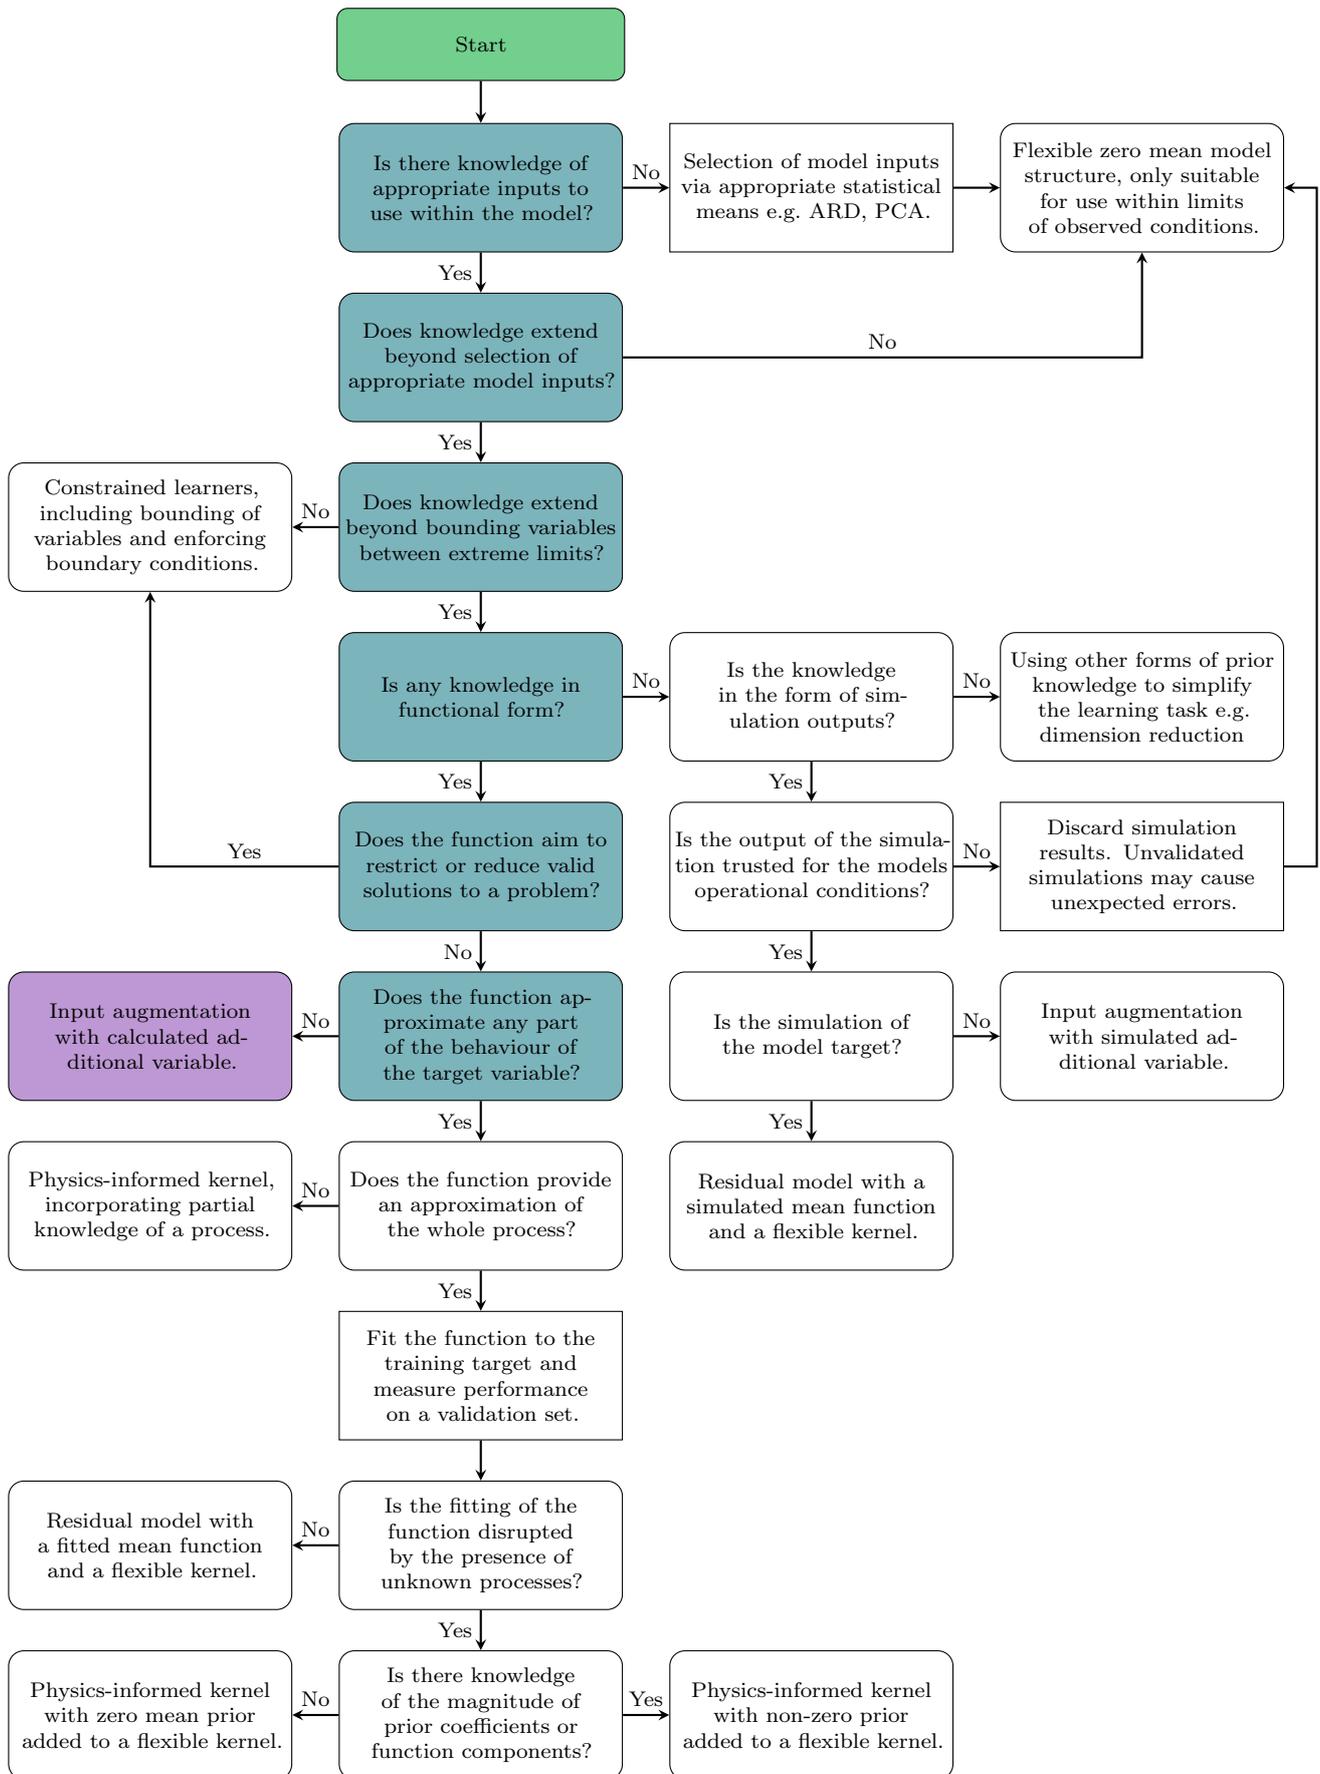


Figure 9.6: A flowchart for the inclusion of the aspects of linear wave theory, utilised in Chapter 6, within a Gaussian process. Teal nodes represent decisions, orange nodes represent actions and purple nodes represent outcomes.

The suggestion of the flowchart again aligned with the findings of the thesis, with the calculated flow conditions close to the monopile found to be a very valuable model input. The constructed models in Chapter 6 were able to achieve better performance across the full range of wave states generated within the wave tank compared with a purely data-based approach. The suspected reason for this performance improvement was that the relationship between the approximated flow conditions and the wave load was an easier learning task than that between incoming wave height and the wave load.

For the navigation of the flowchart in this case, the most challenging aspect was that multiple expressions could potentially be considered when answering each question. To obtain flow conditions close to the monopile, expressions for the velocity potential and free surface were also required. In this case however, answering the questions with respect to any of the utilised expressions would result in the same suggested model architecture. For scenarios where a user might be incorporating multiple instances of prior physical knowledge within the same model, there are two main possible ways to utilise the flowchart; either the prior knowledge is treated as a singular group and questions answered to most reflect the knowledge as a whole, or each piece of knowledge is passed through the flowchart individually. For the case of linear wave theory, the former of these options seemed most appropriate.

9.2.3 Example C: Free vibration of a cantilever beam

Part of the work within Chapter 7 focussed on modelling the response of a cantilever beam in free vibration. The free vibration response of a beam can be found through separation of variables with a solution of the form $Y(x, t) = W(x)T(t)$. The response of a beam in free vibration may be calculated through the superposition of normal modes [244]. The prior knowledge included was the form of the oscillatory, time dependent, behaviour of the beam $T(t)$, whilst no form of the mode shape $W(X)$ was assumed:

Prior knowledge example C

$$Y(x, t) = \sum_{i=1}^{\infty} W_i(x) e^{-\zeta \omega_n^i t} (A_i \cos(\omega_n^i t) + B_i \sin(\omega_n^i t))$$

where A_i and B_i are determined from the initial conditions of the beam. For many sets of known boundary constraints, including cantilever beams, it is possible to

derive an exact analytical expression for the mode shapes $W_i(x)$ [243, 244]. This was not done within the work of Chapter 7, to allow for an increased level of model flexibility.

Navigating the flowchart in Figure 9.7 for this example, takes the following route:

Q1. *Is there knowledge of appropriate inputs to use within the model?*

A1. Yes. The response of the beam $Y(x, t)$ was known to be a function of both position x and time t .

Q2. *Does knowledge extend beyond selection of appropriate model inputs?*

A2. Yes. The form of the oscillatory behaviour of the beam, $T(t)$ was known.

Q3. *Does knowledge extend beyond bounding variables between extreme limits?*

A3. Yes. $T(t)$ approximates how a fixed location on the beam oscillates over time for a given mode.

Q4. *Is any knowledge in functional form?*

A4. Yes. An expression for $T(t)$ is known.

Q5. *Does the function aim to restrict or reduce valid solutions to a problem?*

A5. No. The expression for $T(t)$ does not constrain what values for beam response are or aren't valid.

Q6. *Does the function approximate any part of the behaviour of the target variable?*

A6. Yes. $T(t)$ approximates the oscillatory, time dependent, behaviour of the beam's response.

Q7. *Does the function provide an approximation of the whole process?*

A7. No. The expression for $T(t)$ only provides an approximation of the time dependant component of system behaviour. It does not provide an approximation of the full response $Y(x, t) = W(x)T(t)$.

- **Suggested Model: Physics-informed kernel, incorporating partial knowledge of a process.**

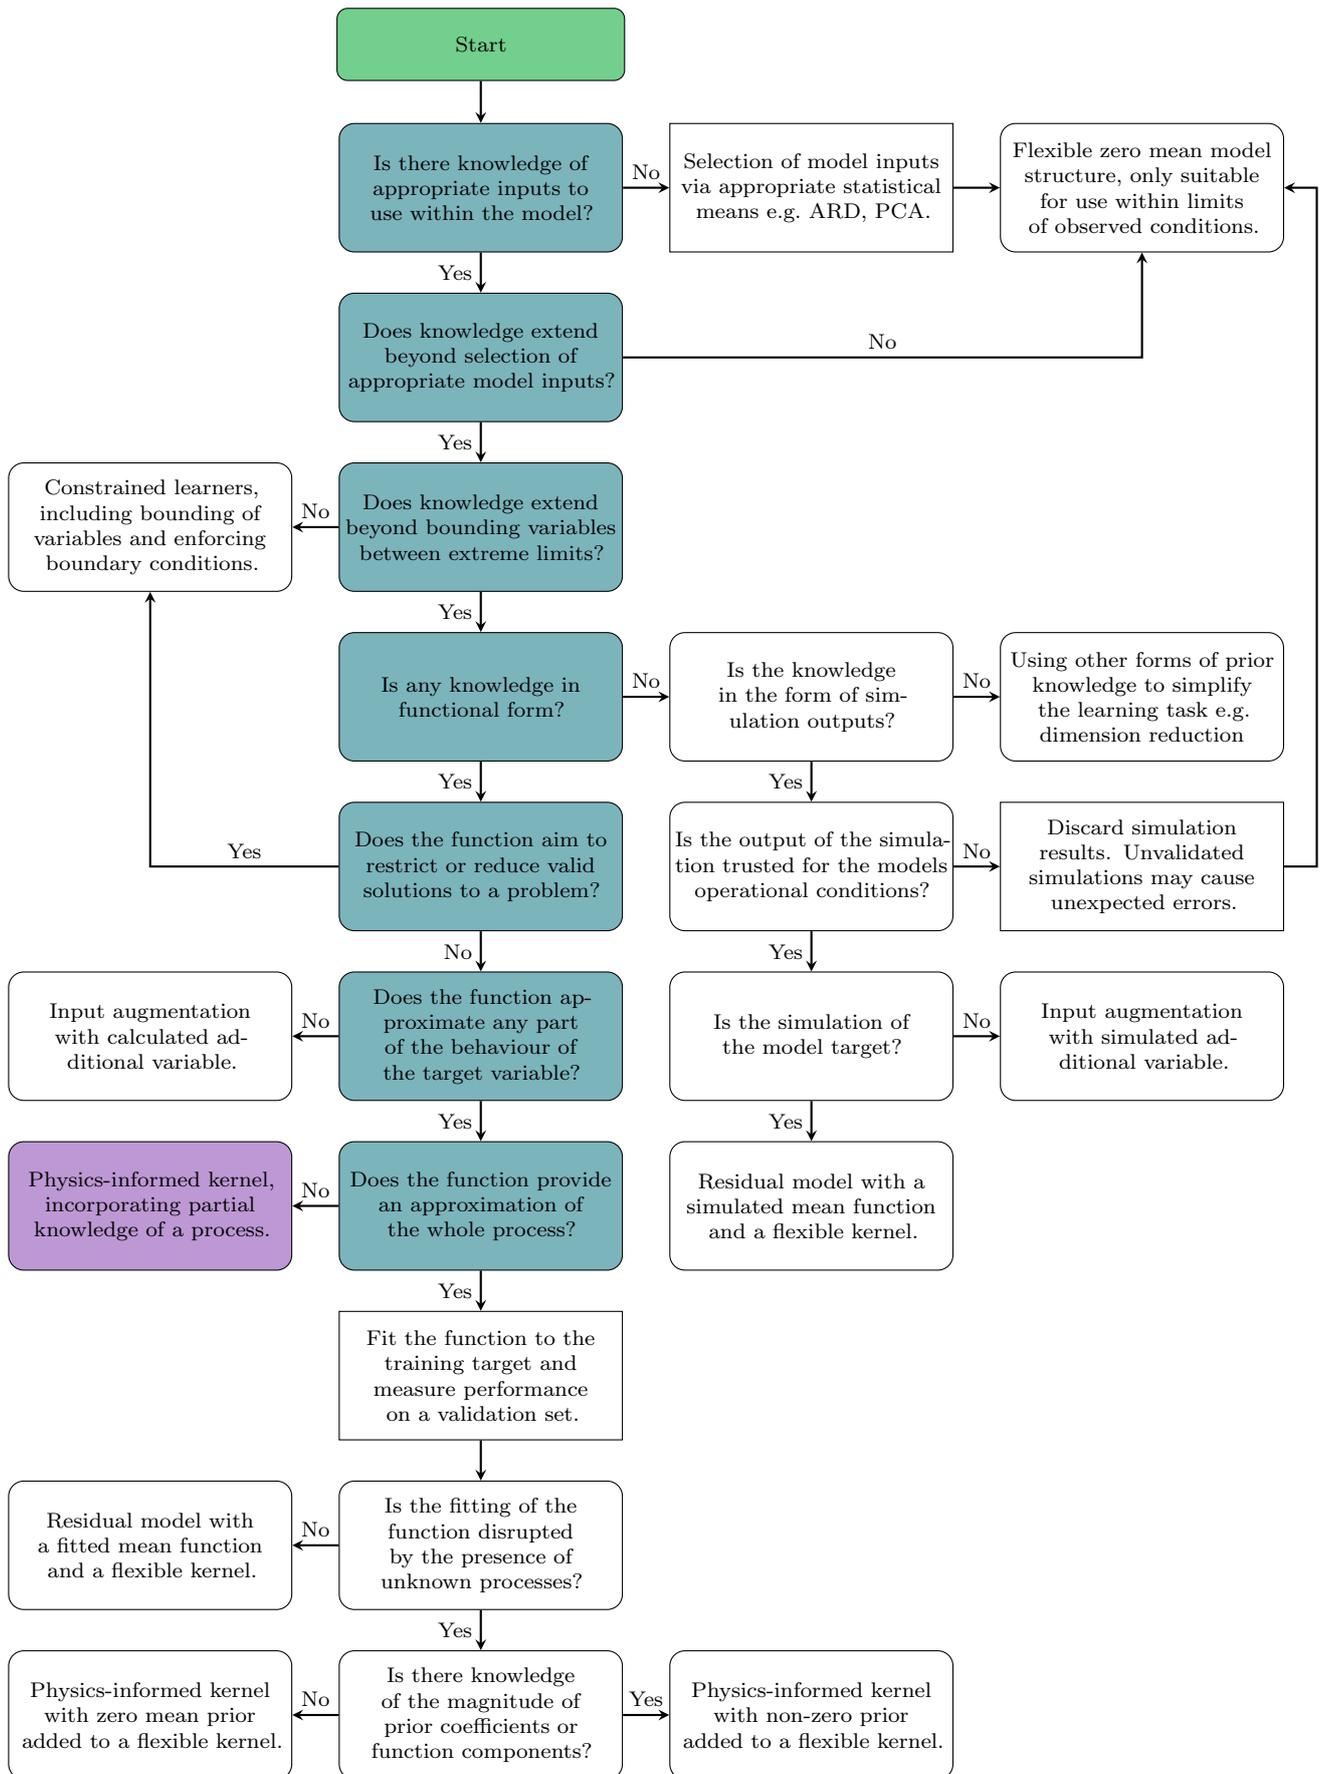


Figure 9.7: A flowchart for the inclusion of the oscillatory behaviour of a cantilever beam $T(t)$ within a Gaussian process. This work was implemented in Chapter 7. Teal nodes represent decisions, orange nodes represent actions and purple nodes represent outcomes.

The suggested model structure by the flowchart for the case of the cantilever beam in free vibration aligned with the findings in Chapter 7. The developed model in this chapter was able to incorporate knowledge of the oscillatory behaviour of a cantilever beam without strict enforcement of boundary conditions specific to the analytical expressions for mode shapes $W(x)$. These included uniform cross-section of the beam, a perfectly fixed base and uniform mass distribution [243, 244].

An important point of note here was that the suggested model structure of ‘*Physics-informed kernel, incorporating partial knowledge of a process.*’ is a broad category. To develop an appropriate PIML model from this point still required a large amount of research and model development effort. Firstly, the physically derived SDOF kernel [149], was utilised to mimic the oscillatory behaviour of the beam. This was then incorporated within a wider multiplicative kernel structure of the form:

$$K(\tau_x, \tau_t) = \sum_{i=1}^N \underbrace{K_{SE}^i(\tau_x)}_{\text{Mode shape}} \underbrace{K_{SDOF}^i(\tau_t)}_{\text{Oscillatory behaviour}} + \underbrace{\sigma_n^2 \delta_{ij}}_{\text{Noise}} \quad (9.2)$$

where the squared exponential kernel K_{SE}^i was used to model the mode shape for a given mode i . The construction of the combined kernel reflected the structure of the solution for free vibration response $Y(x, t) = W(x)T(t)$. Each of these steps required an additional degree of insight to the problem beyond the aid provided by the framework. However, the suggestion of the framework would provide a helpful starting position from which to begin making such developments.

9.3 Conclusions

The common types of physical knowledge that one might wish to incorporate within a physics-informed engineering model were discussed and categorised according to how they were obtained and the form which they take. This is argued as a useful first step a user might take before constructing a PIML model. An improved understanding of how to define a type of physical knowledge should make the answering of questions based upon the knowledge less subjective. A consideration of how difficult, or expensive, any given piece of knowledge is to acquire is considered as a factor for its use within a PIML model.

A framework to aid with decision-making on how best to incorporate a piece of prior knowledge within a PIML model was developed, in the form of a flowchart.

Examples of models developed within the thesis were used to check the alignment of flowchart suggestions with successful modelling decisions. In all cases, without assuming any knowledge of the adopted model structure, the flowchart suggestion aligned with the findings of the thesis.

To date, to the author's knowledge, there are no other examples of attempts to construct a similar framework. Deciding how best to integrate physical knowledge within a machine learner is typically done using the personal expertise of the modeler, researcher or engineer constructing the model. These decisions can be challenging, even for those with appropriate expertise, and efforts to alleviate the difficulty of the decision have the potential to be a valuable resource.

Although several iterations of the framework have been produced, and efforts have been made to test that the suggestions of the framework are sensible, it is not the intention of the author to present the current framework as an absolute truth. It aims to serve as a useful tool to help a user, but never to override an individual's expertise for a given modelling task. There are many nuances to the construction of PIML models, and it is very possible that exceptions to the rules within the framework might be found. An important area for future work is further testing and development of the framework, preferably with those away from the author's own research group. The insight and different perspectives of those within industry or different research groups would be extremely valuable for future versions, and would help to strengthen the current edition.

CONCLUSIONS AND FUTURE WORK

This thesis has developed and applied a range of PIML model architectures within the field of structural dynamics and studied how they performed, the capabilities of changing model structures and potential benefits brought to real-world applications. A framework to aid users with the construction of PIML models for a given modelling scenario concluded the thesis. At the beginning of the thesis, there were two primary goals outlined:

- 1. To develop novel grey-box models for structural dynamics that combine physics and data-based learning with the intention to exploit the benefits of either approach used independently.*
- 2. To explore and formulate the relationship between the type of prior physical knowledge possessed by a user and how best to incorporate it within a model.*

As will now be discussed within the upcoming chapter contributions, the author believes that these aims have been achieved within the thesis. **Chapters 3 to 7** focussed primarily on achieving the first aim, whilst **Chapters 2, 8 and 9** addressed the second. The adoption of PIML within structural dynamics is by no means complete however, with remaining potential challenges then discussed alongside areas for future work.

10.1 Summary of contributions

Chapter 2 introduced a range of physics-informed machine learning methods to the reader, with examples of model structures from literature. Methods were presented in order of progression through a typical black-box machine learning pipeline, from the initial definition of a problem to a complete functioning model, addressing how physics may be incorporated at each stage. The categorising and presentation of methods in this way is novel and will be a helpful resource to a user wishing to modify an existing black-box model architecture. The specific changes made to a typical process, e.g. modification of the cost function, and the resulting effects on modelling capabilities were discussed.

Chapter 3 provided introductions and the author's interpretations of relevant theory from both a data-based and physics-based perspective. Gaussian process models form the data-based component of later developed models and were therefore a major focus. Instead of a traditional mathematical introduction, for which many high quality examples exist within the literature [139, 215, 216], model interpretation through figures, particularly during their training stage was the main aim here. One of the major potential advantages of PIML models, is increased interpretability into model results and an understanding of model behaviours is therefore important.

From a physics-based perspective, linear wave theory, the representation of sea states and empirical wave loading prediction methods were introduced. These are all widely used aspects of physical theory used within offshore applications. It was important to introduce readers to the types of physical knowledge used in the models developed within the thesis. For readers already familiar with physical methods used within offshore applications, this chapter serves to collate the authors understanding of a selection of key methods.

Chapter 4 detailed the two datasets primarily used to develop PIML models within the thesis: a subset from the Christchurch Bay Tower (CBT) project, a heavily instrumented column structure within a real offshore environment; and an experimental study of a monopile within a wave tank devised and carried out by the author. Although both useful for the study of offshore structures, each dataset served its own distinct role. A section of data the Christchurch Bay Tower project provided access to measurements from a real world structure in an offshore environment. The experimental dataset provided a wider range of conditions that were more tightly controlled by the user, allowing for the evaluation of model performance in conditions they hadn't previously seen.

The experiment itself is a novel contribution here, with measurements of wave

load on structures very rare. This is primarily due to the difficulty and expense of creating the bespoke equipment required. Again, the author expresses their thanks to Robin Mills and Michael Dutchman for their work designing and manufacturing the force collar. An effort was made to measure more variables and generate more conditions than would be needed for this thesis, with the hope that this dataset will be useful for the future work of myself and others¹.

Chapter 5 presented the first development of a grey-box model for the prediction of wave loads on a real offshore structure. It was developed using the Christchurch Bay Tower dataset detailed in Chapter 4. The widely used Morison's Equation, an empirical wave loading solution, was incorporated within an autoregressive form of Gaussian process Regression (GP-NARX). Important details of implementation for GP-NARX models are addressed including lag selection, uncertainty propagation and cost function selection, all of which have an impact on final model performance. A specific study focussed on the ability of different grey-box model architectures to extrapolate in circumstances of low data coverage. This was a key expected benefit of incorporating physical knowledge within a machine learner.

The combining of physics-based white-box and data-based black-box modelling techniques in the form of two grey-box architectures was found to offer benefits in predictive performance over either approach used alone. The best performing grey-box model, the residual modelling GP-NARX, achieved a 29.13% and 5.48% relative reduction in NMSE over Morison's Equation and a black-box GP-NARX respectively. It was expected that grey-box models would be of specific help in assisting with extrapolation, an area in which data-based methods typically experience difficulty, and this was indeed found to be the case. Residual modelling was found to offer significant benefits in performance outside the range of observed training conditions, particularly in instances where a pre-established white-box may be available for inclusion into the combined model.

Chapter 6 built upon the grey-box model development of Chapter 5, in which models relied on measurements of flow conditions close to a structure to predict the wave load. Here, aspects of linear wave theory were integrated within a GP-NARX framework to remove this reliance. Utilising physical knowledge, models were developed which relied on only incoming wave height as an input, a commonly available variable on many offshore structures. Models were developed using data from the monopile experiment detailed in Chapter 4. The use of a GP-NARX, Morison's Equation and linear wave theory in combination within a physics-informed model was able to offer increased performance over any of its components used alone. The model was also able to offer a degree of interpretability, by providing a

¹The experimental dataset will be made publicly available in due course.

breakdown of the wave load into drag, inertia and diffraction components. Model results were able to be visualised and interpreted easily, along with quantification of uncertainty.

Chapter 7 focused on the means of incorporation of prior physical knowledge within the kernels of Gaussian processes. Modelling scenarios where one might have partial knowledge of a system or process were investigated through the combination of physics-informed kernels and more flexible ‘data-based’ kernels. Case studies of kernel structures were presented with application to wave loading prediction and a vibrating cantilever beam. Having partial knowledge of a system or process is common within engineering applications and this is discussed with reference to the data vs physics chart in Figure 7.7 and how this might affect one’s position on it.

A kernel was derived for Morison’s Equation, showing how one might use the covariance of physical processes within Gaussian processes. Following a similar approach could be potentially be effective across a wide range of applications. Using the derived Morison kernel in combination with a more flexible Squared Exponential (SE) kernel was found to increase predictive performance over either Morison’s Equation or a (SE) kernel used independently. The decomposition of combined models was found to offer physical insight into the role of each model component. Here, due to the flow conditions, the majority of structure was captured via the physics-based component.

The idea of an optimum level of prior knowledge to be included within a model changing depending on the modelling scenario was highlighted within the context of a cantilever beam in free vibration. The benefits of achieving this optimum level such as improved extrapolation, interpretability and reduced reliance on data collection were discussed. Consequences of being incorrectly placed on the scale of prior knowledge inclusion were also presented. These included reduced flexibility and potential reliance on unrealistic assumptions for a physics-heavy model and poor extrapolation, interpretability and increased demand for data collection for a data-heavy model.

Chapter 8 explored how changing Gaussian process model structure affected the capabilities on properties of a constructed model. A consistent piece of prior knowledge, in the form of a ‘*linear-in-the-parameters*’ function, $f(X) = X\beta$ was incorporated in a range of ways through modification of both mean and covariance functions, with findings compared. An example was considered where a linear function represented some known aspect of a system’s behaviour, with the remainder captured by a flexible kernel; this allowed the maths for all examples to be worked through completely, whilst highlighting differences more clearly. A toy dataset was then used to compare the performance of models in circumstances of limited

training data, an important area of focus for grey-box models. The work of this chapter is early in its field and provided an important platform from which to construct the decision framework presented in Chapter 9.

Chapter 9 investigated important relationships between the form in which knowledge presents itself (e.g. equations, simulations, bounds), the source from which it comes (e.g. domain experts), its cost of accrual and how best to incorporate it within a model. A proposed framework was developed to assist with deciding how best to integrate prior knowledge within Gaussian Process Regression (GPR) models. The aim is to help classify a given piece of prior knowledge and suggest an appropriate model architecture.

Types of physical knowledge that one might wish to incorporate within a physics-informed engineering model were discussed and categorised according to how they were obtained and the form which they take. This was a useful exercise and brought structure to the process of constructing a PIML model. The first stage of the process is argued as understanding and defining the physical knowledge possessed by a user. Only after this, and the selection of a data-based modeller, is it deemed sensible to begin deciding how to integrate the two. This process is also touched on in the beginning of **Chapter 5**, when beginning to construct the first PIML model.

The testing of a framework for PIML model construction concludes both **Chapter 9** and the thesis. The framework, in its current state, presents a concise summary of the authors findings and experience of what has (and hasn't) worked effectively for PIML modelling within engineering applications. Role-playing as a modeller, with access to a given piece of prior knowledge, the author was able to answer yes-no questions (e.g. *Does knowledge extend beyond bounding variables between extreme limits?*) and arrive at a PIML model structure that aligned with the findings of the thesis. For PIML modelling generally, as well as for engineering applications, guidance for model construction based on an authors modelling scenario is currently rare, with the decision based largely on the modellers expertise and judgement. The framework hopes to serve as a useful contribution toward this area, whilst being able to function as a sense check for modellers who might already be confident in their choices.

The contribution of this thesis as a whole aims to progress the adoption and development of PIML models for the fields of structural dynamics, structural health monitoring, offshore structures and wave loading prediction. Though impactful early steps have been taken, including the development of novel PIML models, generation of experimental datasets and creation of a decision framework, the task of industry-ready PIML model adoption is by no means complete. In the coming

section, directions for areas of future work, along with remaining key challenges of the field are introduced.

10.2 Future work

The two primary areas for future work identified relate closely to the two aims of the thesis presented at the beginning of this chapter. Continued PIML model development is an extension of the first aim, whilst decision-making for PIML builds upon the work of the second.

10.2.1 Continued PIML model development

With a key focus of the thesis being the development of PIML model architectures, a natural progression of work is a continuation of this. Introduced here are some potential directions.

Dynamic balancing of physics and data

A theme throughout the thesis was a consideration of achieving ‘the right balance’ between physics and data within a model that reflects a users given modelling scenario. Intuitively, a high availability of physical knowledge, with limited access to measured data would suggest a physics-dominant model, whilst the opposite would be true for a data-dominant case. Particular focus was given to this in Chapter 7 where examples of partial knowledge of a system behaviour were represented via kernel design.

A potentially interesting area of research is how one might be able to tune the relative reliance of physical knowledge and data within a model dynamically. There are many reasons why one might want to change the relative reliance on physical knowledge within a model, an important one of which is the changing validity of a physics-based model. With any physical model, and particularly so with simple ones, assumptions must be made in order to represent the system of interest. The extent to which these assumptions hold affects the performance of the constructed model and care should be taken not to trust the results of models constructed upon invalid assumptions. Assigning a fixed degree of trust within a physical model component may be challenging when a model is required to operate over a range of conditions and allowing this to vary is therefore highly desirable. The presence of regime-switching and localised behaviours provide alternative motivation to vary

the reliance on physical knowledge. If a phenomena is known to occur in specific conditions, for example the dependency of vortex shedding on flow speed [225], one might wish to phase its occurrence in and out of a model.

Some ongoing investigations into switching model structures utilising kernel design have shown to be one potential area for tuning the reliance on physics and data within the same PIML model. Further work on this could look toward the many existing kinds of switching model, including Mixtures of Experts (MOEs) [287, 288], Markov-switching models [289], Treed Gaussian Processes (TGPs) [290] and regime-switching cointegration models [291].

Incorporation of higher fidelity physical knowledge

Through choice, the PIML models developed within this thesis typically relied upon familiar, widely utilised examples of physical knowledge as the physics-based component of the model. These included Morison’s Equation, linear wave theory and the free vibration response of a cantilever beam. Well understood instances of physical knowledge, for which the limitations and assumptions were well known made the suitable scope of PIML model operation more clear.

The inclusion of higher fidelity physics-based components could be one potential way to effectively construct ‘physics-dominant’ model architectures. When a simpler piece of physical prior knowledge is incorporated within a PIML model, it is important to understand the limitations of the assumptions present within the model and understand in which regimes the model is able to extrapolate effectively. For example, one of the assumptions present within the adoption of linear wave theory is that the waves being modelled have small amplitudes, relative to wave length and water depth. For higher frequency, tall waves in shallower water, this assumption may begin to break down. Extensions to linear wave theory, including Stokes Waves [276] have been shown to be more effective at capturing the shape or wind generated offshore waves, which typically have sharper peaks and steeper troughs than those generated by laboratory wave tank paddles [36].

10.2.2 Decision-making for PIML

The framework to aid with the construction of physics-informed Gaussian process models presented in Chapter 9 aims to help tackle an important challenge within the field of PIML - ‘*How should a given piece of physical knowledge be effectively incorporated within a PIML model?*’. Although a start has been made, this remains an ongoing challenge, with some areas of focus for future work.

Progression of the framework

In its current state, the framework navigates a user with access to a given piece of physical knowledge to an appropriate type of model structure. Currently, examples of these model structures include suggestions such as ‘*Residual model with a fitted mean function and a flexible kernel.*’ or ‘*Input augmentation with simulated additional variable.*’. For the cases tested against the content of this thesis, the suggestions of the framework aligned with effective PIML model architectures.

One potential area for extension on the framework, is an ability to provide more specific guidance beyond the suggested model structures. For some suggested architectures, including constrained learners and the incorporation of partial knowledge within the kernel, there are multiple ways in which one could approach this. Further guidance, for example, how specifically one might constrain a GP would be a helpful additional resource. An ability to provide more specific suggestions for the structure of kernels, for instance, whether to use an additive or multiplicative kernel structure would also be a valuable area for exploration.

Although possible, an example of the flowchart suggesting an inappropriate model structure was not found within the work of this thesis, nor did navigation of the flowchart deviate away from effective implementations of PIML model architectures. The consequences of an incorrect navigation of the chart, through differing interpretations of questions or otherwise would be an interesting area for investigation. For example, if the ‘true’ most appropriate model structure was a physics-informed kernel, what would be the consequences of ending up with a suggestion of input augmentation? Firstly, would the suggested model structure be feasible for the available prior knowledge, and if so, how would the capabilities of the final model be affected?

Collaboration with the PIML community and industry

So far, the testing and application of the PIML framework has been done using the authors personal findings and examples of models developed within the thesis. An important extension to this would come from collaboration with others in the PIML community and industrial partners. A key area that this would help to test is the interpretability and navigation of the chart. Although attempts have been made to liaise with supervisors and other members of the research group² when constructing the chart, to test whether its questions are easily interpretable, opinions from outside the group would be very valuable. Through seeing multiple versions, and being familiar within the field, a biasing of how things might be

²Thank you to Matt, Sam, Lizzy and Tim for their readings of many multiple versions of the flowchart

viewed is also very possible.

A rigorous test of the framework would be the presentation of the framework to others in the PIML community to navigate the chart and compare with their findings of existing work. This could be done similarly to the examples in Chapter 9 without any external input from the author. Feedback on whether the flowchart aligned with their selected methods and suggestions for improving the interpretability of the chart would be very valuable. Performing a similar exercise with industrial partners could help make a revision of the flowchart more adapted for use within industrial settings, whether this involves changes to terminology or other modification.

10.3 Challenges for PIML in SHM

10.3.1 Dealing with large quantities of data

One of the major potential benefits of PIML modelling is the ability to utilise physical knowledge to improve model performance in regions of sparse data coverage. A consequence of this is a reduced demand for data collection, with the importance of ensuring every available potential set of conditions has been previously measured reduced. With a purely data-based learner, making predictions on previously unobserved conditions is typically unreliable, and making operational decisions on extrapolative machine learning predictions would be ill-advised.

Although PIML modelling may reduce the demand for data collection, this does by no means suggest that the quantity of data being dealt with would be small. Firstly, it is still beneficial for models to have access to as wide a range of conditions as is feasible. Although possessing an enhanced ability to extrapolate, it remains preferable for PIML models to operate in interpolation wherever possible. For large structures, with dense sensor networks, measurement campaigns can quickly generate extremely large quantities of data, that needs to be managed, stored and handled appropriately. An important topic not focussed on within the content of this thesis how to effectively scale the PIML models developed to work with much higher quantities of data. As is necessary for research and the trialing of new ideas, the models within the thesis were instead developed using smaller scale datasets, that were able to be easily stored on a personal desktop computer. For example, the file size of the complete experimental dataset was of the order $\sim 20\text{GB}$, which could be many magnitudes smaller than the file sizes of complete industrial datasets.

Along with appropriate methods of handling and storing large quantities of data

efficiently, the reliance on Gaussian process regression for PIML models developed within the thesis will require some adjustment of methodology. For a number of data points n within a dataset X , a GP will have requirements of the order:

- $O(n^2)$ for the storage of the covariance matrix $K(X, X)$
- $O(n^3)$ for the computation of the inverse of the covariance matrix $K(X, X)^{-1}$

In the authors experience, with access to moderately powerful computational resources (64GB RAM desktop), the training and storage of data for standard Gaussian process regression became infeasible above $\sim 10,000$ datapoint training set sizes. For effective implementation within industrial applications, one would likely want to exceed this quantity. To achieve this, methods could adopt sparse Gaussian process regression [267], which allows for computation of the order $O(nm^2)$ and storage of $O(nm)$. Hensman [249] showed how a combination of inducing variables and appropriate mini-batching could allow GP regression to be implemented on dataset sizes in the millions of points.

10.3.2 Operator confidence and benefit to the industry

The work of this thesis has highlighted many of the benefits of PIML modelling over either purely physics-based or purely data-based approaches, including: improved performance, enhanced insight into results and reduced reliance on bespoke sensor systems. Within academia, with reference back to the growth chart in Figure 2.1, the popularity and adoption of PIML is growing rapidly. Advances in methodology across a range of applications are continuing to be published. However, a very important distinction to make is the difference in perceptions between academia and industry. A key goal of this thesis is to help progress towards the adoption of PIML methods within real world SHM applications. For this reason, the effective communication of potential PIML benefits to industry is very important.

For the successful adoption of a PIML model within an industrial application, it must be shown that:

- The PIML model should work effectively within the realms of operational conditions and the results that it provides can be validated and trusted.
- The benefit provided by a PIML model (financial, humanitarian, operational or otherwise) outweighs the expense and effort spent to develop it.

Achieving operator confidence within results is a very important step for the adoption of any model, and particularly those containing a machine learning component. For safety critical structures, including bridges, aircraft and many offshore platforms, the decisions informed by the output of a model have severe potential consequences at stake. This might range from a highly expensive false shutdown, to an unnecessary structural failure, to the loss of life. This understandably, produces a degree of inertia when posed with an introduction to any new methodology. One of the potential advantages of PIML modelling over a purely data-based approach is that models are often able to offer a degree of insight into model results. For operators, particularly those familiar with existing physics-based approaches, the insight provided by similar approaches incorporated within a PIML model could be an effective way to communicate information about a models capabilities. The validation of PIML models in real world environments and the effective communication of this remains a key challenge.

An additional consideration when implementing new approaches within industrial settings is a cost-benefit analysis of what might be gained. Assuming the benefits of PIML modelling have been communicated, understood and quantified for a given application, which itself remains a challenge, it is also important to consider the cost of implementation. For the adoption of a PIML model, factors including personnel, hardware, software and changes to existing operational practise would have to be considered. For example, a deviation away from design-life approaches would require an overhaul of many aspects of how an engineering firm might operate. Quantification of both the cost and benefit of a PIML model is an important step before its implementation, to ensure a net-benefit is achieved.

BIBLIOGRAPHY

- [1] G. Holmes, P. Sartor, S. Reed, P. Southern, K. Worden, and E. J. Cross. Prediction of landing gear loads using machine learning techniques. *Structural Health Monitoring*, 15:568–582, 2016.
- [2] H. Sohn, C. R. Farrar, N. F. Hunter, and K. Worden. Structural health monitoring using statistical pattern recognition techniques. *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME*, 123:706–711, 2001.
- [3] I. Antoniadou, N. Dervilis, E. Papatheou, A. E. Maguire, and K. Worden. Aspects of structural health and condition monitoring of offshore wind turbines. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373, 2015.
- [4] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006.
- [5] D. Hawkins. The problem of overfitting. *Journal of Chemical Information and Computer Sciences*, 44:1–12, 2004.
- [6] K. Worden, C. Farrar, and G. Manson. The fundamental axioms of structural health monitoring. *Proceedings of The Royal Society A: Mathematical, Physical and Engineering Sciences*, 463:1639–1664, 2007.
- [7] E. J. Cross. *On Structural Health Monitoring in Changing Environmental and Operational Conditions*. PhD thesis, University of Sheffield, 2012.
- [8] S. Zhang, T. J. Rogers, and E. J. Cross. *Gaussian Process Based Grey-Box Modelling for SHM of Structures Under Fluctuating Environmental Conditions*, pages 55–66. 2021.

- [9] E. J. Cross, T. J. Rogers, D. J. Pitchforth, S. J. Gibson, S. Zhang, and M. R. Jones. A spectrum of physics-informed Gaussian processes for regression in engineering. *Data-Centric Engineering*, 5:e8, 2024.
- [10] L. von Rueden, S. Mayer, K. Beckh, B. Georgiev, S. Giesselbach, R. Heese, B. Kirsch, M. Walczak, J. Pfrommer, A. Pick, R. Ramamurthy, J. Garcke, C. Bauchhage, and J. Schuecker. Informed machine learning - a taxonomy and survey of integrating prior knowledge into learning systems. *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [11] B. Sohlberg and E. Jacobsen. Grey box modelling – branches and experiences. *IFAC Proceedings Volumes*, 41:11415 – 11420, 2008.
- [12] G. Karniadakis, Y. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang. Physics-informed machine learning. *Nature Reviews Physics*, pages 1–19, 2021.
- [13] C. R. Farrar and K. Worden. *Structural health monitoring: a machine learning perspective*. John Wiley & Sons, 2012.
- [14] C. R. Farrar and K. Worden. An introduction to structural health monitoring. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365:303–315, 2007.
- [15] A. Rytter. *Vibrational Based Inspection of Civil Engineering Structures*. PhD thesis, Dept. of Building Technology and Structural Engineering, University of Aalborg, Denmark, 1993.
- [16] K. Worden and J. M. Dulieu-Barton. An overview of intelligent fault detection in systems and structures. *Structural Health Monitoring*, 3:85–98, 2004.
- [17] G. Zhang, Y. Liu, J. Liu, S. Lan, and J. Yang. Causes and statistical characteristics of bridge failures: A review. *Journal of Traffic and Transportation Engineering (English Edition)*, 9:388–406, 2022.
- [18] S. Findlay and N. Harrison. Why aircraft fail. *Materials Today*, 5:18–25, 2002.
- [19] S. P. Baker, J. E. Brady, D. F. Shanahan, and G. Li. Aviation-related injury morbidity and mortality: data from us health information systems. *Aviation, space, and environmental medicine*, 80:1001–1005, 2009.
- [20] M. Lichtveld, S. Sherchan, K. B. Gam, R. K. Kwok, C. Mundorf, A. Shankar, and L. Soares. The deepwater horizon oil spill through the lens of human

- health and the ecosystem. *Current environmental health reports*, 3:370–378, 2016.
- [21] M. Martinez-Luengo and M. Shafiee. Guidelines and cost-benefit analysis of the structural health monitoring implementation in offshore wind turbine support structures. *Energies*, 2019.
- [22] I. N. Giannakeas, Z. S. Khodaei, and M. H. F. Aliabadi. Structural health monitoring cost estimation of a piezosensorized aircraft fuselage. *Sensors (Basel, Switzerland)*, 22, 2022.
- [23] S. Chandrasekaran. *Structural health monitoring with application to offshore structures*. World Scientific, 2019.
- [24] G. Ersdal and E. Hörnlund. Assessment of offshore structures for life extension. In *Proceedings of the International Conference on Offshore Mechanics and Arctic Engineering - OMAE*, volume 5, 2008.
- [25] Q. Liu, Y. Sun, and M. Wu. Decision-making methodologies in offshore wind power investments: A review. *Journal of Cleaner Production*, 295:126459, 2021.
- [26] A. Hughes, R. Barthorpe, N. Dervilis, C. Farrar, and K. Worden. A probabilistic risk-based decision framework for structural health monitoring. *Mechanical Systems and Signal Processing*, 150:107339, 2021.
- [27] P. Stansell. Distributions of freak wave heights measured in the north sea. *Applied Ocean Research*, 26:35–48, 2004.
- [28] M. Martinez-Luengo, A. Kolios, and L. Wang. Structural health monitoring of offshore wind turbines: A review through the statistical pattern recognition paradigm. *Renewable and Sustainable Energy Reviews*, 64:91–105, 2016.
- [29] A. Kharade and S. v. Kapadiya. Offshore engineering: An overview of types and loadings on structures. *Int. J. Struct. Civ. Eng. Res*, 3:16–28, 2014.
- [30] J. Jin and B. Meng. Computation of wave loads on the superstructures of coastal highway bridges. *Ocean Engineering*, 38:2185–2200, 2011.
- [31] K. M. T. Kleefsman, G. Fekken, A. E. P. Veldman, B. Iwanowski, and B. Buchner. A volume-of-fluid based simulation method for wave impact problems. *Journal of Computational Physics*, 206:363–393, 2005.
- [32] J. R. Morison, J. W. Johnson, and S. A. Schaaf. The force exerted by surface waves on piles. *Petroleum Transactions*, 189:149–157, 1950.

- [33] K. Ewans, G. Feld, and P. Jonathan. On wave radar measurement. *Ocean Dynamics*, 64:1281–1303, 2014.
- [34] M. Tucker. *Waves in Ocean Engineering: Measurement, Analysis, Interpretation*. Ellis Horwood, 1991.
- [35] J. Lighthill. *Waves in Fluids*. Cambridge University Press, 1978.
- [36] T. S. Sarpkaya. *Wave Forces on Offshore Structures*. Cambridge University Press, 2010.
- [37] A. Elhanafi. Prediction of regular wave loads on a fixed offshore oscillating water column-wave energy converter using cfd. *Journal of Ocean Engineering and Science*, 1(4):268–283, 2016.
- [38] Y.-H. Lin, J.-F. Chen, and P.-Y. Lu. A cfd model for simulating wave run-ups and wave loads in case of different wind turbine foundations influenced by nonlinear waves. *Ocean Engineering*, 129:428–440, 2017.
- [39] A. Aggarwal, H. Bihs, D. Myrhaug, and M. A. Chella. Characteristics of breaking irregular wave forces on a monopile. *Applied Ocean Research*, 90:101846, 2019.
- [40] M. Zhang, Z.-M. Yuan, S.-S. Dai, and A. Incecik. Development of a Novel Wave-Force Prediction Model Based on Deep Machine Learning Algorithms. In *International Ocean and Polar Engineering Conference*, 2020.
- [41] K. Worden, T. J. Rogers, and E. J. Cross. Identification of nonlinear wave forces using Gaussian process NARX models. In *Nonlinear Dynamics, Volume 1*, pages 203–221, 2017.
- [42] R. Ehsani Moghadam, M. Shafieefar, and H. Akbari. A probabilistic approach to predict wave force on a caisson breakwater based on Bayesian regression and experimental data. *Ocean Engineering*, 249:110945, 2022.
- [43] P. K. Stansby, L. C. Devaney, and T. Stallard. Breaking wave loads on monopiles for offshore wind turbines and estimation of extreme overturning moment. *Iet Renewable Power Generation*, 7:514–520, 2013.
- [44] J. V. D. Tempel, N. F. B. Diepeveen, W. E. D. Vries, and D. C. Salzmann. 15 - Offshore environmental loads and wind turbine design: impact of wind, wave, currents and ice. In *Wind Energy Systems*, pages 463 – 478. Woodhead Publishing, 2011.

- [45] A. Agarwal and A. Jain. Dynamic behavior of offshore spar platforms under regular sea waves. *Ocean Engineering*, 30:487–516, 2003.
- [46] *Assessment of the Morison Equation*. Woodward-Clyde Consultants, Houston, Texas, Defense Technical Information Center, 1980.
- [47] G. (GWEC). Global wind report 2024, 2024.
- [48] H. J. Tulleken. Grey-box modelling and identification using physical knowledge and Bayesian techniques. *Automatica*, 29:285–308, 1993.
- [49] L. von Rueden, S. Mayer, R. Sifa, C. Bauckhage, and J. Garcke. Combining machine learning and simulation to a hybrid modelling approach: Current and future directions. In *International Symposium on Intelligent Data Analysis (IDA)*, pages 548–560, 2020.
- [50] H. Hanachi, W. Yu, I. Y. Kim, J. Liu, and C. Mechefske. Hybrid data-driven physics-based model fusion framework for tool wear prediction. *The International Journal of Advanced Manufacturing Technology*, 101:2861–2872, 2019.
- [51] P. Pillai, A. Kaushik, S. Bhavikatti, A. Roy, and V. Kumar. A hybrid approach for fusing physics and data for failure prediction. *International Journal of Prognostics and Health Management*, pages 2153–2648, 2016.
- [52] M. Raissi, P. Perdikaris, and G. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, pages 686–707, 2019.
- [53] P. Sharma, W. T. Chung, B. Akoush, and M. Ihme. A review of physics-informed machine learning in fluid mechanics. *Energies*, 16, 2023.
- [54] Z. Zhang and C. Sun. Structural damage identification via physics-guided machine learning: a methodology integrating pattern recognition with finite element model updating. *Structural Health Monitoring*, 06 2020.
- [55] X. Jia, J. Willard, A. Karpatne, J. Read, J. Zwart, M. Steinbach, and V. Kumar. Physics-guided machine learning for scientific discovery: An application in simulating lake temperature profiles. 12 2019.
- [56] S.-M. Udrescu and M. Tegmark. Ai feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6, 2020.

- [57] S. Hosseinhashemi, Y. Zhang, C. Thon, and C. Schilde. Process insights with physics-inspired data-driven modeling- example of battery electrode processing. *Journal of Energy Storage*, 73, 2023.
- [58] M. Aramon, G. Rosenberg, E. Valiante, T. Miyazawa, H. Tamura, and H. G. Katzgraber. Physics-inspired optimization for quadratic unconstrained problems using a digital annealer. *Frontiers in Physics*, 7, 2019.
- [59] Z. Wei and X. Chen. Physics-inspired convolutional neural network for solving full-wave inverse scattering problems. *IEEE Transactions on Antennas and Propagation*, 67:6138–6148, 2019.
- [60] P. Lindskog and L. Ljung. Tools for semi-physical modeling. *IFAC Proceedings Volumes*, 27:1199 – 1204, 1994.
- [61] C. Halfmann, O. Nelles, and H. Holzmann. Semi-physical modeling of the vertical vehicle dynamics. In *Proceedings of the 1999 American Control Conference*, volume 3, pages 1707–1711, 1999.
- [62] T. Zubatiuk and O. Isayev. Development of multimodal machine learning potentials: Toward a physics-aware artificial intelligence. *Accounts of Chemical Research*, 54:1575–1585, 2021.
- [63] S. R. Vadyala, S. N. Betgeri, J. C. Matthews, and E. Matthews. A review of physics-based machine learning in civil engineering. *Results in Engineering*, 13:100316, 2022.
- [64] J. Liu, X. Han, Y. Pan, K. Cui, and Q. Xiao. Physics-assisted machine learning methods for predicting the splitting tensile strength of recycled aggregate concrete. *Scientific Reports*, 13, 06 2023.
- [65] E. J. Cross, S. J. Gibson, M. R. Jones, D. J. Pitchforth, S. Zhang, and T. J. Rogers. Physics-informed machine learning for structural health monitoring. In *Structural Health Monitoring Based on Data Science Techniques*, pages 347–367. Springer, 2022.
- [66] R. S. Geiger, K. Yu, Y. Yang, M. Dai, J. Qiu, R. Tang, and J. Huang. Garbage in, garbage out? do machine learning application papers in social computing report where human-labeled training data comes from? In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, page 325–336. Association for Computing Machinery, 2020.
- [67] R. S. Geiger, D. Cope, J. Ip, M. Lotosh, A. Shah, J. Weng, and R. Tang. “Garbage in, garbage out” revisited: What do machine learning application

- papers report about human-labeled training data? *Quantitative Science Studies*, 2:795–827, 2021.
- [68] T. Khoshgoftaar, D. Dittman, R. Wald, and A. Fazelpour. First order statistics based feature selection: A diverse and powerful family of feature selection techniques. In *2012 11th International Conference on Machine Learning and Applications*, volume 2, pages 151–157, 2012.
- [69] H. Peng, F. Long, and C. Ding. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on pattern analysis and machine intelligence*, 27:1226–1238, 2005.
- [70] A. L. Blum and P. Langley. Selection of relevant features and examples in machine learning. *Artificial Intelligence*, 97:245–271, 1997.
- [71] J. Cai, J. Luo, S. Wang, and S. Yang. Feature selection in machine learning: A new perspective. *Neurocomputing*, 300:70–79, 2018.
- [72] D. Luarte, A. K. Myakalwar, M. Velásquez, J. Álvarez, C. Sandoval, R. Fuentes, J. Yañez, and D. Sbarbaro. Combining prior knowledge with input selection algorithms for quantitative analysis using neural networks in laser induced breakdown spectroscopy. *Anal. Methods*, 13:1181–1190, 2021.
- [73] J. Vandekerckhove, D. Matzke, and E.-J. Wagenmakers. Model comparison and the principle of parsimony. *Oxford handbook of computational and mathematical psychology*, pages 300–319, 2015.
- [74] T. des Touches, M. Munda, T. Cornet, P. Gerken, and T. Helleputte. Feature selection with prior knowledge improves interpretability of chemometrics models. *Chemometrics and Intelligent Laboratory Systems*, 240:104905, 2023.
- [75] W. Groves and M. Gini. Optimal airline ticket purchasing using automated user-guided feature selection. In *IJCAI International Joint Conference on Artificial Intelligence*, 2013.
- [76] R. Matthews. Storks deliver babies ($p=0.008$). *Teaching Statistics*, 22:36–38, 2000.
- [77] H. Lu and X. Ma. Hybrid decision tree-based machine learning models for short-term water quality prediction. *Chemosphere*, 249, 2020.
- [78] R. S. Mollel, L. Stankovic, and V. Stankovic. Explainability-informed feature selection and performance prediction for nonintrusive load monitoring. *Sensors*, 23, 2023.

- [79] S. L. Brunton, B. R. Noack, and P. Koumoutsakos. Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics*, 52:477–508, 2020.
- [80] S. Cedillo, A.-G. Núñez, E. Sánchez-Cordero, L. Timbe, E. Samaniego, and A. Alvarado. Physics-informed neural network water surface predictability for 1d steady-state open channel cases with different flow types and complex bed profile shapes. *Advanced Modeling and Simulation in Engineering Sciences*, 9:1–23, 2022.
- [81] M. C. Kennedy and A. O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63:425–464, 2001.
- [82] J. Brynjarsdóttir and A. O’Hagan. Learning about physical parameters: the importance of model discrepancy. *Inverse Problems*, 30:114007, 2014.
- [83] P. Gardner, T. Rogers, C. Lord, and R. Barthorpe. Learning model discrepancy: A Gaussian process and sampling-based approach. *Mechanical Systems and Signal Processing*, 152:107381, 2021.
- [84] Z. E. ASRI, C. Rambour, V. L. GUEN, and N. THOME. Residual model-based reinforcement learning for physical dynamics. In *NeurIPS 3rd Offline RL Workshop*, 2022.
- [85] P. Zhang, Z.-Y. Yin, Y.-F. Jin, J. Yang, and B. Sheil. Physics-informed multifidelity residual neural networks for hydromechanical modeling of granular soils and foundation considering internal erosion. *Journal of Engineering Mechanics*, 148:04022015, 2022.
- [86] H. Chen, Q. Zhang, and Y. Birkelund. Machine learning forecasts of scandinavian numerical weather prediction wind model residuals with control theory for wind energy. *Energy Reports*, 8:661–668, 2022.
- [87] Z. Huang and Z. Chalabi. Use of time-series analysis to model and forecast wind speed. *Journal of Wind Engineering and Industrial Aerodynamics*, 56:311–322, 1995.
- [88] P. Ailliot and V. Monbet. Markov-switching autoregressive models for wind time series. *Environmental Modelling & Software*, 30:92–101, 2012.
- [89] J. Sirignano, J. F. MacArt, and J. B. Freund. Dpm: A deep learning pde augmentation method with application to large-eddy simulation. *Journal of Computational Physics*, 423:109811, 2020.

- [90] H.-P. Wan and W.-X. Ren. A residual-based Gaussian process model framework for finite element model updating. *Computers and Structures*, 156: 149–159, 2015.
- [91] R. Peli, A. Menafoglio, M. Cervino, P. Secchi, and L. Dovera. Physics-based residual Kriging for dynamically evolving functional random fields. *Stochastic Environmental Research and Risk Assessment*, 2022.
- [92] J. Proppe, S. Gugler, and M. Reiher. Gaussian process-based refinement of dispersion corrections. *Journal of chemical theory and computation*, 2019.
- [93] R. G. Adriana Reyes and J. Mateu. Residual Kriging for functional spatial prediction of salinity curves. *Communications in Statistics - Theory and Methods*, 44:798–809, 2015.
- [94] W. Xing, A. Shah, P. Wang, S. Zhe, Q. Fu, and R. Kirby. Residual Gaussian process: A tractable nonparametric Bayesian emulator for multi-fidelity simulations. *Applied Mathematical Modelling*, 97:36–56, 2021.
- [95] S. K. Chakrabarti, A. L. Wolbert, and W. A. Tam. Wave forces on inclined tubes. *Coastal Engineering*, 1:149–165, 1977.
- [96] V. Sundar, V. Vengatesan, G. Anandkumar, and A. Schlenkhoff. Hydrodynamic coefficients for inclined cylinders. *Ocean Engineering*, 25(4):277–294, 1998.
- [97] G. Anandkumar, V. Sundar, K. U. Graw, and H. Kaldenhoff. Pressures and forces on inclined cylinders due to regular waves. *Ocean Engineering*, 22: 747–759, 1995.
- [98] T. Sarpkaya. Resistance in unsteady flow: Search for an in-line force model. *International Journal of Offshore and Polar Engineering*, 10:249–254, 12 2000.
- [99] P. K. Stansby, K. Worden, G. R. Tomlinson, and S. A. Billings. Improved wave force classification using system identification. *Applied Ocean Research*, 14(2):107–118, 1992.
- [100] D. J. Pitchforth, T. J. Rogers, U. T. Tygesen, and E. J. Cross. Grey-box models for wave loading prediction. *Mechanical Systems and Signal Processing*, 159:107741, 2021.
- [101] E. Didenkulova, I. Didenkulova, and I. Medvedev. Freak wave events in 2005–2021: statistics and analysis of favourable wave and wind conditions. *Natural Hazards and Earth System Sciences*, 23:1653–1663, 2023.

- [102] L. S. Bastos and A. O'Hagan. Diagnostics for Gaussian process emulators. *Technometrics*, 51:425–438, 2009.
- [103] B. Pena and L. Huang. Wave-gan: A deep learning approach for the prediction of nonlinear regular wave loads and run-up on a fixed cylinder. *Coastal Engineering*, 167:103902, 2021.
- [104] L. Sun, H. Gao, S. Pan, and J.-X. Wang. Surrogate modeling for fluid flows based on physics-constrained deep learning without simulation data. *Computer Methods in Applied Mechanics and Engineering*, 361:112732, 2020.
- [105] L. Liang, M. Liu, C. Martin, and W. Sun. A deep learning approach to estimate stress distribution: a fast and accurate surrogate of finite-element analysis. *Journal of The Royal Society Interface*, 15:20170844, 2018.
- [106] B. P. Croom, M. Berkson, R. K. Mueller, M. Presley, and S. Storck. Deep learning prediction of stress fields in additively manufactured metals with intricate defect networks. *Mechanics of Materials*, 165:104191, 2022.
- [107] T. E. Fricker, J. E. Oakley, N. D. Sims, and K. Worden. Probabilistic uncertainty analysis of an frf of a structure using a Gaussian process emulator. *Mechanical Systems and Signal Processing*, 25:2962 – 2975, 2011.
- [108] I. Andrianakis, I. R. Vernon, N. McCreesh, T. J. McKinley, J. E. Oakley, R. N. Nsubuga, M. Goldstein, and R. G. White. Bayesian history matching of complex infectious disease models using emulation: A tutorial and a case study on HIV in Uganda. *PLOS Computational Biology*, 11:1–18, 2015.
- [109] I. Vernon, M. Goldstein, and R. Bower. Galaxy Formation: Bayesian History Matching for the Observable Universe. *Statistical Science*, 29:81 – 90, 2014.
- [110] M. Ratto, A. Castelletti, and A. Pagano. Emulation techniques for the reduction and sensitivity analysis of complex environmental models. *Environmental Modelling & Software*, 34:1–4, 2012.
- [111] H. Zhao, F. Amann, and J. Kowalski. Emulator-based global sensitivity analysis for flow-like landslide run-out models. *Landslides*, 18:3299–3314, 2021.
- [112] A. Sarri, S. Guillas, and F. Dias. Statistical emulation of a tsunami model for sensitivity analysis and uncertainty quantification. *Natural Hazards and Earth System Sciences*, 12:2003–2018, 2012.

- [113] E. Borgonovo, W. Castaings, and S. Tarantola. Model emulation and moment-independent sensitivity analysis: An application to environmental modelling. *Environmental Modelling Software*, 34:105–115, 2012.
- [114] M. Kennedy and G. Petropoulos. Chapter 17 - GEM-SA: The Gaussian Emulation Machine for Sensitivity Analysis. In *Sensitivity Analysis in Earth Observation Modelling*, pages 341–361. Elsevier, 2017.
- [115] A. O. Omigbodun, F. Noo, M. McNitt-Gray, W. Hsu, and S. S. Hsieh. The effects of physics-based data augmentation on the generalizability of deep neural networks: Demonstration on nodule false-positive reduction. *Medical Physics*, 46:4563–4574, 2019.
- [116] B. Sohlberg and M. Sernfält. Grey box modelling for river control. *Journal of Hydroinformatics*, 4:265–280, 2002.
- [117] R. Fuentes, E. Cross, A. Halfpenny, K. Worden, and R. Barthorpe. Aircraft parametric structural load monitoring using Gaussian process regression. In *Proceedings of the European Workshop on Structural Health Monitoring*, 2014.
- [118] T. J. Rogers, G. Holmes, E. J. Cross, and K. Worden. On a grey box modelling framework for nonlinear system identification. In *Special Topics in Structural Dynamics, Volume 6*, pages 167–178, 2017.
- [119] I. Brahma, R. Jennings, and B. Freid. Using physics to extend the range of machine learning models for an aerodynamic, hydraulic and combusting system: The toy model concept. *Energy and AI*, 6:100113, 2021.
- [120] L. Ciampiconi, A. Elwood, M. Leonardi, A. Mohamed, and A. Rozza. A survey and taxonomy of loss functions in machine learning, 2023.
- [121] Q. Wang, Y. Ma, K. Zhao, and Y. Tian. A comprehensive survey of loss functions in machine learning. *Annals of Data Science*, 9:187–212, 2022.
- [122] Z. Xiang, W. Peng, X. Liu, and W. Yao. Self-adaptive loss balanced physics-informed neural networks. *Neurocomputing*, 496:11–34, 2022.
- [123] Z. Liu, M. Roy, D. K. Prasad, and K. Agarwal. Physics-guided loss functions improve deep learning performance in inverse scattering. *IEEE Transactions on Computational Imaging*, 8:236–245, 2022.
- [124] S. J. Raymond and D. B. Camarillo. Applying physics-based loss functions to neural networks for improved generalizability in mechanics problems. *ArXiv*, abs/2105.00075, 2021.

- [125] J. Zhang, S. Zhang, and G. Lin. Pagp: A physics-assisted Gaussian process framework with active learning for forward and inverse problems of partial differential equations. *ArXiv*, 2022.
- [126] J. Bai, T. Rabczuk, A. Gupta, L. Alzubaidi, and Y. Gu. A physics-informed neural network technique based on a modified loss function for computational 2d and 3d solid mechanics. *Comput. Mech.*, 71:543–562, 2022.
- [127] M. Diligenti, S. Roychowdhury, and M. Gori. Integrating prior knowledge into deep learning. In *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 920–923, 2017.
- [128] T. Kadeethum, T. M. Jørgensen, and H. M. Nick. Physics-informed neural networks for solving nonlinear diffusivity and biot’s equations. *PLOS ONE*, 15:1–28, 2020.
- [129] R. Laubscher and P. Rousseau. Application of a mixed variable physics-informed neural network to solve the incompressible steady-state and transient mass, momentum, and energy conservation equations for flow over in-line heated tubes. *Applied Soft Computing*, 114:108050, 2022.
- [130] D. H. Wolpert and W. G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1:67–82, 1997.
- [131] T. Candela, C. Goncalves Machado, O. Leeuwenburgh, and J. Ter Heege. A physics-informed optimization workflow to manage injection while constraining induced seismicity: The Oklahoma case. *Frontiers in Earth Science*, 10:1053951, 2022.
- [132] A. AlBahar, I. Kim, X. Wang, and X. Yue. Physics-constrained Bayesian optimization for optimal actuators placement in composite structures assembly. *IEEE Transactions on Automation Science and Engineering*, 20:2772–2783, 2023.
- [133] A. R. Junejo, N. U. Gilal, and J. Doh. Physics-informed optimization of robust control system to enhance power efficiency of renewable energy: Application to wind turbine. *Energy*, 263:125667, 2023.
- [134] C. Xie, S. Patil, T. Moldovan, S. Levine, and P. Abbeel. Model-based reinforcement learning with parametrized physical models and optimism-driven exploration. pages 504–511, 2016.
- [135] D. Q. Ha and J. Carstensen. Human-informed topology optimization: interactive application of feature size controls. *Structural and Multidisciplinary Optimization*, 66, 2023.

- [136] Z. Bai and S. Song. Structural reliability analysis based on neural networks with physics-informed training samples. *Engineering Applications of Artificial Intelligence*, 126:107157, 2023.
- [137] H. Bolandi, G. Sreekumar, X. Li, N. Lajnef, and V. N. Boddeti. Physics informed neural network for dynamic stress prediction. *Applied Intelligence*, 53:26313–26328, 2023.
- [138] S. Moradi, B. Duran, S. Eftekhar Azam, and M. Mofid. Novel physics-informed artificial neural network architectures for system and input identification of structural dynamics pdes. *Buildings*, 13, 2023.
- [139] C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for Machine Learning*. The MIT Press, 2005.
- [140] S. Pfingstl and M. Zimmermann. On integrating prior knowledge into Gaussian processes for prognostic health monitoring. *Mechanical Systems and Signal Processing*, 171:108917, 2022.
- [141] D. An, N. H. Kim, and J.-H. Choi. Practical options for selecting data-driven or physics-based prognostics algorithms with reviews. *Reliability Engineering & System Safety*, 133:223–236, 2015.
- [142] M. Tóth, A. Brown, E. Cross, T. Rogers, and N. D. Sims. Resource-efficient machining through physics-informed machine learning. *Procedia CIRP*, 117:347–352, 2023. 19th CIRP Conference on Modeling of Machining Operations.
- [143] D. Kong, Y. Chen, and N. Li. Gaussian process regression for tool wear prediction. *Mechanical Systems and Signal Processing*, 104:556–574, 2018.
- [144] S. Aye and P. Heyns. An integrated Gaussian process regression for prediction of remaining useful life of slow speed bearings based on acoustic emission. *Mechanical Systems and Signal Processing*, 84:485–498, 2017.
- [145] D. Nguyen-Tuong and J. Peters. Using model knowledge for learning inverse dynamics. In *2010 IEEE international conference on robotics and automation*, pages 2677–2682. IEEE, 2010.
- [146] T. J. Rogers, P. Gardner, N. Dervilis, K. Worden, A. E. Maguire, E. Papathéou, and E. J. Cross. Probabilistic modelling of wind turbine power curves with application of heteroscedastic Gaussian process regression. *Renewable Energy*, 148:1124–1136, 2020.

- [147] L. A. Bull, P. A. Gardner, T. J. Rogers, N. Dervilis, E. J. Cross, E. Papatheou, A. E. Maguire, C. Campos, and K. Worden. Bayesian modelling of multivalued power curves from an operational wind farm. *Mechanical Systems and Signal Processing*, 169:108530, 2022.
- [148] E. Schulz, M. Speekenbrink, and A. Krause. A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions. *Journal of Mathematical Psychology*, 85:1–16, 2018.
- [149] E. J. Cross and T. J. Rogers. Physics-derived covariance functions for machine learning in structural dynamics. *IFAC-PapersOnLine*, 54:168–173, 2021. 19th IFAC Symposium on System Identification SYSID 2021.
- [150] M. Kok and A. Solin. Scalable magnetic field SLAM in 3d using Gaussian process maps. In *2018 21st International Conference on Information Fusion (FUSION)*, pages 1353–1360, 2018.
- [151] M. M. Noack and J. A. Sethian. Advanced stationary and nonstationary kernel designs for domain-aware Gaussian processes. *Communications in Applied Mathematics and Computational Science*, 17, 2022.
- [152] M. Haywood-Alexander, N. Dervilis, K. Worden, E. J. Cross, R. S. Mills, and T. J. Rogers. Structured machine learning tools for modelling characteristics of guided waves. *Mechanical Systems and Signal Processing*, 156:107628, 2021.
- [153] S. Klus, P. Gelß, F. Nüske, and F. Noé. Symmetric and antisymmetric kernels for machine learning problems in quantum physics and chemistry. *Machine Learning: Science and Technology*, 2, 2021.
- [154] A. Grisafi, D. Wilkins, G. Csányi, and M. Ceriotti. Symmetry-adapted machine learning for tensorial properties of atomistic systems. *Physical Review Letters*, 120, 2018.
- [155] D. Duvenaud, J. Lloyd, R. Grosse, J. Tenenbaum, and G. Zoubin. Structure discovery in nonparametric regression through compositional kernel search. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 1166–1174. PMLR, 2013.
- [156] D. Duvenaud. *Automatic model construction with Gaussian processes*. PhD thesis, University of Cambridge, 2014.

- [157] J. Lloyd, D. Duvenaud, R. Grosse, J. Tenenbaum, and Z. Ghahramani. Automatic construction and natural-language description of nonparametric regression models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 28, 2014.
- [158] M. Álvarez, D. Luengo, and N. D. Lawrence. Latent force models. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, pages 9–16, 2009.
- [159] J. Hartikainen and S. Sarkka. Sequential inference for latent force models. *arXiv preprint arXiv:1202.3730*, 2012.
- [160] R. Nayek, S. Chakraborty, and S. Narasimhan. A Gaussian process latent force model for joint input-state estimation in linear structural systems. *Mechanical Systems and Signal Processing*, 128:497–530, 2019.
- [161] T. J. Rogers, K. Worden, and E. J. Cross. On the application of Gaussian process latent force models for joint input-state-parameter estimation: With a view to Bayesian operational identification. *Mechanical Systems and Signal Processing*, 140:106580, 2020.
- [162] T. J. Rogers, K. Worden, G. Manson, U. T. Tygesen, and E. J. Cross. A Bayesian filtering approach to operational modal analysis with recovery of forcing signals. In *ISMA 2018 - International Conference on Noise and Vibration Engineering and USD 2018 - International Conference on Uncertainty in Structural Dynamics*, 2018.
- [163] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82:35–45, 1960.
- [164] C. Guarnizo and M. A. Lopez. Fast kernel approximations for latent force models and convolved multiple-output Gaussian processes, 2018.
- [165] Ø. W. Petersen, O. Øiseth, and E. Lourens. Wind load estimation and virtual sensing in long-span suspension bridges using physics-informed Gaussian process latent force models. *Mechanical Systems and Signal Processing*, 170:108742, 2022.
- [166] S. Cuomo, V. S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, and F. Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what’s next. *J. Sci. Comput.*, 92, 2022.

- [167] Z. K. Lawal, H. Yassin, D. T. C. Lai, and A. Che Idris. Physics-informed neural network (PINN) evolution and beyond: A systematic literature review and bibliometric analysis. *Big Data and Cognitive Computing*, 6, 2022.
- [168] M. Rasht-Behesht, C. Huber, K. Shukla, and G. E. Karniadakis. Physics-informed neural networks (PINNs) for wave propagation and full waveform inversions. *Journal of Geophysical Research: Solid Earth*, 127:e2021JB023120, 2022.
- [169] B. Moseley, A. Markham, and T. Nissen-Meyer. Solving the wave equation with physics-informed deep learning, 2020.
- [170] L. Sun, H. Gao, S. Pan, and J.-X. Wang. Surrogate modeling for fluid flows based on physics-constrained deep learning without simulation data. *Computer Methods in Applied Mechanics and Engineering*, 361:112732, 2020.
- [171] H. Eivazi, M. Tahani, P. Schlatter, and R. Vinuesa. Physics-informed neural networks for solving Reynolds-averaged Navier–Stokes equations. *Physics of Fluids*, 34(7), 2022.
- [172] E. Haghghat, M. Raissi, A. Moure, H. Gomez, and R. Juanes. A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics. *Computer Methods in Applied Mechanics and Engineering*, 379:113741, 2021.
- [173] R. Rodriguez-Torrado, P. Ruiz, L. Cueto-Felgueroso, M. Green, T. Friesen, S. Matringe, and J. Togelius. Physics-informed attention-based neural network for hyperbolic partial differential equations: application to the Buckley–Leverett problem. *Scientific Reports*, 12:7557, 2022.
- [174] S. Buckley and M. Leverett. Mechanism of fluid displacement in sands. *Transactions of the AIME*, 146:107–116, 1942.
- [175] J. Abbasi and P. Østebø Andersen. Physical activation functions (PAFs): An approach for more efficient induction of physics into physics-informed neural networks (PINNs), 2023.
- [176] S. A. Faroughi, P. Datta, S. K. Mahjour, and S. Faroughi. Physics-informed neural networks with periodic activation functions for solute transport in heterogeneous porous media, 2022.
- [177] E. A. Wan and R. Van Der Merwe. The unscented Kalman filter for nonlinear estimation. In *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373)*, pages 153–158. Ieee, 2000.

- [178] S. J. Julier and J. K. Uhlmann. Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92:401–422, 2004.
- [179] F. Gustafsson and G. Hendeby. Some relations between extended and unscented Kalman filters. *IEEE Transactions on Signal Processing*, 60:545–555, 2012.
- [180] R. Frigola, F. Lindsten, T. B. Schön, and C. E. Rasmussen. Bayesian inference and learning in Gaussian process state-space models with particle MCMC. In *Advances in Neural Information Processing Systems*, volume 26, 2013.
- [181] S. Eleftheriadis, T. F. Nicholson, M. P. Deisenroth, and J. Hensman. Identification of Gaussian process state space models. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5315–5325, 2017.
- [182] R. Frigola, Y. Chen, and C. E. Rasmussen. Variational Gaussian process state-space models. *Advances in neural information processing systems*, 27, 2014.
- [183] R. Frigola. *Bayesian Time Series Learning with Gaussian Processes*. PhD thesis, University of Cambridge, 2015.
- [184] A. Svensson, A. Solin, S. Särkkä, and T. Schön. Computationally efficient Bayesian learning of Gaussian process state space models. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 213–221, 2016.
- [185] F. Lindsten, T. B. Schön, and M. I. Jordan. Bayesian semiparametric Wiener system identification. *Automatica*, 49:2053–2063, 2013.
- [186] W. Greblicki and M. Pawlak. *Nonparametric System Identification*. 2008.
- [187] L. Ljung. *System Identification: Theory for the User*. Prentice Hall information and system sciences series. Prentice Hall PTR, 1999.
- [188] T. Rogers, T. Schön, A. Lindholm, K. Worden, and E. Cross. Identification of a Duffing oscillator using particle Gibbs with ancestor sampling. In *Journal of Physics: Conference Series*, volume 1264, page 012051, 2019.
- [189] T. Rogers and T. Friis. A latent restoring force approach to nonlinear system identification. *Mechanical Systems and Signal Processing*, 180:109426, 2022.

- [190] S. Chmiela, A. Tkatchenko, H. E. Sauceda, I. Poltavsky, K. T. Schütt, and K.-R. Müller. Machine learning of accurate energy-conserving molecular force fields. *Science Advances*, 3:e1603015, 2017.
- [191] T. Beucler, S. Rasp, M. Pritchard, and P. Gentine. Achieving conservation of energy in neural network emulators for climate modeling. *arXiv preprint arXiv:1906.06622*, 2019.
- [192] M. R. Jones, T. J. Rogers, and E. J. Cross. Constraining Gaussian processes for physics-informed acoustic emission mapping. *Mechanical Systems and Signal Processing*, 188:109984, 2023.
- [193] J. H. Mclean, M. R. Jones, B. J. O’Connell, E. Maguire, and T. J. Rogers. Physically meaningful uncertainty quantification in probabilistic wind turbine power curve models as a damage-sensitive feature. *Structural Health Monitoring*, 22:3623–3636, 2023.
- [194] N. Wahlström. *Modeling of magnetic fields and extended objects for localization applications*. PhD thesis, Linköping University Electronic Press, 2015.
- [195] C. Jidling, N. Wahlström, A. Wills, and T. B. Schön. Linearly constrained Gaussian processes. *Advances in Neural Information Processing Systems*, 30, 2017.
- [196] C. Jidling, J. Hendriks, N. Wahlström, A. Gregg, T. B. Schön, C. Wensrich, and A. Wills. Probabilistic modelling and reconstruction of strain. *Nuclear Instruments and Methods in Physics Research*, 436:141 – 155, 2018.
- [197] E. J. Cross, T. J. Rogers, and T. J. Gibbons. Grey-box modelling for structural health monitoring: Physical constraints on machine learning algorithms. In *12th International workshop on Structural Health Monitoring*, 2019.
- [198] I. Higgins, S. Racanière, and D. Rezende. Symmetry-based representations for artificial and biological general intelligence. *Frontiers in Computational Neuroscience*, 16, 2022.
- [199] T. E. Smidt. Euclidean symmetry and equivariance in machine learning. *Trends in Chemistry*, 3:82–85, 2021. Special Issue: Machine Learning for Molecules and Materials.
- [200] J. J. Meyer, M. Mularski, E. Gil-Fuster, A. A. Mele, F. Arzani, A. Wilms, and J. Eisert. Exploiting symmetry in variational quantum machine learning. *PRX Quantum*, 4:010328, 2023.

- [201] M. R. Jones, T. J. Rogers, and E. J. Cross. Constraining Gaussian processes for physics-informed acoustic emission mapping. *arXiv preprint arXiv:2206.01495*, 2022.
- [202] L. Liu, S. M. Kuo, and M. Zhou. Virtual sensing techniques and their applications. In *2009 International Conference on Networking, Sensing and Control*, pages 31–36, 2009.
- [203] J. Bilbao, E.-M. Lourens, A. Schulze, and L. Ziegler. Virtual sensing in an onshore wind turbine tower using a Gaussian process latent force model. *Data-Centric Engineering*, 3:e35, 2022.
- [204] S. Yoon. Virtual sensing in intelligent buildings and digitalization. *Automation in Construction*, 143:104578, 2022.
- [205] J. Zou, E.-M. Lourens, and A. Cicirello. Virtual sensing of subsoil strain response in monopile-based offshore wind turbines via Gaussian process latent force models. *Mechanical Systems and Signal Processing*, 200:110488, 2023.
- [206] S. J. Gibson, T. J. Rogers, and E. J. Cross. Distributions of fatigue damage from data-driven strain prediction using Gaussian process regression. *Structural Health Monitoring*, 22:3065–3076, 2023.
- [207] F. Pütz, M. Henrich, N. Fehlemann, A. Roth, and S. Münstermann. Generating input data for microstructure modelling: A deep learning approach using generative adversarial networks. *Materials*, 13, 2020.
- [208] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113:3932–3937, 2016.
- [209] Z. Chen, Y. Liu, and H. Sun. Physics-informed learning of governing equations from scarce data. *Nature Communications*, 12, 2020.
- [210] H. Xu, D. Zhang, and N. Wang. Deep-learning based discovery of partial differential equations in integral form from sparse and noisy data. *Journal of Computational Physics*, 445:110592, 2021.
- [211] A. Pratap Singh, S. Medida, and K. Duraisamy. Machine-learning-augmented predictive modeling of turbulent separated flows over airfoils. *AIAA Journal*, 55, 08 2016.
- [212] L. Gabrielli, S. Tomassetti, S. Squartini, and C. Zinato. Introducing deep machine learning for parameter estimation in physical modelling. In *Proceedings of the 20th international conference on digital audio effects*, 2017.

- [213] Z. Lai, C. Mylonas, S. Nagarajaiah, and E. Chatzi. Structural identification with physics-informed neural Ordinary Differential Equations. *Journal of Sound and Vibration*, 508:116196, 2021.
- [214] A. Solin, M. Kok, N. Wahlström, T. B. Schön, and S. Särkkä. Modeling and interpolation of the ambient magnetic field by Gaussian processes. *IEEE Transactions on Robotics*, 34:1112–1127, 2018.
- [215] D. J. C. Mackay. *Introduction to Gaussian processes*, pages 84–92. Springer, 1998.
- [216] M. Kuß. *Gaussian Process Models for Robust Regression, Classification, and Reinforcement Learning*. PhD thesis, Technische Universität, 2006.
- [217] Y. Goda. *Random Seas and Design of Maritime Structures*. World Scientific, 2000.
- [218] J. V. Wehausen and E. V. Laitone. *Surface Waves*, pages 446–778. Springer Berlin Heidelberg, 1960.
- [219] P. LeBlond and L. Mysak. *Waves in the Ocean*. Elsevier Science, 1981.
- [220] W. J. Pierson Jr. and L. Moskowitz. A proposed spectral form for fully developed wind seas based on the similarity theory of S. A. Kitaigorodskii. *Journal of Geophysical Research (1896-1977)*, 69:5181–5190, 1964.
- [221] K. Hasselmann and D. Olbers. Measurements of wind-wave growth and swell decay during the Joint North Sea Wave Project (JONSWAP). *Ergänzung zur Deut. Hydrogr. Zeitschrift, Reihe A (8)*, 12:1–95, 1973.
- [222] N. Barltrop and A. Adams. *Dynamics of Fixed Marine Structures*. Butterworth-Heinemann, Third edition, 1991.
- [223] D. J. Evans. Analysis of wave force data. *Journal of Petroleum Technology*, 22(03):347–358, 1970.
- [224] G. R. Cook and E. Simiu. Hydrodynamic forces on vertical cylinders and the Lighthill correction. *Ocean Engineering*, 16(4):355–372, 1989.
- [225] T. S. Sarpkaya. *Wave Forces on Offshore Structures*. Cambridge University Press, 2010.
- [226] J. Lighthill. Fundamentals concerning wave loading on offshore structures. *Journal of Fluid Mechanics*, 173:667–681, 1986.

- [227] T. Sarpkaya. On the force decompositions of Lighthill and Morison. *Journal of Fluids and Structures*, 15(2):227–233, 2001.
- [228] K. Worden, P. K. Stansby, G. R. Tomlinson, and S. A. Billings. On wave force analysis using system identification. Technical report, University of Sheffield, 1992.
- [229] J. Bishop. Wave force data from the second Christchurch Bay Tower. In *OTC Offshore Technology Conference*, 1985.
- [230] G. Najafian, R. G. Tickell, R. Burrows, and J. R. Bishop. The UK Christchurch Bay Compliant Cylinder Project: analysis and interpretation of Morison wave force and response data. *Applied Ocean Research*, 22:129 – 153, 2000.
- [231] J. R. Bishop, R. G. Tickell, and K. A. Gallagher. The UK Christchurch Bay Project; a review of results. In *OTC Offshore Technology Conference*, 1980.
- [232] K. Worden, W. Becker, T. J. Rogers, and E. J. Cross. On the confidence bounds of Gaussian process NARX models and their higher-order frequency response functions. *Mechanical Systems and Signal Processing*, 104:188–223, 2018.
- [233] G. Clauss, E. Lehmann, and C. Östergaard. *Offshore structures*. Springer-Verlag, 1992.
- [234] *Environmental Conditions and Environmental Loads, Recommended Practice, DNV-RP-C205*. Det Norske Veritas, 2010.
- [235] R. Burrows, R. Tickell, D. Hames, and G. Najafian. Morison wave force coefficients for application to random seas. *Applied Ocean Research*, 19:183–199, 1997.
- [236] O. M. Faltinsen, J. N. Newman, and T. Vinje. Nonlinear wave loads on a slender vertical cylinder. *Journal of Fluid Mechanics*, 289:179–198, 1995.
- [237] N. Kim, M.-S. Park, and S. Yang. Wave force analysis of the vertical circular cylinder by boundary element method. *KSCE Journal of Civil Engineering*, 11:31–35, 2007.
- [238] L. Arany, S. Bhattacharya, J. H. Macdonald, and S. J. Hogan. Closed form solution of Eigen frequency of monopile supported offshore wind turbines in deeper waters incorporating stiffness of substructure and SSI. *Soil Dynamics and Earthquake Engineering*, 83:18–32, 2016.

- [239] M. Zaaijer *et al.* Design methods for offshore wind turbines at exposed sites (OWTES) - sensitivity analysis for foundations of offshore wind turbines. *Delft, The Netherlands*, 2002.
- [240] M. Damgaard, L. B. Ibsen, L. V. Andersen, and J. K. Andersen. Cross-wind modal properties of offshore wind turbines identified by full scale testing. *Journal of Wind Engineering and Industrial Aerodynamics*, 116:94–108, 2013.
- [241] T. Camp, M. Morris, R. Van Rooij, J. Van Der Tempel, M. Zaaijer, A. Henderson, K. Argyriadis, S. Schwartz, H. Just, W. Grainger, *et al.* Design methods for offshore wind turbines at exposed sites. final report of the OWTES project. Technical report, Garrad Hassan and Partners, 2003.
- [242] C. Leblanc Thilsted and N. Tarp-Johansen. Monopiles in sand - stiffness and damping. In *Proceedings of the European wind energy conference*, 2011.
- [243] R. D. Blevins. *Formulas for Natural Frequency and Mode Shape*. R.E. Krieger, 1984.
- [244] S. S. Rao. *Vibration of Continuous Systems*. John Wiley & Sons, Ltd, 2006.
- [245] J. Greig and I. Jason. Wave generation software manual. Technical report, Edinburgh Designs, 2012.
- [246] E. E. Bachynski, T. Kristiansen, and M. Thys. Experimental and numerical investigations of monopile ringing in irregular finite-depth water waves. *Applied Ocean Research*, 68:154–170, 2017.
- [247] J. M. Esandi, E. Buldakov, R. Simons, and D. Stagonas. An experimental study on wave forces on a vertical cylinder due to spilling breaking and near-breaking wave groups. *Coastal Engineering*, 162:103778, 2020.
- [248] G. Casella and E. I. George. Explaining the Gibbs sampler. *The American Statistician*, 46:167–174, 1992.
- [249] J. Hensman, N. Fusi, and N. Lawrence. Gaussian processes for big data. In *Uncertainty in Artificial Intelligence - Proceedings of the 29th Conference, UAI 2013*, pages 282–290, 2013.
- [250] H. He and W. Siu. Single image super-resolution using Gaussian process regression. In *CVPR 2011*, pages 449–456, 2011.
- [251] T. J. Rogers, P. Gardner, N. Dervilis, K. Worden, A. E. Maguire, E. Papa-theou, and E. J. Cross. Probabilistic modelling of wind turbine power curves with application of heteroscedastic Gaussian process regression. *Renewable Energy*, 148:1124–1136, 2020.

- [252] J. Kocijan, A. Girard, B. Banko, and R. Murray-Smith. Dynamic systems identification with Gaussian processes. *Mathematical and Computer Modelling of Dynamical Systems*, 11:411–424, 2005.
- [253] L. Piroddi, M. Farina, and M. Lovera. Polynomial NARX model identification: A Wiener–Hammerstein benchmark. *IFAC Proceedings Volumes*, 42:1074 – 1079, 2009.
- [254] M. Farina and L. Piroddi. Simulation error minimization–based identification of polynomial input–output recursive models. *IFAC Proceedings Volumes*, 42: 1393 – 1398, 2009.
- [255] B. J. N. Blight and L. Ott. A Bayesian approach to model inadequacy for polynomial regression. *Biometrika*, 62:79–88, 1975.
- [256] E. Guilmineau and P. Queutey. A numerical simulation of vortex shedding from an oscillating circular cylinder. *Journal of Fluids and Structures*, 16: 773–794, 2002.
- [257] A. M. M. Wood and C. A. Fleming. *Coastal Hydraulics*. Macmillan Education, Limited, 1981.
- [258] K. Worden, R. J. Barthorpe, E. J. Cross, N. Dervilis, G. R. Holmes, G. Manson, and T. J. Rogers. On evolutionary system identification with applications to nonlinear benchmarks. *Mechanical Systems and Signal Processing*, 112: 194 – 232, 2018.
- [259] T. J. Rogers. *Towards Bayesian System Identification: With Application to SHM of Offshore Structures*. PhD thesis, University of Sheffield, 2019.
- [260] L. Juan, B. Feng, and W. Xu. Particle swarm optimization with particles having quantum behavior. In *Proceedings of the 2004 Congress on Evolutionary Computation, CEC2004*, volume 1, pages 325 – 331, 2004.
- [261] H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19:716–723, 1974.
- [262] K. P. Burnham and D. R. Anderson. *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. Springer New York, 2002.
- [263] C. M. Hurvich and C.-L. Tsai. Regression and time series model selection in small samples. *Biometrika*, 76:297–307, 1989.

- [264] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6:461–464, 1978.
- [265] K. P. Burnham and D. R. Anderson. Multimodel inference: Understanding AIC and BIC in model selection. *Sociological Methods & Research*, 33:261–304, 2004.
- [266] E. Snelson and Z. Ghahramani. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems*, pages 1257–1264. MIT press, 2006.
- [267] M. Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, pages 567–574, 2009.
- [268] I. Young. On the measurement of directional wave spectra. *Applied Ocean Research*, 16:283–294, 1994.
- [269] N. N. Panicker and L. E. Borgman. *Enhancement of Directional Wave Spectrum Estimates*, pages 258–279. American Society of Civil Engineers, 1974.
- [270] Ø. W. Petersen, O. Øiseth, and E. Lourens. Full-scale identification of the wave forces exerted on a floating bridge using inverse methods and directional wave spectrum estimation. *Mechanical Systems and Signal Processing*, 120:708–726, 2019.
- [271] *MIROS WAVEFUSION Datasheet, The unique dual sensor for measuring directional wave spectra*. MIROS group, 2023.
- [272] A. Aggarwal, C. Pâkozdi, H. Bihs, D. Myrhaug, and M. Alagan Chella. Free surface reconstruction for phase accurate irregular wave generation. *Journal of Marine Science and Engineering*, 6:105–128, 2018.
- [273] H. N. Southgate and D. G. Goldberg. An efficient computational model for wave refraction and diffraction using finite differences. Technical report, Hydraulics Research Wallingford, 1989.
- [274] H. Li and E. E. Bachynski. Experimental and numerical investigation of non-linear diffraction wave loads on a semi-submersible wind turbine. *Renewable Energy*, 171:709–727, 2021.
- [275] C. Williamson. In-line response of a cylinder in oscillatory flow. *Applied Ocean Research*, 7:97–106, 1985.

- [276] J. Stoker. *Water Waves: The Mathematical Theory with Applications*. Wiley, 2011.
- [277] J. Mercer and A. R. Forsyth. Functions of positive and negative type, and their connection with the theory of integral equations. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 83:69–70, 1909.
- [278] S. Bochner, M. Functions, S. Integrals, H. Analysis, M. Tenenbaum, and H. Pollard. *Lectures on Fourier Integrals. (AM-42)*. Princeton University Press, 1959.
- [279] A. Wilson and R. Adams. Gaussian process kernels for pattern discovery and extrapolation. In *International Conference on Machine Learning (ICML)*, 2013.
- [280] W. T. Thomson and M. D. Dahleh. *Theory of Vibrations with Applications*. Prentice Hall, 1997.
- [281] D. Lombardi, S. Bhattacharya, and D. Muir Wood. Dynamic soil-structure interaction of monopile supported wind turbines in cohesive soil. *Soil Dynamics and Earthquake Engineering*, 49:165–180, 2013.
- [282] M. Martinez-Luengo, P. Causon, A. B. Gill, and A. J. Kolios. The effect of marine growth dynamics in offshore wind turbine support structures. In *Progress in the Analysis and Design of Marine Structures*, pages 889–898. CRC Press/Balkema, 2017.
- [283] M. Tarpø, B. Nabuco, C. Georgakis, and R. Brincker. Expansion of experimental mode shape from operational modal analysis and virtual sensing for fatigue analysis using the modal expansion method. *International Journal of Fatigue*, 130:105280, 2020.
- [284] A. W. Max. Inverting modified matrices. In *Memorandum Rept. 42, Statistical Research Group*, page 4. Princeton Univ., 1950.
- [285] M. R. Jones, D. J. Pitchforth, and E. J. Cross. Gaussian process kernels for partial physical insight. In *Proceedings of the 11th European Workshop on Structural Health Monitoring (EWSHM 2024)*, 2024.
- [286] World sea temperature 2023, North Sea Water Temperature. <https://www.seatemperature.org/north-sea>, 2023. [Online; accessed 02-March-2023].

-
- [287] T. Baldacchino, E. J. Cross, K. Worden, and J. Rowson. Variational Bayesian mixture of experts models and sensitivity analysis for nonlinear dynamical systems. *Mechanical Systems and Signal Processing*, 66-67:178–200, 2016.
- [288] T. Baldacchino, K. Worden, and J. Rowson. Robust nonlinear system identification: Bayesian mixture of experts using the t-distribution. *Mechanical Systems and Signal Processing*, 85:977–992, 2017.
- [289] C.-J. Kim. Dynamic linear models with markov-switching. *Journal of Econometrics*, 60:1–22, 1994.
- [290] K. Worden and E. J. Cross. On switching response surface models, with applications to the structural health monitoring of bridges. *Mechanical Systems and Signal Processing*, 98:139–156, 2018.
- [291] H. Shi, K. Worden, and E. J. Cross. A regime-switching cointegration approach for removing environmental and operational variations in structural health monitoring. *Mechanical Systems and Signal Processing*, 103:381–397, 2018.