

# On Novel Machine Learning Approaches for Acoustic Emission Source Localisation: A Probabilistic Perspective



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

by

M. R. Jones

Department of Mechanical Engineering

University of Sheffield

March 2023



---

## ACKNOWLEDGEMENTS

Firstly, a massive thank you goes to Professor Lizzy Cross, who has been a fantastic and inspiring supervisor, and without whom this thesis would not have been possible. Lizzy, thanks for your continued support and friendship. I would also like to thank Professor Keith Worden for his guidance and invaluable insight that he so regularly provides. Additionally, I would like to extend this thanks to Dr Tim Rogers, whose contribution sparked many ideas used in this work, as well as always finding the time to help me out. It's really appreciated.

For the collection and provision of the experimental data sets used throughout this thesis, I would like to thank Dr James Hensman, Dr Mark Eaton, Dr Robin Mills, Professor Gareth Pierce and Dr Ian Martinez. I am also grateful for the many interesting chats I had with Dr Ramon Fuentes related to the work during the first year of my PhD.

I would like to express my gratitude for being able to work at such a brilliant and inspiring place as the DRG. There are so many lovely and talented people within the group. It has, and continues to be, such a pleasure to be a part of. Particular thanks goes to Marcus, Aidan, Max, Chandy, Chris, Tristan, Tina, Dan and Julian for their friendship.

Thanks goes to my parents and brother, who have always given me their full support and encouragement. Finally, a special thank you goes to Gina, whose patience, care and understanding ensured I had the support necessary to complete the thesis, particularly during lockdown times.



## Highlights:

- A Bayesian framework for probabilistically localising damage in complex structures using acoustic emission is developed.
- The use of both Gaussian process-based and physics-informed machine learning models are presented for the task of source localisation.
- Localisation is performed in the context of SHM and CM.
- An introduction to acoustic emission theory is provided, in addition to an overview of the current state-of-the-art acoustic emission localisation techniques.
- The need for a probabilistic acoustic emission localisation framework is motivated.
- An introduction to Gaussian process regression from the perspective of acoustic emission localisation is presented.
- A novel framework for probabilistically localising acoustic emission events in complex structures is proposed.
- The approach enables a likelihood of emission location to be quantified across the structure of interest.
- It is shown that the proposed method outperforms state-of-the-art methods from the literature over a number of testing scenarios.
- A  $\Delta T$  forward model is developed that more appropriately models the uncertainty on the onset times.
- It is shown that the onset time features are heteroscedastic, which may be handled with a local (input-dependent) noise process in the observation model.
- Predictions from individual sensor pairs are weighted according to how well a particular area of the structure is covered by the pairing, improving the overall prediction accuracy and confidence.

- The burden of data collection for  $\Delta T$  maps is alleviated through the use of physics-informed Gaussian processes.
- Physical insight is included into the  $\Delta T$  learning process by constraining the covariance function to known boundary conditions.
- The constrained model improves predictive accuracy where the availability of training data are reduced.
- The localisation framework is extended to non-Euclidean spaces, allowing localisation in spherical structures such as bearings.
- Acoustic emission localisation is investigated in a scaled-down wind turbine gear box bearing.
- Good localisation performance is achieved in the presence of noisy  $\Delta T$  features.

---

# ABSTRACT

With the objective of making engineering infrastructure safer and more cost-effective to operate and maintain, the use of automated strategies for monitoring damage in structures and high value assets are becoming increasingly common. A critical component in the assessment of a structure's condition is the localisation of defects, with a promising solution the monitoring of acoustic emissions, a technique concerned with passively listening to ultrasonic signals generated by damage mechanisms. With that said, a significant barrier to a more widespread adoption of techniques of this nature are their use in structures with intricate geometrical features and anisotropic materials. In these structures, propagation paths are complex, material parameters often unknown, with stochasticity and a deficiency in complete physical understanding introducing sources of uncertainty that are often unaccounted for.

The work contained in this thesis develops and extends a probabilistic framework for localising acoustic emissions in complex structures, handling uncertainty in a principled manner through Bayesian inference. A forward mapping of expected arrival time information is first learnt through the use of Gaussian process regression. For an event with an unknown origin, it is shown that these maps can be used to quantify a likelihood of emission location, providing probable damage locations on the structure. Next, the use of a heteroscedastic noise model is presented, allowing predictions made by the localisation model to be locally-weighted such that sensors contribute to the prediction relative to the quality of coverage offered, returning a more accurate, confident and robust localisation methodology. On the topic of the practicality of implementing the proposed approach, the inclusion of physical insight is considered within a grey-box framework to constrain the Gaussian process to abide by known physical laws. It is demonstrated that the constraints improve performance where the availability of training data reduces, increasing the feasibility of implementing the

developed methodology. Finally, localisation is extended to cases where the geometry is not most appropriately characterised in Euclidean space, such as for roller-element and many other types of bearings. It is demonstrated how localisation may also be performed in a condition monitoring setting, as well as demonstrating the ability of the method to handle measurements that are contaminated with significant noise levels.

---

# PUBLICATIONS

## Author Publications to date

### Journal papers

M.R. Jones, T.J. Rogers, K. Worden, E.J. Cross, 2022. A Bayesian methodology for localising acoustic emission sources in complex structures. *Mechanical Systems and Signal Processing*. doi.org/10.1016/j.ymssp.2021.108143

### Submitted journal papers

M.R. Jones, T.J. Rogers, E.J. Cross, 2022. Constraining Gaussian processes for physics-informed acoustic emission mapping. *Mechanical Systems and Signal Processing*. doi.org/10.48550/arXiv.2206.01495

J. Mclean, M.R. Jones, B.J. O'Connell, T.J. Rogers, 2022. Physically Meaningful Uncertainty Quantification in Probabilistic Wind Turbine Power Curve Models. Submitted to *Structural Health Monitoring*. arXiv:2209.15579

### Book chapters

E.J. Cross, S.J. Gibson, M.R. Jones, D.J. Pitchforth, S.Zhang, T.J. Rogers, 2022. Physics-Informed Machine Learning for Structural Health Monitoring. *Structural Health Monitoring Based on Data Science Techniques*. doi.org/10.1007/978-3-030-

## Conference papers

M.R. Jones, T.J. Rogers, K. Worden, E.J. Cross, 2022. Enhanced uncertainty quantification for acoustic emission localisation. To appear in the *Proceedings of ISMA 2022 - International Conference on Noise and Vibration Engineering and USD 2022 - International Conference on Uncertainty in Structural Dynamics*.

C.A. Lindley, M.R. Jones, T.J. Rogers, R. Dwyer-Joyce, N. Dervilis, K. Worden, 2022. A nonparametric Bayesian approach for the detection of acoustic emission events in time series signals. To appear in the *Proceedings of ISMA 2022 - International Conference on Noise and Vibration Engineering and USD 2022 - International Conference on Uncertainty in Structural Dynamics*.

M.R. Jones, T.J. Rogers, E.J. Cross, 2022. Constraining Gaussian processes for the inclusion of boundary conditions in physics-informed structural health monitoring. To appear in the *Proceedings of the 11th International Conference on Structural Health Monitoring of Intelligent Infrastructure*.

M.R. Jones, T.J. Rogers, E.J. Cross, 2022. Physical covariance functions for dynamic systems with time-dependent parameters. *Proceedings of the 10th European Workshop on Structural Health Monitoring*.

R.E. Scott, M.R. Jones, T.J. Rogers, 2022. Bayesian Change-point Modelling for Reference-Free Damage Detection with Acoustic Emission Data. *Proceedings of the 10th European Workshop on Structural Health Monitoring*.

M.R. Jones, T.J. Rogers, I.E. Martinez, E.J. Cross, 2021. Bayesian localisation of acoustic emission sources for wind turbine bearings. *Proceedings of SPIE 11593, Health Monitoring of Structural and Biological Systems XV*.

M.R. Jones, T.J. Rogers, K. Worden, E.J. Cross, 2021. Heteroscedastic Gaussian processes for localising acoustic emission. *Proceedings of IMAC XXXIX*.

M.R. Jones, T.J. Rogers, K. Worden, E.J. Cross, 2020. Constraining Gaussian processes for grey-box acoustic emission source localisation. *Proceedings of ISMA 2020 - International Conference on Noise and Vibration Engineering and USD 2020 - International Conference on Uncertainty in Structural Dynamics*.

M.R. Jones, T.J Rogers, N. Dervilis, K. Worden, E.J Cross, 2019. Assessing the likelihood of damage at the start of a structural health monitoring campaign. *Proceedings of 12th International Workshop on Structural Health Monitoring, Stanford.*

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	A brief overview of acoustic emission . . . . .	3
1.1.1	Methodology for AE localisation . . . . .	3
1.2	Localisation within the SHM framework . . . . .	6
1.3	Towards probabilistic source localisation . . . . .	7
1.4	Motivation: Current challenges of damage localisation within the SHM framework . . . . .	10
1.4.1	Thesis contribution . . . . .	11
<b>2</b>	<b>AE localisation from a data-driven perspective</b>	<b>13</b>
2.1	Acoustic emission: Basic theory . . . . .	13
2.1.1	The source . . . . .	15
2.1.2	Propagation . . . . .	17
2.1.3	Detection . . . . .	22
2.2	Features of AE waveforms . . . . .	23
2.3	Time of arrival extraction . . . . .	25
2.4	Traditional approaches for source localisation . . . . .	28
2.5	Advances in AE localisation . . . . .	32
2.6	Towards data-driven localisation . . . . .	34
2.7	Experimental case study . . . . .	37
2.8	Overview of the thesis . . . . .	39
<b>3</b>	<b>Gaussian processes for probabilistic AE localisation</b>	<b>41</b>
3.1	An introduction to Gaussian process regression . . . . .	41
3.1.1	Kernel selection . . . . .	46
3.1.2	Learning hyperparameters . . . . .	49
3.1.3	Performance metrics . . . . .	50
3.2	AE source localisation using Gaussian processes . . . . .	51
3.2.1	Outline of localisation framework . . . . .	51

3.2.2	Learning $\Delta T$ maps . . . . .	52
3.2.3	Source location likelihood . . . . .	54
3.3	Results and discussion . . . . .	60
3.3.1	Single source location prediction . . . . .	64
3.3.2	Sensitivity to training grid size . . . . .	64
3.3.3	Comparison with existing methods . . . . .	66
3.4	Conclusions . . . . .	68
<b>4</b>	<b>Heteroscedastic Gaussian processes for enhanced uncertainty quantification</b>	<b>71</b>
4.1	Spatially-dependent noise in $\Delta T$ values . . . . .	72
4.2	Heteroscedasticity in Gaussian process models . . . . .	75
4.3	Learning $\Delta T$ maps with a heteroscedastic model . . . . .	79
4.4	Uncertainty in $\Delta T$ maps with a heteroscedastic model . . . . .	82
4.5	Localisation of AE events . . . . .	86
4.6	Conclusions . . . . .	95
<b>5</b>	<b>Physical constraints for physics-informed onset time mapping</b>	<b>97</b>
5.1	Introduction . . . . .	98
5.2	Constrained Gaussian processes . . . . .	100
5.3	Physical constraints for acoustic emission time of flight mapping . . .	103
5.4	Results and discussion . . . . .	106
5.4.1	Sparse training data available across the whole structure . . .	109
5.4.2	Training data from partial structure coverage . . . . .	113
5.5	Conclusions . . . . .	120
<b>6</b>	<b>AE localisation for wind turbine bearings</b>	<b>123</b>
6.1	AE localisation within bearings . . . . .	124
6.2	Generalising to non-Euclidean domains . . . . .	126
6.3	Case study . . . . .	127
6.3.1	Exploration of the data . . . . .	129
6.4	Localising AE events in a bearing . . . . .	132
6.4.1	Learning $\Delta T$ maps . . . . .	133
6.4.2	Source localisation . . . . .	136
6.5	Conclusions . . . . .	143
<b>7</b>	<b>Conclusions and future work</b>	<b>145</b>
7.1	Localisation in complex structures . . . . .	146

7.2	Towards a more-informative noise model . . . . .	147
7.3	Physics-informed machine learning for AE localisation . . . . .	148
7.4	Localisation beyond Euclidean geometries . . . . .	149
7.5	Future work . . . . .	150
<b>A</b>	<b>Basic Bayesian probability theory</b>	<b>155</b>
	<b>Bibliography</b>	<b>157</b>

# INTRODUCTION

The capture and use of acoustic emission (AE) measurements has long been a popular technique in the non-destructive evaluation community for assessing and monitoring damage [1, 2]. Likened to “mini earthquakes”, the release of AE can be considered as a form of microseismic activity within a structure, generated as materials undergo failure mechanisms [3]. Within the last decade or so, the use of these measurements have become more routinely adopted by structural health monitoring (SHM) practitioners, which is a field of research concerned with the development of tools for continuously monitoring and assessing the structural integrity of civil infrastructure and high-value assets [4]. SHM is commonly approached from a data-based perspective, where the general objective is to first acquire some data that characterises the response of a structure of interest. From these data, features sensitive to damage are extracted, and used as inputs into some statistical model/machine learner.

In traditional SHM, measurements that represent the lower frequency behaviour of a structure (typically up to 1KHz), more commonly referred to as vibration-based techniques [5], have largely been favoured due to the ability to readily extract dynamic properties such as structural natural frequencies and modal shapes, which have been demonstrated to be sensitive to damage [4]. Acoustic emissions, however, characterise response mechanisms at much higher frequencies (in the rough range of 50KHz to 1MHz), and are much more sensitive to smaller scale changes in the structure. Consequently, these types of measurements are able to offer a more sensitive indicator of damage, allowing defects to be detected sooner (assuming an appropriate sensor setup). Another advantage offered by AE is the ability to localise

signals, allowing damage to not only be detected, but also tracked to its origin, albeit with an increased burden on data collection and handling. Such a functionality is key in the pursuit of improved structural health monitoring strategies that provide increasing insight into the damage state of the structure, where the desire to make more informed operation and maintenance necessitates SHM practitioners to go beyond simply detecting the presence of damage.

Although there is much literature concerned with the use of acoustic emission for localising damage, many of these approaches are only valid for simple structures where AE signals are transmitted without obstruction and at a constant speed. In many real engineering structures, these idealised conditions are not valid, and so there is still the critical challenge of ensuring that defects may be accurately and reliably localised in structures representative of those found in the field. Take an aircraft wing as an example; present on the wing will be many geometrical features such as bolts, rivets, tighteners etc. As an elastic wave travels through these types of features, the behaviour of the transmitted signal will be altered. If the structure is made up of composite material, such as in a reinforced fiberglass wind turbine blade, due to the nature of the composition of layers, material properties such as wave speed are directionally dependent, introducing multiple wave speeds which may also be unknown. More details on these challenges are provided in Chapters 2.5 and 2.6.

Another area that has had limited attention directed towards it is quantifying a measure of uncertainty within AE localisation, with the output of many localisation models that are currently available in the literature exclusively deterministic. In a physical world where the collection of data is itself inherently stochastic, with the form of the subsequently developed model never exact, it is of the author's belief that some attempt should be made to capture this uncertainty appropriately. Quantification of uncertainty is also essential when one adopts a risk-based approach to SHM [6, 7]. This thesis, therefore, attempts to address these limitations and develops a novel, probabilistic framework for acoustic emission localisation for use on complex structures.

Within this chapter, a short introduction into acoustic emission is given, before providing the reader with an overview of the research field of SHM, and how the task of (acoustic emission) localisation falls within the remit of the SHM framework. Next, the transition to a probabilistic view of localisation will be discussed, before concluding the chapter with the key contributions made by this thesis.

## 1.1 A brief overview of acoustic emission

In engineering structures, materials will often undergo sudden stress redistributions, causing a rapid release of small packets of energy. Consider the initiation of a crack; as the crack grows, localised strain energy within the material is released [3], generating discrete bursts of energy which then manifest as high frequency elastic stress waves that will propagate through a structure. These packets of energy are what are known as acoustic emissions, where the duration of one burst is known as an “event”. If piezoelectric sensors are attached to the medium in which these waves travel, the emission can be captured in a digitised manner, with the overall process of generating and recording acoustic emissions referred to as acoustic emission monitoring.

Although crack growth serves as a typical AE source, any irreversible change within the internal structure of a material can generate acoustic emission. Given that many physical mechanisms related to structural damage will exhibit this type of change (this will be discussed in Chapter 2), acoustic emission monitoring can be very useful for identifying damage in engineering components and structures, particularly due to its capabilities in detecting small scale defects. However, to afford this increased sensitivity requires high data sampling rates, which can lead to the generation of millions of data samples after only several seconds of acquiring data. Although an important topic in acoustic emission research, the data storage problem is outside the scope of this thesis, with the assumption made that one has access to distinct events related to acoustic emission activity. Examples of how one may extract events from raw time signatures can be found in interesting works such as [8–10].

Another advantage of AE monitoring is the ability to localise the origin of the source, given an appropriate sensor setup. A brief overview of how localisation is achieved is given below.

### 1.1.1 Methodology for AE localisation

In a basic testing scenario, imagine the generation of an AE signal at the centre of a solid plate. Given that the emissions will be emitted radially in all directions, there will be contour lines of constant propagation time, located at a radial distance of  $ct$  (where  $c$  is wave speed and  $t$  is time travelled) from the centre of the source. Figure

1.1 provides a visualisation of propagating AE wavefronts, with time of arrival (TOA) contours represented by the dotted lines.

If one were to mount three piezoelectric sensors in the area surrounding the source, then the observed arrival time of the source at sensor  $i$  will form a single contour line at a distance  $D_i = ct_i$ . There will then be a unique intersection point of the contour lines located at the source position, which can be expressed analytically through simple trigonometry. Figure 1.2 shows a graphical representation of this process, where the blue highlighted region marks the intersection of the contours, and therefore the position of the AE event.

Now, although Figure 1.2 serves as an intuitive visualisation of the basic principles of source localisation, in real structures localisation faces many more challenges than in this simple example. When working directly with time of arrival, an exact time at which the source was generated is required. As acoustic emission is a *passive* technique rather than *active*<sup>1</sup>, then, generally, this information will not be available. Additionally, it is assumed that the wave speed,  $c$ , is constant and known, as well as the existence of a direct propagation path between the sources and sensor(s). To address the first point is fairly trivial. Instead of dealing with absolute arrival times, one can move to working with the *difference*-in arrival times between sensor pairings, removing the dependence on time of source generation. The localisation problem is still solved through the calculation of a common intersection point of contours, except now these contours are representative of constant difference-in arrival times between sensor pairs. The remaining assumptions are not necessarily able to be dealt with in such a straightforward manner, requiring a more detailed treatment, which will be given in the following chapter.

---

<sup>1</sup>Passive sources refer to those that have been generated by the medium itself, such as the growth of a defect. An active source is one that has been pulsed within the solid medium by some mechanism that is separate to the structure, such as a transducer. More details provided in Chapter 2.

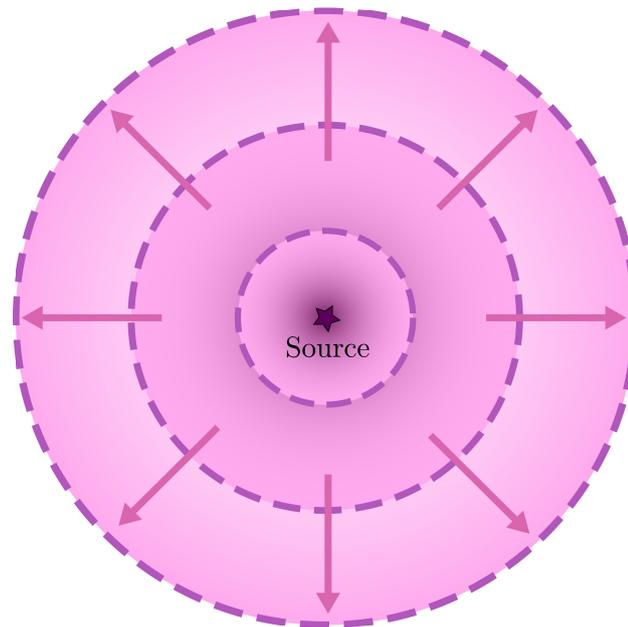


Figure 1.1: Propagation of acoustic emission waves from source.

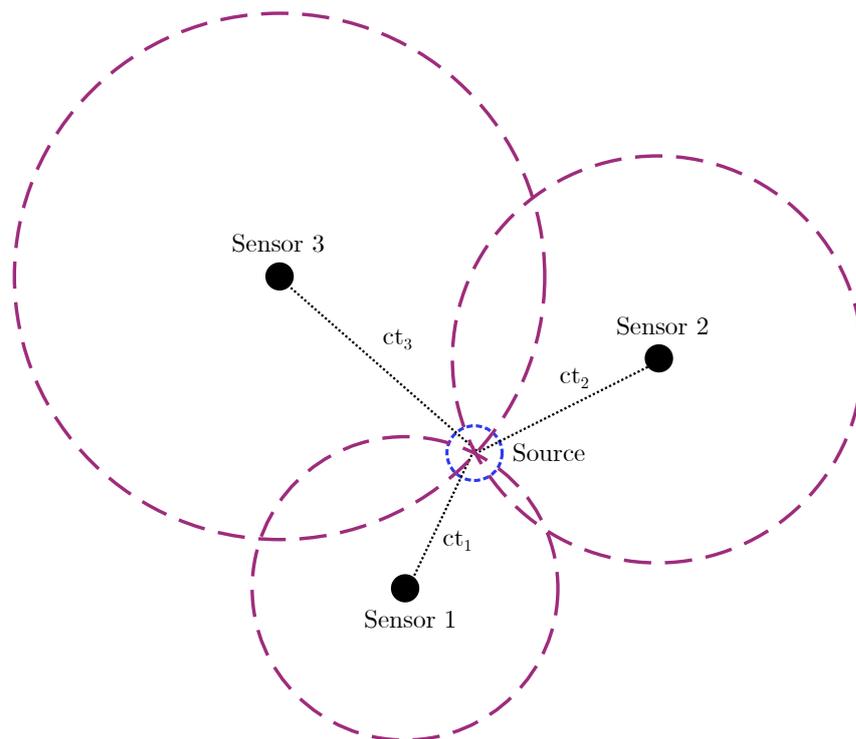


Figure 1.2: Basic illustration of acoustic emission localisation.

## 1.2 Localisation within the SHM framework

Structural health monitoring is defined as the process of implementing a strategy to assess the health or condition of a structure, generally for civil infrastructure and other high value engineering assets such as aircraft [4]. The distinct difference between non-destructive evaluation (NDE) and structural health monitoring is that whilst NDE is carried out retrospectively, SHM is generally conducted in an online manner, with the aim that an SHM system provides feedback whilst a structure or engineering component is in active operation, ideally for its entire operational life. It is also useful to note the close connection between SHM and the field of system identification [11], which in a broad sense, is a topic concerned with understanding and modelling dynamical systems from measurement data, with some interesting examples found in [12–15].

Structural health monitoring is often viewed as a multi-level process, defined in a hierarchical sense by the seminal work of Rytter [16], and later extended by Worden and Dulieu-Barton [17], as:

1. Detection - does damage exist within the system?
2. Localisation - where is the damage?
3. Classification - what type of damage is present?
4. Assessment - how severe is the damage?
5. Prognosis - what is the remaining useful life of the system?

Owed to the increased maturity of sensing hardware, it is now possible to collect vast amounts of data that capture the dynamic response of structures of interest. It is therefore possible to utilise these data to develop models for health monitoring purposes. These approaches that view SHM from a *data-driven* perspective treat the damage identification paradigm as a problem of pattern recognition [4, 18], where a relationship between the physical structure and extracted data is learnt in order to build a statistical model. Ultimately, one wishes to extract some finite amount of data, and associate it with a class label. The number and physical meaning of these labels will be dependent upon which level of the damage identification hierarchy is of interest, along with the desired depth of insight one wishes to obtain from

the monitoring algorithm; for example, whether the problem is one of classifying a structure as damaged or undamaged (binary-class problem), or formally defining multiple types of damage, in addition to an undamaged state (greater than two-class problem).

The task of damage identification is most commonly treated as a two-class problem at the first level of the hierarchy, where a health monitoring practitioner is concerned with assessing whether damage is present in a structure or not. Much progress has been made in developing damage detection strategies, with examples of learning in either an *unsupervised* or *semi-supervised* manner [19–24] (none or partial training labels), removing environmental and operational trends [25–27], as well as learning online [28, 29]. As progress is made further up the hierarchy, the complexity of the task grows, relying on more informative training data. For instance, in order to classify damage, then examples of data pertaining to each state of interest is needed, requiring much more of a comprehensive training data set than in the case of damage detection, where only data from an undamaged condition is necessary. However, with additional complexity comes increased knowledge regarding the nature of the damage that has been identified. For instance, understanding the type of damage that has occurred in a structure gives more insight to an engineer than if only the presence of damage is identified. It is, therefore, desirable for progress to be made beyond the damage detection stage, where much of current SHM-related research has been directed, and towards the upper levels of the hierarchy, to reach what is perhaps the “holy grail” of SHM; estimating (accurately) remaining useful life.

It is generally not the case that, other than for the problem of damage detection, that a particular level of the hierarchy can be approached in isolation, with each stage usually requiring that all lower-level information is available [4]. As such, in the quest to develop more insightful and reliable health monitoring algorithms, it is crucial that each stage of the damage assessment problem, such as localisation, is sufficiently addressed.

### 1.3 Towards probabilistic source localisation

When modelling physical processes, there are many potential sources of *uncertainty* that may arise and influence the problems that SHM researchers and practitioners are concerned with. For instance, the phenomena that interact with our engineering

structures such as wind and wave loading are often stochastic, with acquired measurements that characterise the structures' response to these inputs contaminated by some level of noise. Where a model fails to account for certain physical effects, either by simplifying the true underlying mechanics or through a lack of data that spans the full operational envelope a structure is exposed to, also results in the introduction of uncertainty. Despite there being many examples of sources of uncertainty within engineering systems, a significant number of SHM solutions remain deterministic. A convenient approach for quantifying a measure over the uncertainty present is to work within a *probabilistic* framework, where rather than assuming some deterministic model form, probability distributions are used to represent the belief that exact values are not known, but instead there exists a distribution over a range of possible values, with the occurrence of a particular value quantified by a probability. Through the use of a probabilistic-based workflow, uncertainty can be handled in a statistically principled and rigorous manner, whilst also providing an operator with a deeper insight into the predictions made by an SHM model.

For the source localisation problem, under the use of a deterministic model, a location prediction will generally be a fixed, single position value. But what happens in the scenario that a prediction is made in a region where training data is limited, or sensor coverage is poor? How is the noise that is inevitably present on any measured data accounted for? Here, by adopting a probabilistic model, these sources of uncertainty may be fed into the model in a principled manner using fairly simple probability theory.

In the world of probability theory, there are two schools of thought surrounding how a probability distribution should be formed and interpreted; *Bayesian* and *frequentist* statistics. In frequentist inference, a statistician will treat the probability of an event as the limiting value of the frequency of occurrence in a sequence of trials. In other words, frequentist probabilities are constructed purely from the evidence provided by observed data. On the other hand, the Bayesian approach allows the inclusion of any pre-existing beliefs held about the event to be included through the use of a *prior* distribution. The prior is then combined with observed data to obtain a *posterior* distribution; a probability that represents an update on any subjective beliefs held after observing data. For an engineer, the Bayesian workflow is particularly attractive due to the fact this prior belief is often available in the form of some physical insight held about the problem at hand, obtained from our engineering understanding. Additionally, a Bayesian treatment allows uncertainty in the model

to be expressed, with the probability of an event treated as a random variable instead of a single, real number, as with a frequentist approach. Confidence in the model can be increased as more data are observed, with the ability to recursively update the posterior. This formulation lends itself particularly well to SHM in the sense that once a model has been constructed, it is desirable to be able to continuously update the model as extra data becomes available.

To implement Bayesian theory, as the name suggests, relies on the use of Bayes' theorem, which is written as,

$$p(\mathbf{y}|\mathbf{x}) = \frac{p(\mathbf{x}|\mathbf{y})p(\mathbf{y})}{p(\mathbf{x})}. \quad (1.1)$$

At its core, Bayes' theorem is simply the application of the product rule<sup>2</sup>. For the interested reader, a derivation can be found in Appendix A.

In Bayes' theorem, one generally begins with the prior,  $p(\mathbf{y})$ , which as stated above, encodes any available information available to the modeler before any data are observed.  $p(\mathbf{x}|\mathbf{y})$  is the likelihood, which is the probability of observing the data given a model, whilst  $p(\mathbf{x})$  is the marginal likelihood, often referred to as the model evidence, providing a measure for how likely the data are under all possible models. These terms are then combined to compute the posterior distribution,  $p(\mathbf{y}|\mathbf{x})$ , which represents our updated belief regarding a given quantity having observed some data.

The main source of computational burden when computing the posterior is evaluation of the model evidence, which includes solving an integral term; notorious for having no known polynomial time algorithm for exact calculation. To overcome this problem, *conjugate priors* are regularly employed such that the evaluation of integrals as part of assessing the model evidence can be computed in closed form through simple parameter updates. Priors that enable this condition are said to be conjugate to the likelihood, where the resulting posterior is of the same family of the prior. As the Normal distribution is conjugate to itself, it is common for Gaussian priors and likelihoods to be employed. When a Gaussian assumption is not appropriate, one can turn to numerical techniques such as Markov chain Monte Carlo sampling, or approximate inference techniques such as those available from the field of variational inference. The latter case will be demonstrated in Chapter 4.

---

<sup>2</sup>Some text books may cite the chain rule instead of the product rule. The chain rule is just the repeated application of the product rule, and can be viewed as a generalisation of the product rule.

## 1.4 Motivation: Current challenges of damage localisation within the SHM framework

The automated localisation of damage in structures is a critical ingredient in the path towards predictive or condition-based monitoring of high value structures, paving the way for progression up Rytter's hierarchy. The development of accurate and robust localisation methods is therefore imperative. Although acoustic emission serves as a useful tool in the localisation problem, there is still significant challenge in developing localisation models, whether analytical, numerical, or data-based, in structures that are representative of those found in the engineering world. Whilst simple models work well where wave propagation is simple, for example, where wave speed is constant with an unobstructed propagation path, extending beyond these conditions is challenging.

Over the course of the last decade, there has been notable progression made towards developing localisation methods in complex structures; some examples include [30–32]. A unifying theme across these works is that all the authors adopt a data-driven approach as a means for overcoming the challenge of localising AE in structures that contain complex propagation paths, for instance, on a landing gear component [30] and riveted panels [32]. However, in these and other related works, there has been either partial or no attempt at rigorously quantifying an uncertainty or confidence associated with subsequent predictions, where the statistical model does not appropriately reflect how uncertainty manifests itself in the localisation problem. It is the author's belief that such a functionality is essential, handling processes that are inherently stochastic that arise in the AE localisation problem, as well as providing greater insight into the predictions made. By approaching the localisation task probabilistically, one can quantify uncertainty arising from sources such as onset time estimation, lack of data and the sensor measurements themselves, in terms of both measurement noise and how well a particular sensor may cover parts of the structure, with attempts made to understand and account for such behaviour leading to more robust and rich models. The ability to handle uncertainty in models that can generalise to a wide range of geometries and materials will mark significant advancement in localisation models for SHM.

Another limitation that exists with current methods that take a data-driven approach to acoustic emission localisation is that they require the collection of an extensive set

of artificial acoustic emission measurements. For many structures, particularly those that are large in area with many geometrical features, this requirement results in a cumbersome data collection campaign and limits the practicality of the methods. Not only do many datum points need to be collected to generate a suitable training set, but the process of initiating artificial AE is itself very time-consuming due to the often manual nature of generating the elastic waves (see Chapter 2.1.1 for more details). To improve the viability of deploying these types of localisation models for predictive and condition-based maintenance, it is crucial that the methods can be used where training data is sparse and few.

### 1.4.1 Thesis contribution

In this thesis, a novel, probabilistic framework for localising acoustic emission in complex structures is developed. By first learning a relationship between source location and difference-in arrival times ( $\Delta T$ ) for multiple sensor pairs, the likelihood of emission location of AE events with an unknown origin can be calculated. By quantifying this likelihood across the surface of a structure, a map is returned that indicates probable areas of source location, providing both a maximum likelihood source location estimate, as well as an intuitive visualisation of the associated confidence. A transition to a heteroscedastic model is then made such that an input-dependent noise may be captured, allowing for a more expressive model that can account for the dependence on prediction accuracy arising from the locations of particular sensors. Next, consideration is given to how cumbersome and often impractical data acquisition requirements can be addressed through the use of combining data-driven learning with physical knowledge; specifically, boundary condition constraints. Finally, a condition monitoring case study is investigated, where the localisation procedure is applied to a roller element bearing, presenting a solution for localisation outside of Euclidean spaces.

In the next chapter, an introduction to acoustic emission theory will be given, as well as providing an overview of the current literature concerned with acoustic emission localisation techniques, and the need for a probabilistic acoustic emission localisation framework.



# AE LOCALISATION FROM A DATA-DRIVEN PERSPECTIVE

### Highlights:

- *An introduction to acoustic emission theory is provided, in addition to an overview of the current state-of-the-art acoustic emission localisation techniques.*
- *The need for a probabilistic acoustic emission localisation framework is motivated.*

## 2.1 Acoustic emission: Basic theory

The discovery of acoustic emissions is generally attributed to Kaiser [33], who, in 1950, found that when a material was loaded to a previously unseen stress, acoustic waves within the ultrasonic frequency range are emitted. In practice, this range is usually around 50Khz to 1Mhz. These acoustic waves are the manifestation of rapid releases of energy from a localised source [34]. Acoustic emission differs from other ultrasonic testing methods in that it is a passive approach; the waves are generated through a release of energy within the material itself, as opposed to being projected into the structure through the use of a transducer, as is the case with active approaches such as guided wave testing [35–38].

As described by Vallen [39], there are 7 basic principles to AE testing, collectively defined as the *AE Process Chain*. This chain is detailed as follows:

1. Acoustic emission mechanism is generated, either on application of dynamic load or excited via machine/system operation.
2. Elastic energy associated with source mechanism is released.
3. Elastic energy (the waves) propagates through a solid medium to a receiving sensor.
4. Either a single or array of sensors detect the waves, converting mechanical energy into an electrical signal.
5. The signal is interpreted by some acquisition equipment that is then subsequently extracted.
6. Extracted data are plotted and displayed.
7. All recorded data are evaluated.

As the AE community continues developing methods to process and interpret acoustic emission data, it is desirable to remove a dependence on an expert simply plotting raw data to make condition-related assessments, and instead, feed that raw data into models (whether physics or data-based), where the output of a given model can provide deeper insight than is possible with raw data alone. Not only will an operator have a better understanding about what the recorded data means, but they will also be better positioned to make maintenance and planning decisions, potentially in an autonomous manner. With these considerations in mind, steps 6 and 7 ought to be revised. In the context of SHM, an update to the latter part of the process chain may read something along the lines of:

6. The acquired data are fed into a (task-specific) model.
7. Outputs of the model are utilised for damage assessment (any single step or combination of from the SHM hierarchy). These outputs are either interpreted directly by an operator, or used as a prerequisite for an autonomous decision-making agent.

In terms of which of these principles govern the actual physical behaviour of acoustic emissions, these are steps 1/2, 3 and 4; the source, propagation and detection [40]. The remainder of Chapter 2.1 will consider each of these components individually.

### 2.1.1 The source

Although there exists a general definition for acoustic emission that unifies the type of behaviour all AE sources generate, the exact mechanism will vary depending on the type of the source, of which there are many.

When defining an AE signal, potentially the most fundamental characterisation that can be made is if the signal is transient or continuous. For typical damage mechanisms such as crack growth, fibre breakage or plastic deformation, these types of sources will produce small bursts of energy that are classed as transient AE [41, 42]. Continuous sources, however, will emit a continual level of AE over a sustained period of time, with less well defined start and end times. For sources continuous in nature, it is common for these types of emissions to be categorised as noise of the system [43], although in some monitoring cases, such as rotating machinery [44, 45] and fluid leakage [46], they can be used to give an indication of the overall damage level of a structure.

Some common examples of acoustic emission sources include:

- Crack formation and growth
- Plastic deformation
- Fatigue (crack growth)
- External impact
- Fluid leakage
- Composite fibre cracking and failure
- Composite delamination
- Wear
- Corrosion
- Lubricant failure

In addition to those listed, there are also numerous ways to generate artificial AE sources as a means for replicating the physical behaviour of real emissions. Artificial sources are very useful for determining the range of sensitivity that a sensor in a particular position has, testing the quality of the couplant between the sensor and the structure being monitored, as well as for building a database of acoustic emission signals from known positions, which can be very useful for localisation (see Chapter

2.6). It is often also possible to assess the accuracy of many types of AE models using artificial sources.

Perhaps the most common form of artificial AE is that of the Hsu-Nielsen (H-N) source [47, 48], consisting of breaking a pencil lead at the surface of a structure of interest. This approach is often employed by AE practitioners as it is a very cheap and convenient way to generate artificial AE quickly. However, due to the variance induced through human operation, it can be difficult to guarantee consistency between different H-N tests. Therefore, to aid with reproducibility, a small disc that sits at the head of the pencil can be placed against the surface. This technique is employed to acquire the data set that is used in Chapter 6.

A more repeatable approach is to use an actuated piezoelectric transducer to transmit an ultrasonic wave. Here, one can essentially use a transducer to both send and receive AE signals, with the ability to specify the signal that is pulsed from the transducer. Inducing artificial sources in this way is more time consuming than the H-N method, however, requiring a transducer to be mounted at any location one wishes to excite AE at, practically limiting the number of locations where emissions can be generated. However, for situations in which artificial sources only need to be excited at a small number of locations, pulsing transducers provide highly repeatable sources across a range of frequencies and types of excitation. Care should be taken that the frequency bandwidth of the emitted signal matches what is physically expected from the source one is trying to simulate.

A final example of artificially generating a source is by way of laser thermoelastic expansion. By pulsing a laser at the surface of a structure, the surface is rapidly heated and then cooled, creating an elastic stress wave that mimics that of an AE event. A distinct advantage here is that the laser can be attached to an electro-mechanical positioning device, making the process of generating the sources far more autonomous. This approach is employed to generate the artificial sources contained in the experimental case study detailed in Chapter 2.7, which is the data set used in Chapters 3 - 5.

It should be mentioned that each mechanism discussed above for inducing artificial AE is limited to excitation at the surface of a structure. For some damage mechanisms such as fatigue cracking in bearings, defects may form below the surface [49, 50], resulting in the generation of subsurface acoustic emissions. A potential option one has here is to embed a transducer capable of pulsing an ultrasonic signal inside

of a structure, for instance, during the lay-up period if working with composites. Mechanical compression of the structure has also been suggested by [50] as a means for inducing subsurface AE, however, generating consistent and repeatable AE mechanisms in this manner is challenging. Additionally, as the specimen is physically compressed, some level of damage may be introduced into the structure that may be undesirable if the objective is to only induce artificial sources.

### 2.1.2 Propagation

Once a source mechanism has generated an AE signal, the way in which the wave will propagate, such as its velocity, frequency, and direction in which oscillations of the wave move in relation to the propagation direction depend on factors such as the material properties and geometry of the medium. Where the propagation medium is an infinite body, there are only two types of propagating waves; *transverse* and *longitudinal* waves. Collectively, these two propagation modes are referred to as *bulk waves*, and exist when there is no interaction between the wave and a boundary. For instance, human speech produces bulk waves.

In a longitudinal wave, the direction of oscillatory motion is in parallel with the propagation direction, with their motion driven by compression forces, and thus, are sometimes referred to as pressure waves. The velocity of longitudinal waves is given as,

$$c_L = \sqrt{\frac{\Lambda + 2\mu}{\rho}}, \quad (2.1)$$

where  $\Lambda$  and  $\mu$  are Lamé constants, and  $\rho$  is the density of the material. Equivalently, the velocity of transverse waves is expressed as,

$$c_T = \sqrt{\frac{\mu}{\rho}}, \quad (2.2)$$

where the particle motion of the wave is perpendicular to direction of travel, and therefore requires the support of shear forces to propagate.

Where a surface is introduced, e.g. semi-infinite media, these two waves modes now no longer exist independently, resulting in an extra wave mode; a *Rayleigh wave*

[51]. Often termed surface waves due to their existence on the surface of elastic media, these waves will travel with longitudinal velocity  $v$  and transverse velocity  $w$ , expressed as,

$$v = A(re^{-qz} - 2sqe^{-sz}) \cos k(x - ct), \quad (2.3)$$

$$w = Aq(rq^{-qz} - 2e^{-sz}) \sin k(x - ct), \quad (2.4)$$

where  $k = \frac{2\pi}{\lambda}$  is the wavenumber,  $\lambda$  is the wavelength and  $A$  is the amplitude.  $x$  and  $z$  index the plane of motion, whilst  $q$ ,  $s$  and  $r$  are some constants defined in terms of ratios between the material bulk wave speed,  $c$ , and the transverse and longitudinal wave speeds,

$$q = \sqrt{1 - \left(\frac{c}{c_L}\right)^2}, \quad (2.5)$$

$$s = \sqrt{1 - \left(\frac{c}{c_T}\right)^2}, \quad (2.6)$$

$$r = 2 - \left(\frac{2}{c_T}\right)^2. \quad (2.7)$$

Through incorporation of known boundary solutions, a characteristic equation that constrains the possible solutions for the velocities (or wavenumbers) can be formed, from which it is possible to identify  $c_R$ , the Rayleigh wave velocity, by finding the real root [36].

When an additional boundary is added such that there are now two boundaries in parallel, the final wave type that will be discussed in this introductory chapter is formed; *Lamb waves* [52]. In the presence of two surfaces, the rate at which waves are reflected and interact increases, and forms additional surface waves, known as Lamb waves. These types of waves will exist as either a symmetric ( $S_n$ ) or asymmetric ( $A_n$ ) wave, where  $n$  refers to the particular order of the wave mode. In an isotropic medium, the characteristic equation for the symmetric modes is given as,

$$\frac{\tan(qh)}{\tan(ph)} = -\frac{4k^2pq}{(q^2 - k^2)^2}, \quad (2.8)$$

and for asymmetric modes,

$$\frac{\tan(qh)}{\tan(ph)} = -\frac{(q^2 - k^2)^2}{4k^2pq}, \quad (2.9)$$

where  $p$  and  $q$  equal,

$$p = \frac{f}{\sqrt{c_L^2 - c^2}}, \quad q = \frac{f}{\sqrt{c_T^2 - c^2}}, \quad (2.10)$$

and  $f$  is the central frequency and  $h$  is the plate half-thickness. With the introduction of a frequency term into the characteristic equations, it can be said that Lamb waves are *dispersive*; the speed at which the Lamb waves propagate, unlike Bulk and Rayleigh waves, is dependent upon the frequency. Solutions to these equations can be computed numerically, which leads to plots of the wave velocity with respect to frequency (or more commonly, frequency-thickness). These plots are termed *dispersion curves*.

Typically, dispersion curves are presented in relation to *wavenumber*, *phase velocity* or *group velocity*, and are readily interchangeable with frequency and wavenumber through the following equations,

$$c_{phase} = \frac{f}{k}, \quad c_{group} = \frac{df}{dk}. \quad (2.11)$$

At low frequency-thickness values, only the zeroth order modes ( $S_0$  and  $A_0$ ) exist. As frequency-thickness is increased, more modes appear, up to a theoretical limit of infinity. The zeroth order modes are often termed the fundamental modes due to the fact that they are the only modes which exist across all values of frequency-thicknesses.

To visualise the presence of dispersion in AE waveforms, Figures 2.1 and 2.2 plot the raw signal and corresponding time-frequency behaviour of two AE events, initiated in an aluminum plate by a laser pulse at two different distances away from the receiving sensor. Inspecting the two figures, it can be seen that the evolution of the frequency content throughout time significantly differs between the two signals. The cause driving this behaviour is the dispersive nature of the waves, with different segments of the waveform travelling at different speeds. For Figure 2.2, the event is positioned

further away from the sensor than Figure 2.1, with the signal having to propagate a greater distance before being received. Where the signal travels further, the received waveform is initially composed of a low magnitude signal, which is followed by an increase in amplitude and the arrival of the majority of the frequency content.

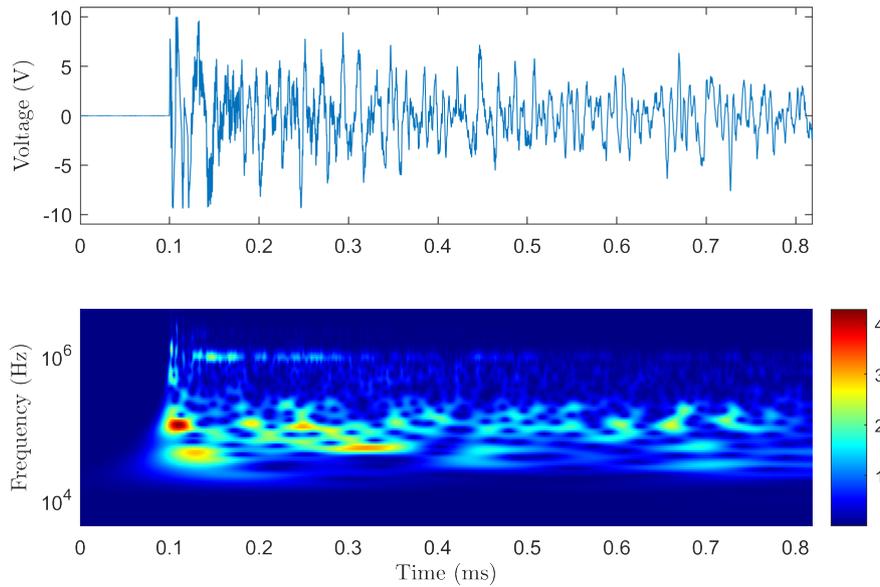


Figure 2.1: Time signal and frequency content of an AE signal initiated in an aluminium plate around 5mm away from the sensor.

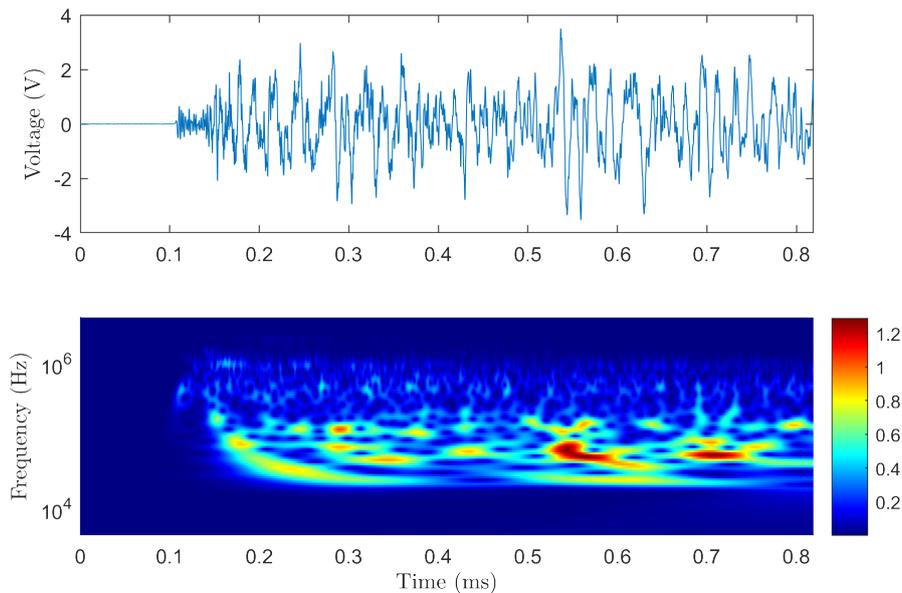


Figure 2.2: Time signal and frequency content of an AE signal initiated in an aluminium plate around 300mm from the sensor.

This behaviour demonstrates a separation between the arrival of the two fundamental wave modes, where the lower amplitude, faster travelling  $S_0$  mode arrives some time before the  $A_0$  mode. Where the event waveform has to travel less distance in Figure 2.1, there is no visible distinction between the arrival of the  $A_0$  and  $S_0$  modes, with less opportunity for dispersion to take effect.

With knowledge of the plate depth and a frequency range of interest, it is possible to use dispersion curves to visually estimate mode arrival times, as well as identify which wave modes will be physically present. An example of a dispersion curve for a 1.2mm thick aluminium plate is shown in Figure 2.3. However, whilst these curves are a useful tool, the characteristic equations do not account for indirect propagation paths, and so their use is somewhat limited for structures with complex geometrical features.

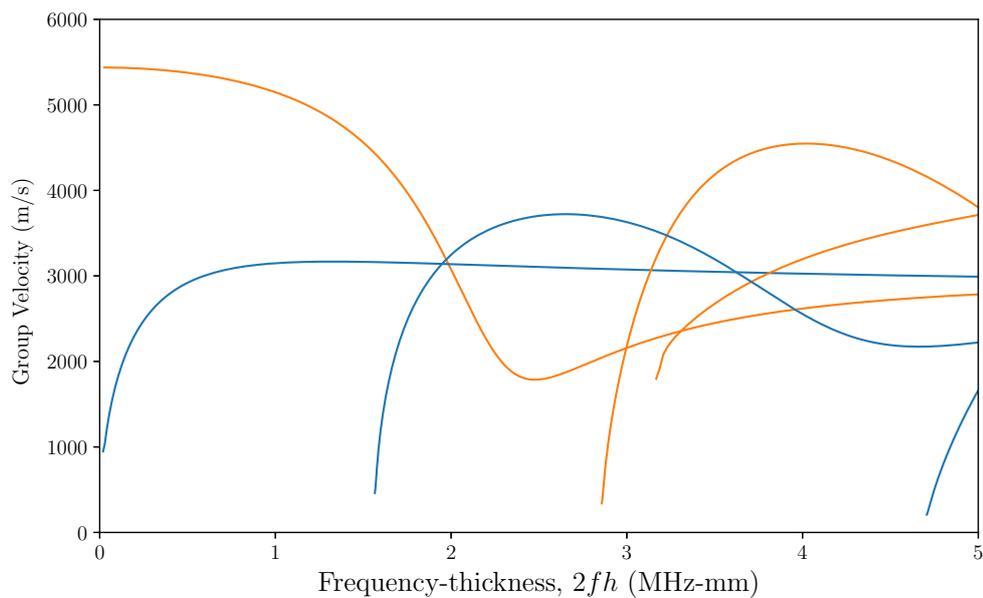


Figure 2.3: Dispersion curves for a 1.2mm thick aluminium plate, created using DISPERSE [53]. Orange lines are the symmetric modes, blue lines are the antisymmetric modes, with their order corresponding to the order of appearance along the direction of the  $x$  axis.

A final point to make regarding wave propagation is on the loss of energy as the wave travels outward from the origin of the source. The physical mechanism responsible for this phenomenon is known as *attenuation*. A primary cause of attenuation is geometric spreading, where as the wave propagates outwards, the wave front becomes larger, resulting in an increased spatial distribution of energy. As the

wave propagates, the energy that the wave possesses will also be absorbed into the propagation medium, with higher frequencies absorbed at a greater rate than lower frequencies. Absorption does not occur exclusively in the propagation medium. During the interaction between a boundary and the wave, some energy will be lost to any adjacent media. Diffraction and reflection at any boundary will also cause a loss in the peak amplitude of the incident wave. Finally, should the propagating wave be dispersive, as the higher frequency content of the signal will travel faster than that at a lower frequency, the resultant wavefront will spread, with the peak amplitude reduced. This behaviour can be seen in Figures 2.1 and 2.2.

All the theory that has been detailed in this section so far is only relevant to isotropic, homogenous materials, meaning that wave velocity has been assumed to be equal in all directions, with the propagation medium continuous (e.g. no discontinuities in the material from geometrical features). Although there has been considerable attention given to elastic wave theory in both anisotropic and inhomogeneous media, it is out of scope of this introductory chapter. For a full and in-depth treatment of more complex elastic wave theory, the interested reader is referred to the texts of Rose [54, 55] and Lowe [56].

### 2.1.3 Detection

The final step in the characterisation of acoustic emission signals is detecting/sensing the elastic waves. Figure 2.4 shows a typical configuration of an AE sensor. AE sensors generally operate by converting surface displacements caused by the motion of stress waves into electrical energy. A piezo-electric material such as Lead Zirconate Titanate (often shorted to PZT) is often used, which when deformed, produces a voltage proportional to the deformation it has undergone. The ratio between the output voltage and input displacement defines the sensitivity of the sensor.

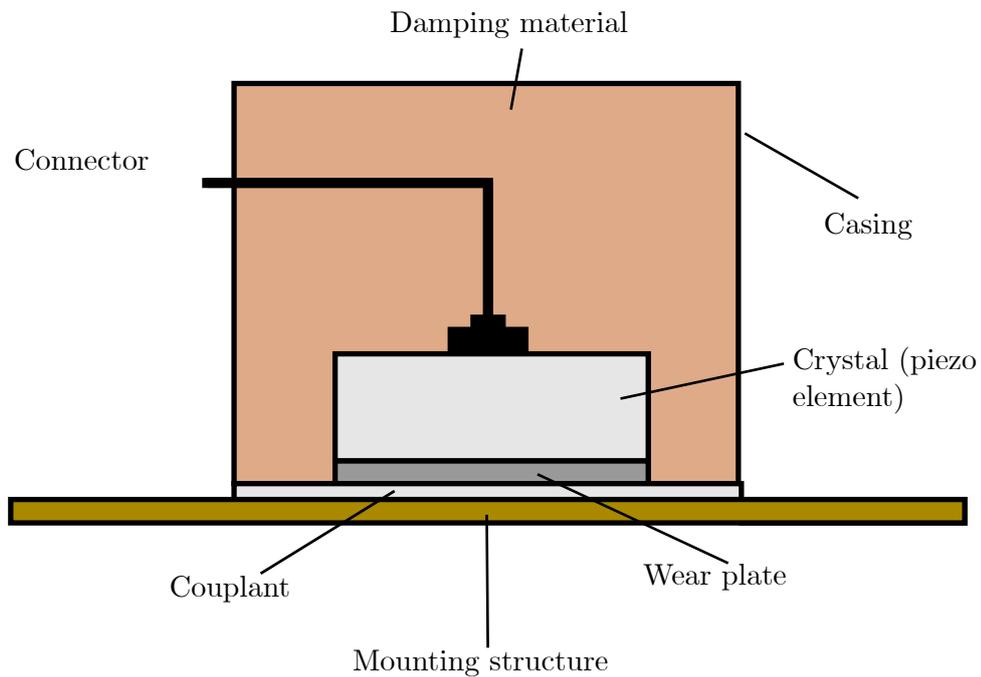


Figure 2.4: Schematic of a typical PZT sensor.

AE sensors can usually be categorised as broadband or resonant. In the case of a resonant sensor, the corresponding frequency spectra will exhibit a relatively sharp narrowband response around a particular (resonant) frequency. Therefore, when using this type of sensor, one should be careful there is no undesired frequency modulation. Matching the resonant frequency of a narrowband sensor with the frequency range of interest, can, however, be used to improve sensitivity. Alternatively, broadband sensors can be used, which have a much flatter frequency response. Sensitivity will generally be lower, however.

A final consideration to make with sensing hardware is the process of bonding the sensor to the material under test. Ensuring a good quality of contact between the sensor and the surface is paramount to obtaining maximal energy transfer between the sensor and the specimen. Common couplants include adhesives and grease.

## 2.2 Features of AE waveforms

For transient AE events, it is possible to encapsulate the behaviour of the waveform with a number of characteristic features. Some common examples include:

- **Peak Amplitude:** Maximum amplitude of the burst waveform.
- **Duration:** Total time of the event.
- **Rise Time:** Time taken to reach maximum amplitude from the start of the event.
- **Decay time:** Time between maximum amplitude and end of event.
- **MARSE:** Measured area under the rectified signal envelope. Can be viewed as the energy of the signal or a relative amplitude with respect to the duration of the event.
- **Power:** Energy of the waveform normalised by duration.
- **Ring Down Counts:** Number of times that the signal of an event crosses the detection threshold, and provides some indication of the frequency content of the signal. Also called counts.
- **First Threshold Crossing:** Time that an AE burst crosses a user-defined threshold. Sometimes used as time of arrival of event (at a receiving sensor).

A visual representation of a selection of these features is provided in Figure 2.5. In the case of continuously emitted AE, it is much harder to extract individual bursts from these types of emissions. Instead, it is common to work with statistical features of the entire raw time signal, such as root-mean-square, variance and other higher order statistical moments.

It is also possible to obtain insight into the frequency content of the acoustic emission signals by shifting the captured time signals into the frequency domain using tools such as short time Fourier transforms or wavelet transforms. The work considered here requires only the use of time domain-based features, and so a detailed analysis of features derived in the frequency domain is out of scope. However, more details on this perspective can be found in [57, 58].

Of critical importance to the work in this thesis is the time of arrival of the AE events. Given that an AE waveform will be emitted from the origin of the source before propagating to a nearby sensor where it is then received, the travel time is dependent upon the distance between the position of the source and the sensor. Therefore, the arrival time provides a measure that is sensitive to the position of the source, and so can be used as a feature in localisation models.

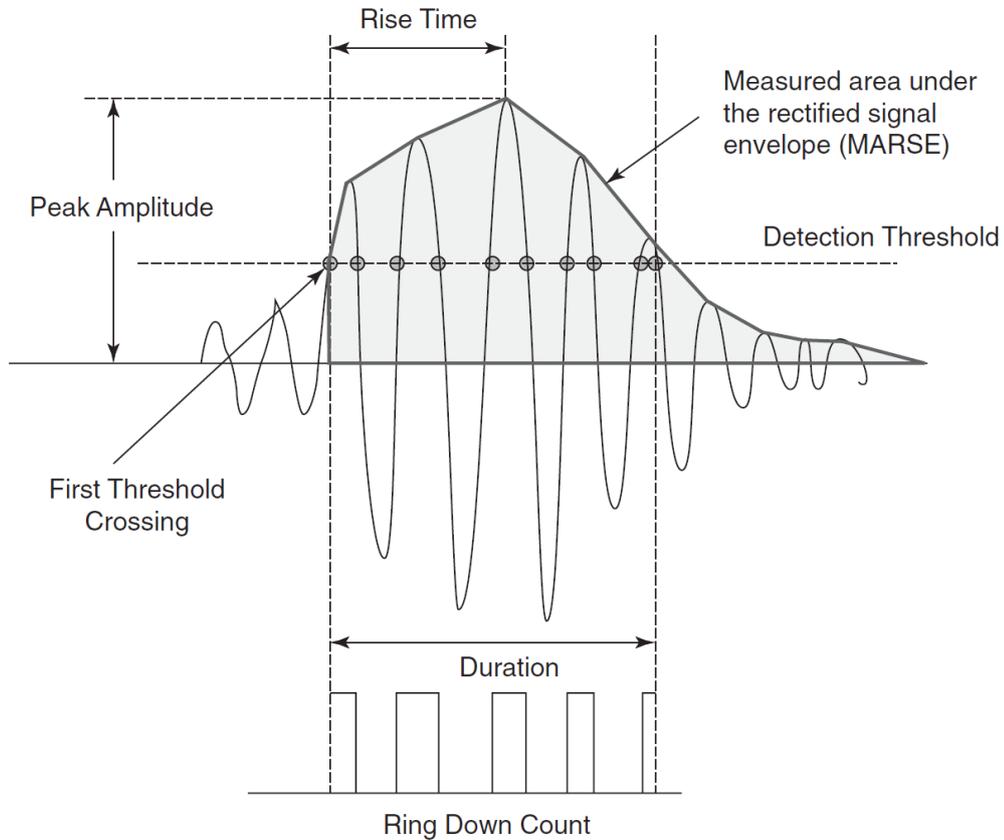


Figure 2.5: Some common features extracted from an AE event [4].

## 2.3 Time of arrival extraction

The identification of the arrival time of a waveform does, in itself, pose a challenge. For an incoming AE waveform, accurately detecting the time that the event arrives at a receiving sensor is of critical importance to being able to retrieve the origin of the source. It is, therefore, imperative that there are tools available to extract onset times so that localisation of the received signals is viable, regardless of whether the eventual localisation model is analytical, numerical or data-driven.

Perhaps the simplest approach for identifying arrival times is the use of a hard threshold, where the point at which the amplitude of the AE data crosses a predetermined value is set as the arrival time. Although straightforward to implement, in most practical situations, it will be unclear as to what the value of the threshold should be. One possible suggestion is to fix the threshold at  $n$  standard deviations of the data. For instance, if the data is assumed to be Gaussian, setting a threshold

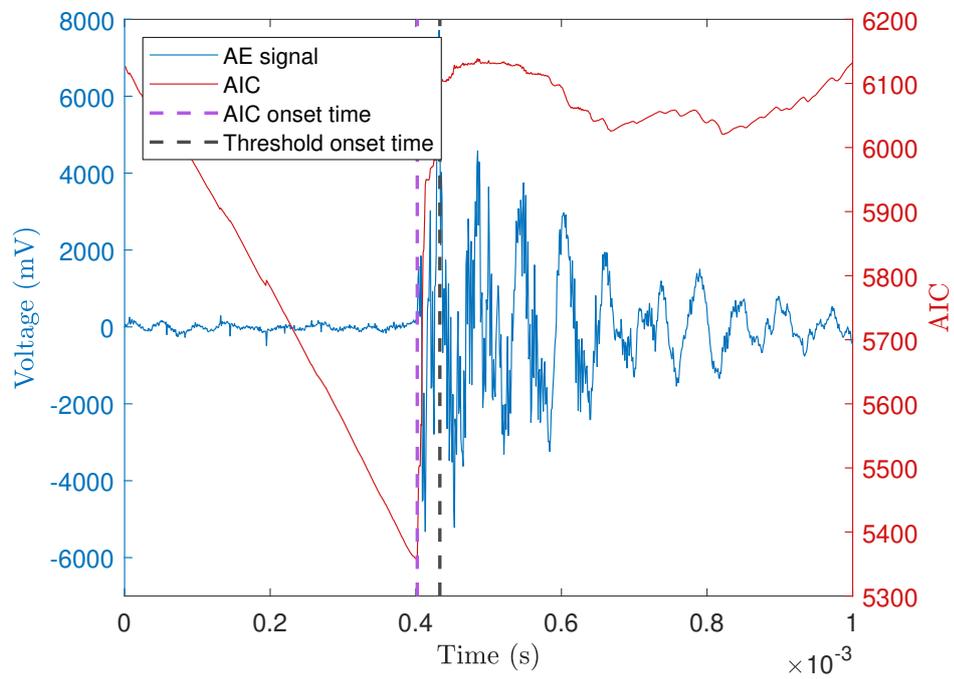
of three standard deviations corresponds to the expectation that only  $\sim 0.3\%$  of the measurements in a given sample of data are from an emission event. However, it is unlikely that one will have prior knowledge of how many hits are expected to be contained in incoming data, and so the use of a hard threshold can be somewhat arbitrary. Even where a reasonable estimate can be made, hard threshold approaches often miss the arrival of the faster but lower amplitude  $S_0$  mode [59], potentially leading to significant arrival time errors, particularly when there exists a big gap in the arrival time of the two wave modes.

In this thesis, an onset picking approach that is based on the use of the Akaike Information Criterion (AIC) is chosen, where the incoming signal is modelled as an autoregressive (AR) process. Initially developed as an approach to predict the optimal order of an AR model in the analysis of seismic signals, it was demonstrated in the works of [60, 61] that a time series may be divided into two local segments, each modelled as an AR process. Drawing an equivalence between this analysis of seismic time series data and the AE onset time problem, a recorded AE signal can first be thought of as containing some low-level noise produced by the acquisition equipment. As an AE event occurs and is received at a sensing transducer, the signal content becomes very different. The signal just before and just after the onset of an AE event can, therefore, be considered as two separate time series, where the AIC picker is used to determine the transition between the two AR processes. Mathematically, the AIC picker is defined as [62, 63]

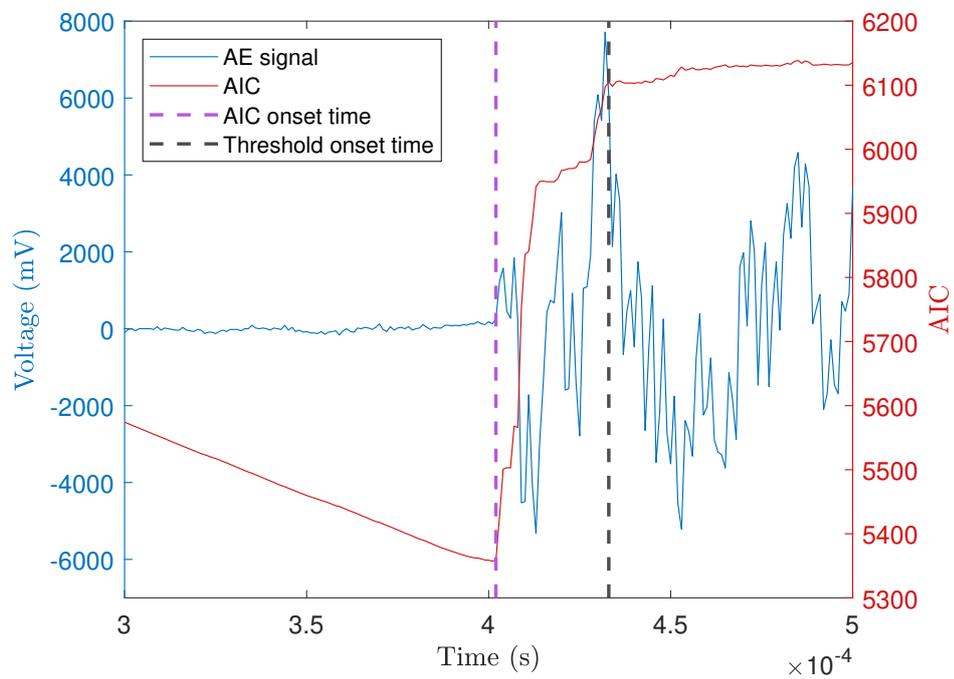
$$AIC(t) = t \log_{10}(\mathbb{V}(x[1 : t])) + (T - t - 1) \log_{10}(\mathbb{V}(x[t : T])). \quad (2.12)$$

Calculating equation (2.12) at every time step,  $t$ , the entropy of the signal is calculated before and after each time step. At the point of the onset of the event, the signal up to  $t$  will be composed of uncorrelated noise, yielding a high entropy. Conversely, the portion of the signal after  $t$  will be highly correlated and therefore of low entropy. By calculating  $AIC(t)$  across the sample time window  $t = 0 : T$ , a minimum will be returned at the onset of the acoustic burst.

Figure 2.6 demonstrates the use of the AIC picker to identify the onset of an AE burst, as well as that obtained by a fixed threshold set at the 99.9th percentile of the full signal duration. Comparing both methods of onset identification, it can be seen that the hard threshold detects the event later than the AIC, and misses the initial arrival of the low amplitude part of the signal. Although there does exist other



(a) Onset detection for an AE burst.



(b) Zoomed in on onset time identified by both a hard threshold and AIC.

Figure 2.6: A comparison between using a 99.9th percentile threshold and the AIC picker for identifying the onset of an acoustic emission burst.

onset time identification methods such as the use of a floating threshold [64] and the combination of AIC with some form of filtering [65, 66], onset time identification is not directly the topic of this thesis. It is, however, an important problem that will demand more attention as localisation methods continue to develop.

## 2.4 Traditional approaches for source localisation

For isotropic and homogenous materials, a very common approach for localisation is that of triangulation, where the common intersection point of  $\Delta T$  contours from multiple sensor pairs provides a unique source location. It should be stated that sometimes these triangulation approaches are actually examples of trilateration, where from some known distance (in this case, a distance between sensors), the angle of the source (where exactly on the  $\Delta T$  contour the source originated from) is obtained. Triangulation is the opposite, where a known angle is used to obtain an unknown distance. However, the AE literature generally uses the term triangulation to imply both techniques, and so that convention will be continued here. The method of using hyperbolae of  $\Delta T$  times that will be introduced later in this chapter is an example of triangulation, where known angles can be used to identify a distance that defines a hyperbola of constant  $\Delta T$  intersecting the source location.

One simple localisation technique that uses  $\Delta T$  times is that discussed in [67], and is visualised through the diagram in Figure 2.7. Having observed an acoustic emission event, the method first requires that the sensor that receives the signal first is noted. In the case of Figure 2.7, this is sensor 1. It can then be said that the source must travel an additional distance of  $d_{1,2}$  and  $d_{1,3}$  to reach sensors 2 and 3. Where the wave speed is known, these two distances can be calculated from the recorded  $\Delta T_{1,2}$  and  $\Delta T_{1,3}$ , which are then used to draw two constant radii circles around sensors 2 and 3 respectively. Following this step, both circles may be extended by a radial distance of  $h_1$ , which is also drawn from the centre of sensor 1. By extending  $h_1$ , a common intersection point between the three circles centred on each sensor can be found, which is then the location of the source.

Although simple and quite intuitive, this method is somewhat laborious, requiring single step iterative updates to identify the location of the source. It is also a very manual process, requiring the user to arbitrarily increase the size of  $h_1$  until some common intersection point emerges. Where step sizes in the increase in  $h_1$  are

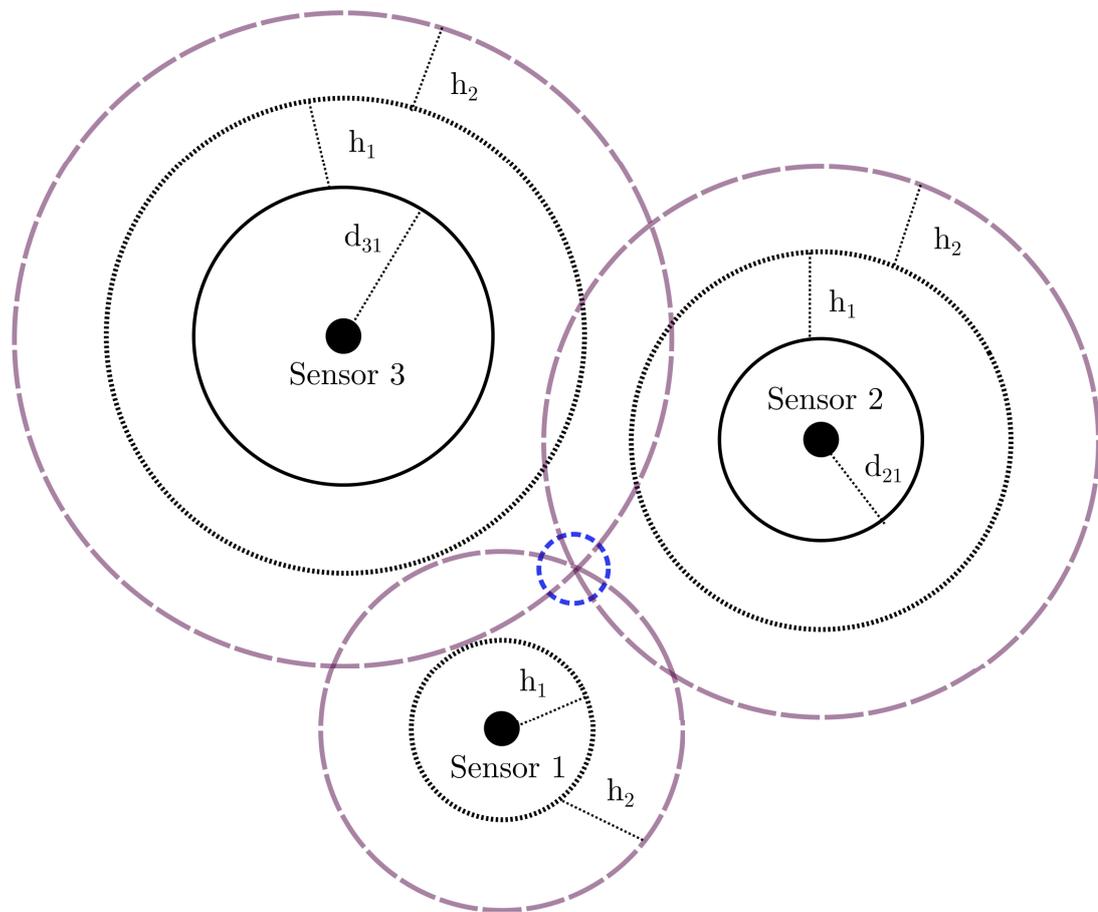


Figure 2.7: Basic triangulation technique.

too large, significant errors may be introduced, where one may overstep the true intersection point. A more general, and arguably more principled approach can be retrieved from the work of Tobias [68], where  $\Delta T$  values are used to define hyperbolae of potential source locations for each sensor pairing. A unique source location may then be found through consideration of only the  $\Delta T$  times, removing the need to define an incremental distance to progress from one of the sensors, and, thus, is more convenient.

For two sensors, a hyperbola of constant  $\Delta T$  may be defined analytically by first considering that,

$$\Delta T_{(i,j)}c = d_j - d_i, \quad (2.13)$$

and,

$$h = d_i \sin \theta, \quad (2.14)$$

where  $h$  is the opposite of the angle between sensor  $i$  and the source location, as illustrated in Figure 2.8, with  $d$  equivalent to  $r$ . By applying the trigonometric law of cosines across the full triangle created between the two sensors and source, the following expression is returned,

$$d_i^2 = d_j^2 - D^2 + 2d_i D \cos \theta, \quad (2.15)$$

which can be partly factorised into,

$$d_i^2 = d_j^2 - (D - d_i \cos \theta)^2 + d_i^2 \cos^2 \theta. \quad (2.16)$$

Given that  $h^2 = d_i^2 \sin^2 \theta$ , then,

$$d_i^2(1 - d_i^2 \cos^2 \theta) = d_i^2 \sin^2 \theta = h^2, \quad (2.17)$$

and, therefore,

$$h^2 = d_j^2 - (D - d_i \cos \theta)^2, \quad (2.18)$$

which by expanding out and, again, substituting  $\sin^2 \theta + \cos^2 \theta = 1$ , returns

$$d_i^2 \sin^2 \theta = d_j^2 - (D - d_i \cos \theta)^2. \quad (2.19)$$

Finally, from  $d_j = d_i + \Delta T_{(i,j)}c$ , a final expression for  $d_i$  can be stated as,

$$d_i = \frac{1}{2} \frac{D^2 - \Delta T_{(i,j)}^2 c^2}{\Delta T_{(i,j)} c + D \cos \theta}, \quad (2.20)$$

allowing for an analytical expression of the radii for constant  $\Delta T$  times. For localisation in  $D$  dimensions,  $D + 1$  unique hyperbolae are generally needed to satisfy uniqueness, requiring  $D + 1$  unique sensor pairs for  $D \leq 3$ . The above process can then be repeated for each pairing, where consideration of  $D + 1$  hyperbolae removes the unknown angle from equation (2.20), and allows identification of a common intersection point.

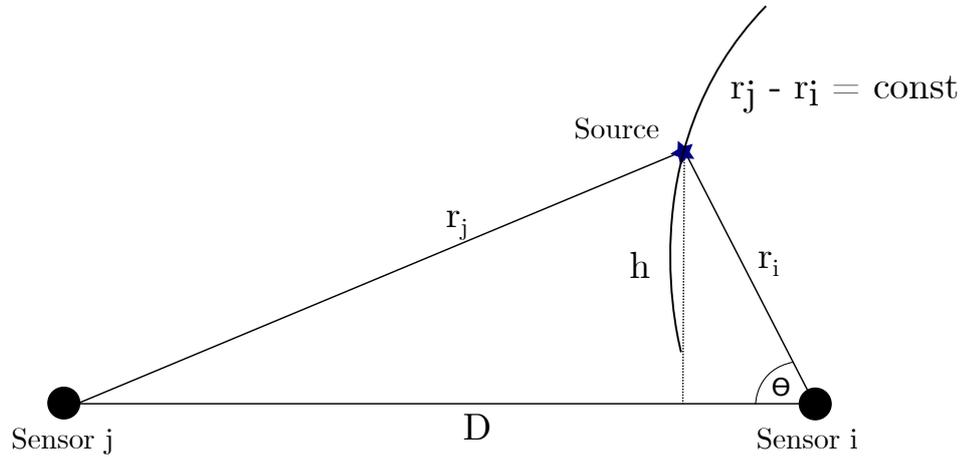


Figure 2.8: Hyperbola of constant  $\Delta T$  for two arbitrary sensors.

For three sensors, the basic triangulation technique introduced earlier in the chapter can now be reformulated from the perspective of  $\Delta T$  hyperbolae. The radii of the lines of constant  $\Delta T$ s can be calculated for each of the three sensor pairings, and resulting in the overlap of each of these hyperbolae, as shown in Figure 2.9. The more sensors that are added, the more hyperbolae that one can call upon, improving the overall accuracy of the localisation.

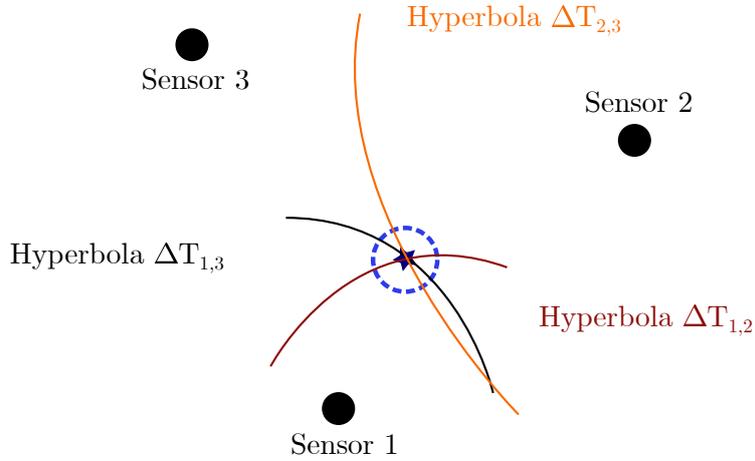


Figure 2.9: Hyperbola of constant  $\Delta T$  for two arbitrary sensors.

It should be noted that this sensor number requirement generally exists for methods that consider the arrival time of only the fastest travelling wave mode. For approaches that exploit the dispersive nature of AE signals, extracting the differences in arrival time of multiple wave modes allows events to be localised with one sensor, and fall into the category of single-sensor AE localisation, with examples including [69–72]. These types of methods require Lamb waves to form, and so are generally only relevant for plate-like structures where the propagation distance is sufficient such that the effects of dispersion emerge. Although single sensor methods reduce the complexity of the acquisition system, one is required to have much more prior knowledge regarding the propagation behaviour of the emissions, whilst also succumbing to similar limitations encountered in the single mode approach. In this thesis, methods that require only information regarding the arrival time of the first arriving wave mode will be considered, and so modal-based localisation methods fall out of the scope of this work. However, progress made in unimodal localisation will likely be transferable to the multimodal case.

## 2.5 Advances in AE localisation

Another way in which source localisation may be posed is as an optimisation problem, where the origin of an event can be estimated by minimising the difference between each recorded  $\Delta T$  and a calculated value originating from a trial source position,  $\mathbf{x}$ . Transforming the task of localisation to one of an optimisation then opens up the fast evolving field of numerical optimisation to handle such problems. Mathematically,

approaching localisation from this perspective equates to,

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \{Z\} \quad (2.21)$$

where,

$$Z = \sum_i (\Delta T_{i,observed} - \Delta T_{i,predicted}) \quad (2.22)$$

with  $i$  expressing the sensor pair index. Given that  $\Delta T_{i,observed}$  can be obtained experimentally and  $\Delta T_{i,predicted}$  can be expressed as,

$$\Delta T_{i,predicted} = \frac{1}{C} \|\mathbf{x} - \mathbf{S}_{(i,1)}\| - \|\mathbf{x} - \mathbf{S}_{(i,2)}\| \quad (2.23)$$

where  $\|\mathbf{x} - \mathbf{S}_{(i,\cdot)}\|$  represents the Euclidean distance between the trial location  $\mathbf{x}$  and the location of the  $i$ th sensor pairing  $\mathbf{S}_{(i,\cdot)}$ . Using a numerical solver of choice to minimise equation (2.22) then returns a source origin prediction. Although trivial to calculate the relevant gradients of equation (2.23), a gradient-free heuristic optimiser may be used so that equation (2.23) may be directly minimised [73].

Defined in the literature as the TOA method, this approach is still largely favoured by industry and is the standard localisation procedure detailed in the NDT handbook [74]. Despite being simple to implement, the TOA approach relies on two assumptions: firstly, that there exists a direct path between the source and sensors, and secondly, that the wave velocity profile is uniform in all directions. For anisotropic media such as composite materials, or components that contain geometrical features such as thickness changes, bolts and holes, these assumptions become largely invalid. For example, composites of different materials will induce angular dependent wave speeds, whilst geometric irregularities will result in phenomena such as scattering and wave mode conversion. It is also very likely that paths of direct propagation will be obstructed. Given that the majority of engineering components will contain at least one of these inhomogeneities, there has been a significant amount of literature dedicated to developing AE localisation strategies that are not constrained to these assumptions.

A number of initial studies focused on developing methods that are robust to aniso-

tropic media, incorporating an angular-dependent velocity term into the minimisation process [75–77]. However, knowledge of the wave-velocity profile will not always be available *a-priori*. To account for this scenario, approaches that allow a directionally-dependent velocity term to be solved in parallel with the event location have also been suggested [67, 78]. In relation to the localisation of AE in structures where a direct propagation path is obstructed, one strategy considered by a number of authors is to define a correlation operator that can quantify the similarity of two AE events through the concept of time-reversal [79, 80]. The correlation between a real AE event and a set of artificial events can then be assessed, where the origin of the artificial excitation that returns the highest similarity is determined to be the location of the real source. Ciampa and Meo [81] implemented a similar strategy for localisation in complex geometries, and by ensuring a diffused wavefield, were able to localise an acoustic impact with only a single sensor. In general, single sensor localisation approaches leverage the multi-modal nature of an AE burst to extract multiple arrival times for a single waveform, with other examples including far field separation of Lamb wave modes [82], a finite-element approach to back-propagate waveforms [70], as well as the use of reverberation and reflection patterns [83].

## 2.6 Towards data-driven localisation

Although many of the previously referenced works have been successfully applied to anisotropic materials, there still persists a challenge in using many of these methods on structures where direct propagation between source and sensors may be blocked. Regularly termed complex structures in the AE literature, these structures contain geometrical features that obstruct the propagation path of an emitted waveform. It can therefore be difficult to analytically express how a source travels to a receiving position. In an attempt to address this limitation, a number of data-driven methods have been developed, where instead of attempting to directly characterise the behaviour of the AE waves through physical laws, measurements that capture the phenomena of interest are first obtained and then used to learn a statistical model. The use of such a technique in the context of AE source location was first suggested as part of the Delta  $T$  mapping approach developed by Baxter et al. [30]. The method views AE localisation as a problem of spatial mapping, whereby extracting the onset time of a series of artificial events at a number of distributed sensors, a map that quantifies the expected  $\Delta T$  information across a structure of interest is

constructed for each available sensor pair. As a real AE event occurs, the recorded  $\Delta T$  values are then matched to those on the maps, allowing the identification of an estimated source location. The approach has since been applied in a number of challenging environments, such as an aircraft landing gear component [84], as well as composite materials [85]. An extension to the original methodology was also offered by Al-Jumaili et al. [86], reducing the required level of input from an operator. This reduction was achieved by the use of a hierarchical clustering algorithm to automatically discard erroneous training events, as well as the implementation of a minimisation scheme so that a source location estimate could be obtained in an autonomous manner.

To offer an increased spatial resolution, both the initial and extended Delta  $T$  methodologies linearly interpolate between the training grid points to generate the  $\Delta T$  maps. However, for complex structures, the presence of phenomena such as wave mode conversion and internal reflections will result in a highly nonlinear wavefield. As a consequence of this, a linear interpolation scheme will not be sufficient to capture the complexity of the waveform behaviour between grid points, particularly as the density of the training grid is reduced. Hensman et al. [31] recognised this, and instead proposed the use of Gaussian process (GP) regression, a Bayesian non-parametric technique for nonlinear regression. There, rather than using a minimisation scheme, source localisation is achieved by directly learning an inverse mapping from  $\Delta T$  to source location. In addition to providing an improved interpolation strategy, a probabilistic output is also returned, where each prediction is associated with a degree of uncertainty - which, as discussed already in this thesis, is an attractive property of Bayesian solutions. Although directly modelling the inverse problem returned promising results, the approach of [31] is not reflective of the general approach to regression problems where the inputs to the model are set as the independent variables. Constructing the regression problem in this manner also prevents proper treatment of the uncertainty on  $\Delta T$  values.

A number of other authors have since adopted various Bayesian approaches as part of AE localisation strategies. To localise AE in a concrete column, Schumacher et al. [87] formulated a probabilistic time of arrival model, placing prior distributions over the propagation parameters of the model. Combining the priors with observed data, posterior distributions over the model parameters can be inferred. To locate new AE events, a predictive model is constructed from the joint distribution between the posterior of the parameters and the likelihood of the parameters given a new event,

with inference in both the parameter estimation and predictive step computed using a Markov chain Monte Carlo sampling scheme. Similarly, [88] consider a probabilistic interpretation of AE source localisation, but treat the problem dynamically, transforming the time of flight equations into a set of first order differential equations that describe the spatial evolution of a propagating AE signal. Such a formulation naturally lends itself to a state space representation, where the hidden states to be identified are the positional coordinates of the AE event, as well as the wave propagation speed. The system is solved recursively using an extended Kalman filter, with the methodology demonstrated by localising AE in an aluminium plate. It should be noted that by constructing the filtering problem, the resulting state distributions that are being recovered grow in dimensions, unlike in the approach of [87], where distributions over propagation parameters are treated statically. Sen et al. [89] also approach the localisation task through a state space formulation, considering the use of an ensemble Kalman filter (EnKF) as well as the particle filter. In comparison to standard Kalman filtering, the EnKF replaces full, closed form updates of the states with an ensemble of draws from the full state distribution and offers considerable speed up at inference time. The particle filter is similar to the EnKF in that it is also a type of sequential Monte Carlo (SMC), except where sampling from an ensemble is replaced with a set of particles that are weighted according to the likelihood that they represent the sampled state probability distribution. The advantage offered by the particle filter over the EnKF is that state distributions are not restricted to being Gaussian, allowing for more accurate AE localisation predictions of the two SMC approaches, albeit with an increased computational cost [89].

Each of these probabilistic approaches enable uncertainties arising in the parameters that govern the propagation behaviour of the AE waves to be captured, thus allowing variability that may exist in material properties, time of arrival estimation error, measurement noise, and other sources of uncertainty to be incorporated into the model. However, in all of these works, the limiting assumptions that the structure is both isotropic and homogenous are made, rendering the approaches invalid for many propagation media encountered in engineering such as composites and those that contain geometrical features. MCMC and particle filtering schemes are also notorious for large computational overheads, and can often take a long time to converge to the posterior solution.

Although another example of probabilistic localisation in an isotropic homogenous material, Madarshahian et al. [90] adopted an interesting interpretation of the onset

time challenge, performing Bayesian model selection to determine the most likely onset time from two candidate models. The most probable onset time is subsequently fed into the localisation scheme, improving performance over the case where only a single identification method is used.

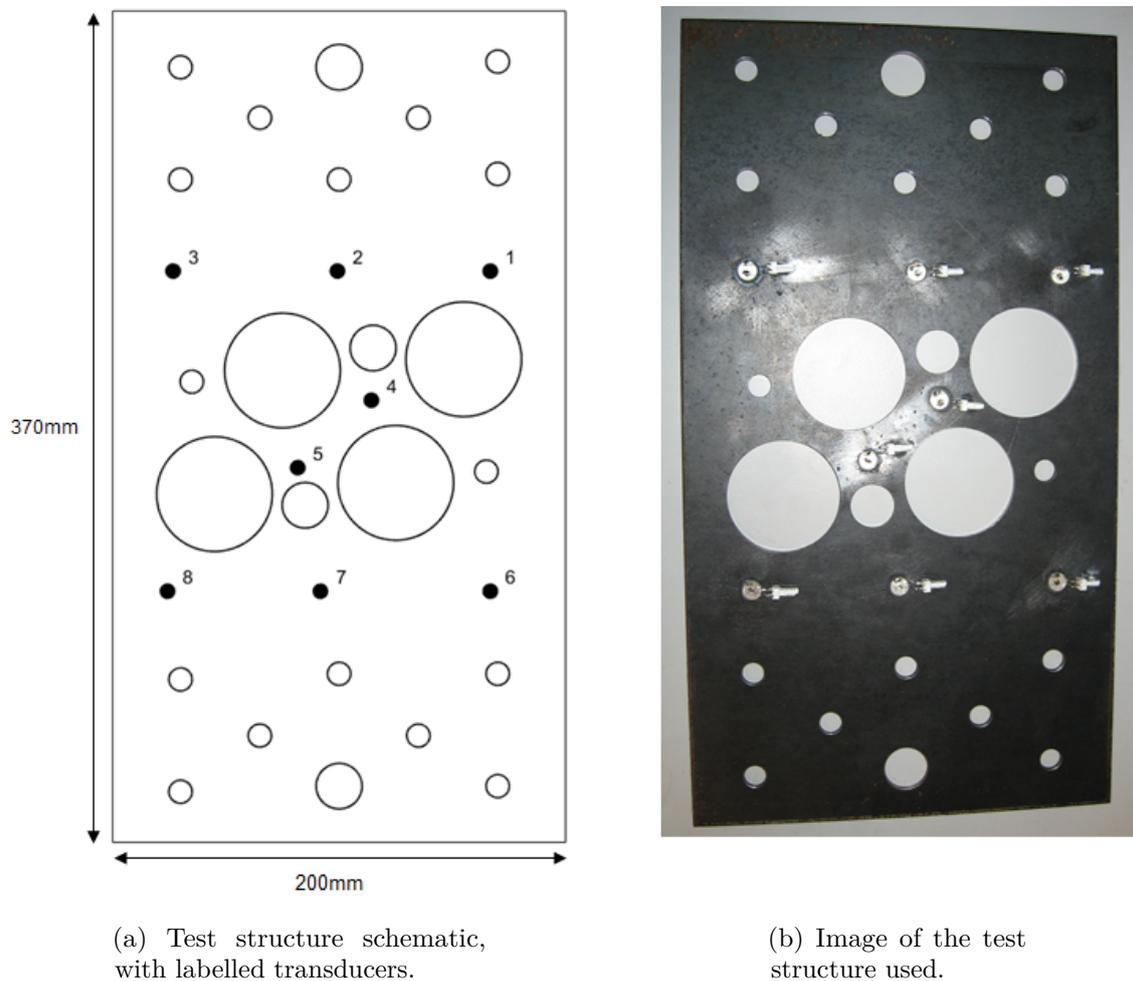
In an attempt to remove the isotropic assumption, Niri et al. [91] extended their original work to allow for anisotropic media with the use of an unscented Kalman filter for localisation in a carbon-fibre reinforced plate. However, this work requires either that the velocity profile is already known, or that sensors be placed in similarly positioned clusters of no less than two sensors, increasing the minimum sensor requirement for localisation in two dimensions from three to six (although the work of [67] suggests that this should be nine sensors in clusters of three). The method also still requires a direct propagation path between sources and receivers, and so does not extend to complex structures.

Whilst these works are important contributions, it is clear that there still remains the significant challenge of locating AE sources in complex structures from a probabilistic perspective, warranting the need for a new localisation framework.

## 2.7 Experimental case study

For the work that has been completed in Chapters 3 - 5, an experimental data set obtained during the work of Hensman et al. [31] is used to demonstrate various aspects of some of the methods developed in this thesis. This subsection, therefore, provides the motivation for the use of this data, as well as a full description of the experimental set-up, data acquisition and data processing that was performed. Although a complete overview is given, the reader may still wish to refer to Section 3 of the original paper should further detail be needed [31].

The interest in the generation of the data set was to collect a series of acoustic emission measurements across a structure that is representative of the challenge of localising acoustic emission in realistic engineering components and structures. To this end, a 200 x 370 x 3mm mild steel sheet with a series of holes cut through the full volume of the plate was manufactured. Figure 2.10 provides a schematic and image of the plate. The presence of the holes that are located across the entirety of the plate induce a number of complex phenomena into the propagation of the waves, including



(a) Test structure schematic, with labelled transducers.

(b) Image of the test structure used.

Figure 2.10: Complex structure used to demonstrate the methodology. Recreated from [31].

scattering and wave mode conversion, creating a highly challenging wavefield to capture and accurately model. For large areas of the plate, a direct propagation path to many of the sensor pairs is also blocked, adding further complexity.

To collect acoustic emission measurements from the structure, a grid with a spacing of 5mm (excluding where holes are located) was first constructed across the surface of the plate, where at each nodal location, a high-power laser pulse was used to artificially generate an acoustic emission source through thermoelastic expansion. Eight Sonox P5 piezoceramic transducers were mounted to the specimen using cyanoacrylate and used to capture the ultrasonic response to each excitation, recorded with an eight-channel Physical Acoustics DISP AE acquisition unit at a sampling frequency of 10MHz. A filter with a bandwidth of 100Hz - 3MHz was also applied,

which is the frequency range one is generally interested in when capturing acoustic emissions. The locations of these sensors are indicated by the black circles on Figure 2.10(a). To extract time of arrival values, the AIC onset picking technique as detailed above was implemented. At every node, a  $\Delta T$  is then calculated for each sensor pairing, creating a 28 dimensional  $\Delta T$  vector, arising from the existence of 28 unique pairing combinations. Given that there are 2278 excitation sites, the final data set corresponded to a 2278 x 28 matrix (2 dimensional array with 2278 rows and 28 columns).

## 2.8 Overview of the thesis

In the following section, a brief summary of the contribution made by each chapter is provided.

### Chapter 3

*Gaussian processes for probabilistic AE localisation:* A probabilistic localisation framework suitable for complex structures is developed using Gaussian process regression. Details of the proposed probabilistic localisation framework are given, as well as the relevant Gaussian process theory. A case study consisting of a complex structure is then investigated, with the proposed method offering a number of gains when compared with leading models available in the current literature, including the quantification of uncertainty, proper treatment of measurement noise, and an improved overall predictive accuracy, particularly as the number of training measurements used in the learning process is reduced.

### Chapter 4

*Heteroscedastic Gaussian processes for enhanced uncertainty quantification:* A forward model of  $\Delta T$  that more appropriately models the associated uncertainty with the onset times is developed. By transitioning to a heteroscedastic noise model, it is possible to capture a locally varying noise process, where the contribution of an individual sensor pair to an overall prediction is effectively weighted according to the quality of the sensor pair coverage in a given location. The corresponding emission

likelihood surfaces are generally tighter about the true event location, leading to more accurate maximum likelihood location estimates.

**Chapter 5**

*Physical constraints for physics-informed onset time mapping:* To reduce the amount of training data required to learn the localisation model, this chapter considers a physics-informed machine learning approach for  $\Delta T$  mapping. Physical insight is incorporated into the  $\Delta T$  learning process through constraining the Gaussian process kernel function. These constraints ensure that any draws made from the GP prior abide by the associated boundary conditions, improving predictive performance where training data is limited.

**Chapter 6**

*AE Localisation for wind turbine bearings:* The localisation framework is extended to non-Euclidean space, allowing its use in spherical geometry such as bearings. Localisation on a scaled-down roller element bearing is investigated, with the method demonstrating good localisation performance in the presence of highly variable (noisy)  $\Delta T$  features.

**Chapter 7**

*Conclusions:* Concluding remarks and discussion on future direction that the work may take.

# GAUSSIAN PROCESSES FOR PROBABILISTIC AE LOCALISATION

### Highlights:

- *An introduction to Gaussian process regression from the perspective of acoustic emission localisation is presented.*
- *A novel framework for probabilistically localising acoustic emission events in complex structures is proposed.*
- *The approach enables a likelihood of emission location to be quantified across the structure of interest.*
- *It is shown that the proposed method outperforms state-of-the-art methods from the literature over a number of testing scenarios.*

## 3.1 An introduction to Gaussian process regression

At the core of much of the work developed in this thesis is the use of *Gaussian processes* (GPs), which provide a Bayesian machine learning approach for solving regression problems [92]. They are largely characterised by their ability to quantify uncertainty on predictions and capture nonlinear relationships in the presence of

noisy data. In a physical world where there are many sources of uncertainty, this makes GPs particularly relevant for application within engineering and SHM, of which some other examples can be found in [93–95].

There are many ways in which GP regression can be motivated. Here, we will begin by introducing Bayesian linear regression, the simplest regression model available in a Bayesian framework, before discussing how this model may be generalised into a more flexible and expressive regressor; Gaussian process regression.

To begin, let us discuss the basis of any regression problem, in which we are concerned with mapping a series of  $D$ -dimensional observations  $X = (\mathbf{x}_i)_{i=1}^N$  ( $N \times D$  matrix) to a corresponding vector of targets  $\mathbf{y} = (y_i)_{i=1}^N$ . In engineering, lots of tasks often boil down to problems of this nature. Perhaps we want to build a surrogate model of some expensive experiment (either real or numerical) [96, 97], or maybe we want to learn the relationship between wind speed and power output of a wind turbine [93, 98] such that electrical energy generation at a given time may be predicted. Likewise, in SHM, these types of mappings consistently arise when developing damage detection strategies. For example, predicting the loading or response of a structure from some observable inputs [94, 99, 100].

In a Bayesian linear regression model, the assumption is made that the targets are corrupted by an additive noise,  $\epsilon$ . To further simplify, the noise is assumed to follow the form of a zero mean Gaussian white noise  $\epsilon$ , with variance  $\sigma_n^2$ . Mathematically, this equates to modelling functions of the form,

$$y = f(\mathbf{x}) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma_n^2), \quad (3.1)$$

where,

$$f(\mathbf{x}) = \mathbf{x}^\top \mathbf{w}, \quad (3.2)$$

with  $\mathbf{w}$  representing a vector of weights corresponding to the parameters of the model. Following a Normally distributed noise that is i.i.d (independent and identically distributed), an expression for the likelihood of the output can be derived in closed form,

$$p(\mathbf{y}|X, \mathbf{w}) = \prod_{i=1}^N p(y_i|\mathbf{x}_i, \mathbf{w}) = \mathcal{N}(X^T \mathbf{w}, \sigma_n^2 \mathbb{I}), \quad (3.3)$$

which is simply the product of the individual targets likelihoods. To construct a Bayesian formulation first requires that a prior over the weights is specified. As discussed in Chapter 1.3, by ensuring that the prior is from the family of distributions conjugate to the chosen likelihood, further inference can be computed in a closed form manner [101]. Given that the Normal distribution is conjugate to itself (which is the assumed form of the likelihood), a Gaussian prior can be selected, parameterised with a zero mean and covariance matrix  $\Sigma_w$ ,

$$\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \Sigma_w). \quad (3.4)$$

Through the use of Bayes' rule, the posterior over the weights of the model may then be expressed as,

$$p(\mathbf{w}|\mathbf{y}, X) = \frac{p(\mathbf{y}|X, \mathbf{w})p(\mathbf{w})}{p(\mathbf{y}|X)}, \quad (3.5)$$

where the marginal likelihood,  $p(\mathbf{y}|X)$ , can equivalently be written as the integral of the product of the likelihood and prior with respect to the weights,

$$p(\mathbf{y}|X) = \int p(\mathbf{y}|X, \mathbf{w})p(\mathbf{w})d\mathbf{w}. \quad (3.6)$$

Following the assignment of a Normal prior and likelihood, the posterior distribution too is Gaussian, of whose mean and covariance is given by,

$$p(\mathbf{w}|X, \mathbf{y}) \sim \mathcal{N}(\bar{\mathbf{w}}, \Sigma_{w*}), \quad (3.7)$$

$$\bar{\mathbf{w}} = \sigma_n^{-2}(\sigma_n^{-2}XX^T + \Sigma_w^{-1})^{-1}X\mathbf{y}, \quad (3.8)$$

$$\Sigma_{w*} = (\sigma_n^{-2}XX^T + \Sigma_w^{-1})^{-1}. \quad (3.9)$$

For more details, see [92], page 9, or more generally, [102]. In general, the interest is not in the values of  $\mathbf{w}$  themselves, but those of the target values,  $y_*$ , that arise from

previously unseen observations with inputs  $\mathbf{x}_*$ . For instance, returning to our earlier example of wind turbine power curve modelling, a common task is to make power output predictions at wind speeds not yet observed, given previous wind speed and wind turbine power output measurements. To make predictions of this nature with the linear model that has been derived, it is possible to average over all of the possible weights, scaled by their relative posterior probability. Note the difference between a non-Bayesian approach, where a fixed parameter value would be picked. To perform this averaging across all possible parameter values, it is possible to marginalise over  $\mathbf{w}$ , and calculate the posterior predictive distribution  $p(y_*|\mathbf{x}_*, X, \mathbf{y})$ ,

$$p(y_*|\mathbf{x}_*, X, \mathbf{y}) = \int p(y_*|\mathbf{x}_*, \mathbf{w})p(\mathbf{w}|X, \mathbf{y})d\mathbf{w}, \quad (3.10)$$

returning the following Gaussian distribution,

$$p(y_*|\mathbf{x}_*, X, \mathbf{y}) = \mathcal{N}(\sigma_n^{-2}\mathbf{x}_*^T\zeta^{-1}X\mathbf{y}, \mathbf{x}_*^T\zeta^{-1}\mathbf{x}_*), \quad (3.11)$$

where  $\zeta = \sigma_n^{-2}XX^T + \Sigma_w^{-1}$ , which is also equal to the inverse of the posterior variance of the weights. By marginalising over all possible weights, the predictive posterior is only conditioned on previously seen data, and not the model parameters themselves.

At present, the regression model is limited in the sense that it only has the capacity to represent linear relationships. To improve the expressivity of the model, it is possible to project the inputs into a higher dimensional space using a basis function expansion, replacing  $\mathbf{x}$  with some transformation  $\Omega(\mathbf{x})$ . For example, polynomial regression of order  $n$  may be implemented by setting  $\Omega(\mathbf{x}) = [1; \mathbf{x}; \mathbf{x}^2; \dots; \mathbf{x}^n]$ . To represent increasingly complex functions,  $n$  will be required to become larger, which to directly compute, will become exceedingly computationally expensive. Fortunately, one does not have to deal with this direct computation, with a call to the projected input vector being equivalently represented by a kernel function,  $k(\cdot, \cdot)$ . Transitioning to this kernel representation firstly relies on the projected inputs being computed as an inner product. Through use of the *kernel trick* [92, 103], all inner products can subsequently be replaced with a kernel function, where in certain cases, the kernelised representation is equivalent to an infinite basis set, allowing an infinitely complex model to be computed in finite time.

From a kernel perspective, a transition is made away from placing a distribution over

the weights and towards a distributions over possible functions. Therefore instead of capturing uncertainty in the parameters of the model, uncertainty in the functions, and, therefore, their evaluations at points of interest, are handled directly. In this sense, the Gaussian process can be viewed as a distribution over possible functions that represent our process of interest, where each realisation is a potential function generated by that GP. The GP therefore acts as a prior over  $f(\mathbf{x})$  in equation (3.1), and can be defined by,

$$f \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \quad (3.12)$$

It can be seen that the GP is fully specified by its mean function  $m(\cdot)$ , and covariance (kernel) function  $k(\cdot, \cdot)$ . Throughout this work, a mean function of zero will be used, and is therefore omitted from this point onwards. Although, given the mean is simply an offset of the predictions, it is trivial to include in the predictive equations, either as a constant or a parametric function of the relevant inputs. The covariance function then introduces a measure of similarity between data points depending upon how close they are in the input space. For instance, we expect two inputs  $\mathbf{x}_1$  and  $\mathbf{x}_2$  that are near each other to return similar function evaluations. Loosely, the covariance function can be thought of as governing how “closeness” in the input space relates to the target space.

Following from the general motivation provided earlier in the chapter for constructing regression problems, given some initial training data, the Gaussian process can be used to make predictions at a set of unseen testing points,  $X_*$ , by constructing a joint distribution between the training targets,  $\mathbf{y}$ , and the unobserved latent functions,  $\mathbf{f}_*$  as,

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{f}_* \end{bmatrix} = \mathcal{N} \left( \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} K(X, X) + \sigma_n^2 \mathbb{I} & K(X, X_*) \\ K(X_*, X) & K(X_*, X_*) \end{bmatrix} \right). \quad (3.13)$$

Through the standard machinery for conditioning a multivariate Gaussian distribution, the conditional distribution over  $\mathbf{f}_*$  can then be obtained [92] as,

$$p(\mathbf{f}_* | X, \mathbf{y}, X_*) = \mathcal{N}(\mathbb{E}[\mathbf{f}_*], \mathbb{V}[\mathbf{f}_*]), \quad (3.14)$$

where the posterior mean and variance are,

$$\mathbb{E}[\mathbf{f}_*] = K_{X_*X}(K_{XX} + \sigma_n^2\mathbb{I})^{-1}\mathbf{y} \quad (3.15)$$

$$\mathbb{V}[\mathbf{f}_*] = K_{X_*X_*} - K_{X_*X}(K_{XX} + \sigma_n^2\mathbb{I})^{-1}K_{XX_*}. \quad (3.16)$$

where  $K(X, X) = K_{XX}$ , which equals the covariance matrix between  $X, X$ .

In this work, the author wrote custom scripts for Gaussian process regression in MATLAB, based on Williams and Rasmussen [92], with the code for the predictive equations summarised in algorithm 1.

---

**Algorithm 1** Predictions with Gaussian process regression

---

**Input:**  $X_*$  (testing inputs),  $X$  (training inputs),  $\mathbf{y}$  (targets),  $\{k, \theta\}$  (kernel function),  $\sigma_n^2$  (noise parameter),

**Output:**  $\mathbf{y}_*$  (predictive mean),  $\mathbb{V}[\mathbf{y}_*]$  (predictive variance),  $\log p(\mathbf{y}|X)$  (log marginal likelihood).

$$\begin{aligned} L &= \text{chol}(K_{XX} + \sigma_n^2\mathbb{I}), && \triangleright \text{lower triangular cholesky} \\ \alpha &= L^\top \setminus (L \setminus \mathbf{y}), \\ \mathbb{E}[\mathbf{y}_*] &= (K_{X_*X})^\top \alpha, \\ \mathbf{v} &= L \setminus K_{X_*X_*} \\ \mathbb{V}[\mathbf{y}_*] &= K_{X_*X_*} - \mathbf{v}^\top \mathbf{v} \\ \log p(\mathbf{y}|X) &= -\frac{1}{2}\mathbf{y}^\top \alpha - \Sigma \log L - \frac{n}{2} \log 2\pi \end{aligned}$$


---

The  $A \setminus B$  operator refers to the solution of  $Ax = B$  for  $x$ , and is directly available in MATLAB, with equivalent implementations commonly available as part of linear algebra libraries in other languages.

### 3.1.1 Kernel selection

In a Gaussian process, the kernel provides a measure of similarity between inputs. The role of the kernel can be likened to governing what functions are likely to arise from draws of the prior, determining the family of functions that may be modelled. The definition of the kernel, therefore, allows us to encode any prior beliefs that may be held about the type of functions wished to be captured by the model. In an engineering context, this functionality is particularly useful, as it is common from physical insight that we often have some knowledge about how the phenomena we

are interested in modelling behaves. For example, will my data exhibit periodic behaviour? Or, is it desirable to smooth over low amplitude perturbations about a mean? Will there be a global linear trend within the data? These are all assumptions that can easily be encoded into the predictive outputs through selection of the kernel.

To be a valid covariance function requires the function to be positive semidefinite. The *gram matrix*<sup>1</sup> corresponding to the covariance function therefore has eigenvalues that are both real and non-negative. Many common kernels also exhibit stationarity, and are, therefore, functions of pairwise distance between inputs, e.g.  $\mathbf{x} - \mathbf{x}'$  rather than  $\mathbf{x}$ . Often this distance is defined as a Euclidean distance. However, there exists other spaces on which Gaussian processes and their respective kernels may be defined on, such as Riemannian manifolds [104] (a generalisation of Euclidean space), and will be discussed in more detail from a health monitoring perspective in Chapter 6.

Perhaps the most common covariance function is that of the squared-exponential (SE), also sometimes referred to as the exponentiated quadratic due to its mathematical form. The kernel is given by,

$$k_{SE}(\mathbf{x}, \mathbf{x}') = \sigma_f \exp\left(\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2l^2}\right) \quad (3.17)$$

where  $\sigma_f$  is the signal variance that controls the scale of the function, with  $l$  the characteristic length scale, specifying how far away pairs of inputs may be such that they still influence one another. A slightly less formal, but perhaps more intuitive way to think about the length scale is that it determines the “wiggleness” of the function draws, where a larger length scale creates smoother functions, whilst shorter values return very sharp, rapidly varying function evaluations, as illustrated in Figure 3.1. These parameters of the covariance function are more generally defined as the hyperparameters of the model and will be discussed in more detail in the proceeding section.

Draws arising from a squared-exponential function are, in general, very smooth. This property arises from the infinite differentiability of the SE function, which Stein [105] regards as often too unrealistic when modelling physical processes.

Another form of covariance functions are those that are given by the Matérn class of

---

<sup>1</sup>The gram matrix corresponds to the the matrix  $K$  that results from kernel function evaluations between pairs of inputs at indices  $i, j$ . Therefore,  $K_{i,j} = k(\mathbf{x}_i, \mathbf{x}'_j)$

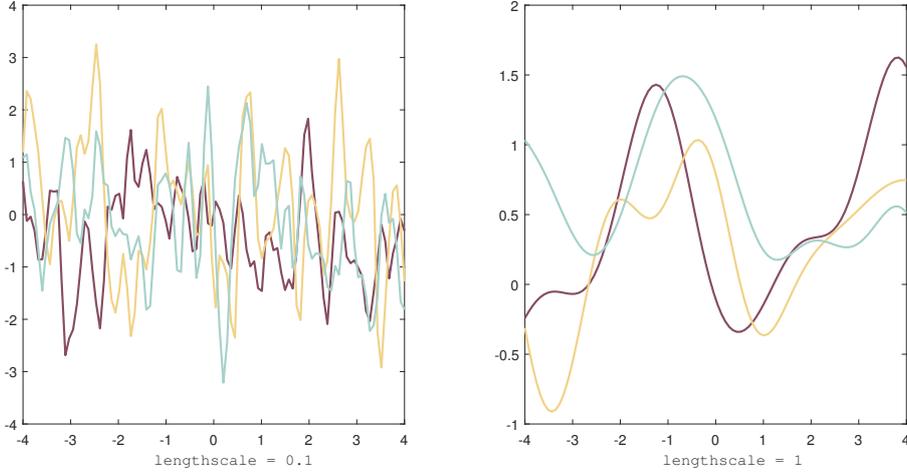


Figure 3.1: Draws from GP prior with different lengthscales. Squared exponential function used in both cases.

kernels, defined as,

$$k_{MA}(\mathbf{x}, \mathbf{x}') = \sigma_f \frac{2^{1-\nu}}{\Gamma(\nu)} \left( \frac{\sqrt{2\nu} \|\mathbf{x} - \mathbf{x}'\|}{l} \right)^\nu K_\nu \left( \frac{\sqrt{2\nu} \|\mathbf{x} - \mathbf{x}'\|}{l} \right) \quad (3.18)$$

where  $K_\nu$  is the modified Bessel function and  $\nu$  is a positive parameter that defines the order of differentiability of the function and, therefore, roughness of the subsequent draws. It is common for  $\nu$  to be set to  $p + 0.5$ , where  $p$  is a positive integer. For example, when  $\nu = 3/2$ , equation (3.18) simplifies to,

$$k_{\nu=3/2}(\mathbf{x}, \mathbf{x}') = \left( 1 + \frac{\sqrt{3} \|\mathbf{x} - \mathbf{x}'\|}{l} \right) \exp \left( -\frac{\sqrt{3} \|\mathbf{x} - \mathbf{x}'\|}{l} \right) \quad (3.19)$$

The kernels that are employed in this work will generally be of either of the two forms discussed so far, and so further discussion on kernel selection is out of the scope of this thesis. For more insight into the kernel selection process and other examples of covariance functions, the interested reader may consult Chapter 4 of [92] and the work of [106].

### 3.1.2 Learning hyperparameters

As discussed in the preceding section, with the specification of the covariance function comes the requirement to define a number of hyperparameters. The inclusion of the prefix ‘hyper’ indicates that these are parameters of the prior, rather than explicit parameters of the model. These hyperparameters will shape the exact behaviour of the kernel function draws, and can be tuned in order to maximise the likelihood of the observed data.

We are also generally interested in making predictions on noisy data. As such, it is necessary to include a noise term that represents the amplitude of noise we believe the true function evaluations have been corrupted with. As seen above, it is standard practice to treat this noise as a zero mean Gaussian with variance  $\sigma_n^2$  (although in Chapter 4 the non-constant case is considered), which can also be included into the vector of total hyperparameters,  $\boldsymbol{\theta}$ , that are to be determined.

Due to the Bayesian framework in which Gaussian process regressions resides, we have access to the marginal likelihood, providing a means of quantifying how likely the observed data are given all possible latent functions. This term can be expressed conveniently in log form as,

$$\log p(\mathbf{y}|X, \boldsymbol{\theta}) = -\frac{1}{2}\mathbf{y}^T(K_{XX} + \sigma_n^2\mathbb{I})^{-1}\mathbf{y} - \frac{1}{2}\log |(K_{XX} + \sigma_n^2\mathbb{I})| - \frac{N}{2}\log(2\pi). \quad (3.20)$$

The log marginal likelihood can be seen as self regularising, satisfying Occam’s Razor [107] in that one should seek simple solutions, where the term  $\frac{1}{2}\log |(K_{XX} + \sigma_n^2\mathbb{I})|$  penalises overly complex models;  $-\frac{1}{2}\mathbf{y}^T(K_{XX} + \sigma_n^2\mathbb{I})^{-1}\mathbf{y}$  is then a measure of how well the model fits the training data. By maximising equation (3.20) with respect to the covariance hyperparameters, an ‘optimal’ set of hyperparameters may be identified, where optimal is used in the sense that the chosen hyperparameters maximise the marginal likelihood. This problem therefore falls into the remit of the field of numerical optimisation, from which we can call upon a vast suite of tools to assist us [73, 108, 109]. Where access is available to the gradient of the marginal likelihood with respect to the hyperparameters, simple gradient-based optimisation may be used, such as stochastic gradient descent. Alternatively, heuristic, gradient-free techniques may also be employed, which do not require the gradients of equation (3.20) to

be specified, and instead, allow direct use of the marginal likelihood term. As per standard numerical optimisation conventions, this maximisation can be translated to a minimisation of the negative log marginal likelihood,

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} -\log(\mathbf{y}|X, \boldsymbol{\theta}) \quad (3.21)$$

As there is no free lunch [110] in search and optimisation problems, it is likely that many optimisation schemes will, on average, perform similarly across all such problems. The authors' preference here is an example of population-based optimisation - the Quantum-Behaved Particle Swarm (QPSO) [111]. This approach has the key advantage of being globally convergent [112], and has been demonstrated to perform well in engineering optimisation challenges, of which an example can be found in [113].

### 3.1.3 Performance metrics

In order to quantify a measure of error, the normalised mean squared error (nMSE) is employed regularly throughout the thesis:

$$nMSE = \frac{100}{N\sigma_y^2}(\mathbf{y}_* - \mathbf{y})^T(\mathbf{y}_* - \mathbf{y}), \quad (3.22)$$

where  $N$  is the number of points,  $\sigma_y^2$  is the variance of the true targets,  $\mathbf{y}_*$  are the model predictions, with  $\mathbf{y}$  the true targets. A score of 0 would be returned if the model predictions perfectly aligned with the true targets, whilst a score of 100 is identical to predicting at the mean. As many of the predictions that are made return a distribution, it is also possible to assess the negative probability of a prediction under the model, also referred to as the loss of the model. Standardising this measure with respect to the mean and variance of the training set, the mean standardised log loss (MSLL) can be expressed as,

$$MSLL = \frac{1}{N} \sum_k \{-\log p(\mathbf{y}_{*,k}|X, \mathbf{y}, \mathbf{x}_{*,k}) + \log p(\mathbf{y}_{*,k}; \mathbb{E}(\mathbf{y}_k), \mathbb{V}(\mathbf{y}_k))\}, \quad (3.23)$$

where  $k$  indexes a particular test point. The log loss can be interpreted as the

negative likelihood of the predictions relative to those made under the trivial model, i.e. the mean and variance of the training observations. As such, a larger negative MSSL reflects more favourable predictive distributions. Note that equation (3.23) is only valid where the assumed form of the data and subsequent predictions is Gaussian.

## 3.2 AE source localisation using Gaussian processes

In this section, a probabilistic approach for localising acoustic emission is presented that allows probable locations of damage within complex structures to be identified. The method, termed Likelihood of Emission Location (LoEL), firstly proposes the use of Gaussian process regression to learn the relationship between source location and  $\Delta T$  for a number of sensor pairs across the surface of a structure of interest. As an acoustic emission event with an unknown origin is observed, a mapping is then generated that quantifies the likelihood of the emission location over the structure. The approach allows surfaces of probable emission locations in complex structures to be identified, with uncertainty on  $\Delta T$  values able to be handled and incorporated into the localisation process in a statistically principled manner. To demonstrate and measure the success of the proposed approach, the data set as detailed in Chapter 2.7 is used. This framework is also the basis of the work contained in the proceeding chapters, and will therefore be referred to throughout the duration of the thesis.

### 3.2.1 Outline of localisation framework

In the proposed methodology, the first objective is to employ a Gaussian process to learn the relationship between locations on a test structure and the corresponding  $\Delta T$  values across a number of sensor pairs. For each unique sensor pair, the likelihood that a given  $\Delta T$  originated from locations across the structure is assessed, generating a contour of potential source locations. A marginal likelihood across all sensors pairs can then be calculated, returning a map of the structure that represents the total likelihood of source location given the model for each sensor pair. Regions of high likelihood are interpreted as being more likely to be the origin of the event, where a single-point prediction of the source location can be identified at the location with

the highest marginal likelihood. In addition to the standard benefits of taking a Bayesian approach, such as predictions that represent a distribution and the ability to naturally incorporate uncertainty, the method offers a number of distinct advantages. By assessing the source location likelihood across the structure, a metric is provided that quantifies how probable is the AE event at a given location. By assessing this likelihood over the surface of the structure, multiple possible damage locations are able to be flagged, as well as providing an operator with a greater level of insight than single point predictions. Secondly, by using coordinate values as inputs to the GP, the noise-free input assumption that standard Gaussian process regression requires is more closely upheld than in the inverse case considered by Hensman et al. [31]. There is also the flexibility to deal with targets that have an associated level of uncertainty, which as discussed extensively, is an important functionality for modelling AE arrival times. Lastly, it will be shown that the proposed strategy returns a favourable accuracy as the density of the training grid is lowered, particularly in comparison to other leading AE localisation methods. Reducing the number of training measurements will be particularly advantageous for large structures, where generating a fine grid of artificial sources is very time-consuming.

### 3.2.2 Learning $\Delta T$ maps

Following the data acquisition process as detailed in Chapter 2.7, the data set that has been gathered needs to be processed to reflect the relationship that is being modelled. Considering a single artificial AE event, there will be eight captured waveforms, each corresponding to an individual sensor. To extract the  $\Delta T$  information, the arrival time of the waveform at each sensor is first estimated through application of the AIC picker (see Chapter 2.3). As an example, Figure 3.2 demonstrates the identification of the onset time of an event at each of the 8 sensors. The onset times may then be used to calculate  $\Delta T$ s across all sensor pairs. Given that the use of eight individual sensors results in 28 unique pairing combinations, for each AE event, there will be a corresponding  $\Delta T$  vector containing 28 entries.

Having established a suitable data set, the process of learning a relationship between source location and  $\Delta T$  can begin, where a GP regressor model is learnt for each sensor pair; that is, a functional mapping  $(x, y) \rightarrow \Delta T_j$  for each  $j$  is learnt, where  $j$  indexes the sensor pair with  $x, y$  the spatial coordinates. The kernel selected for each of the models is that of a Matérn 3/2 kernel, following the demonstration of its

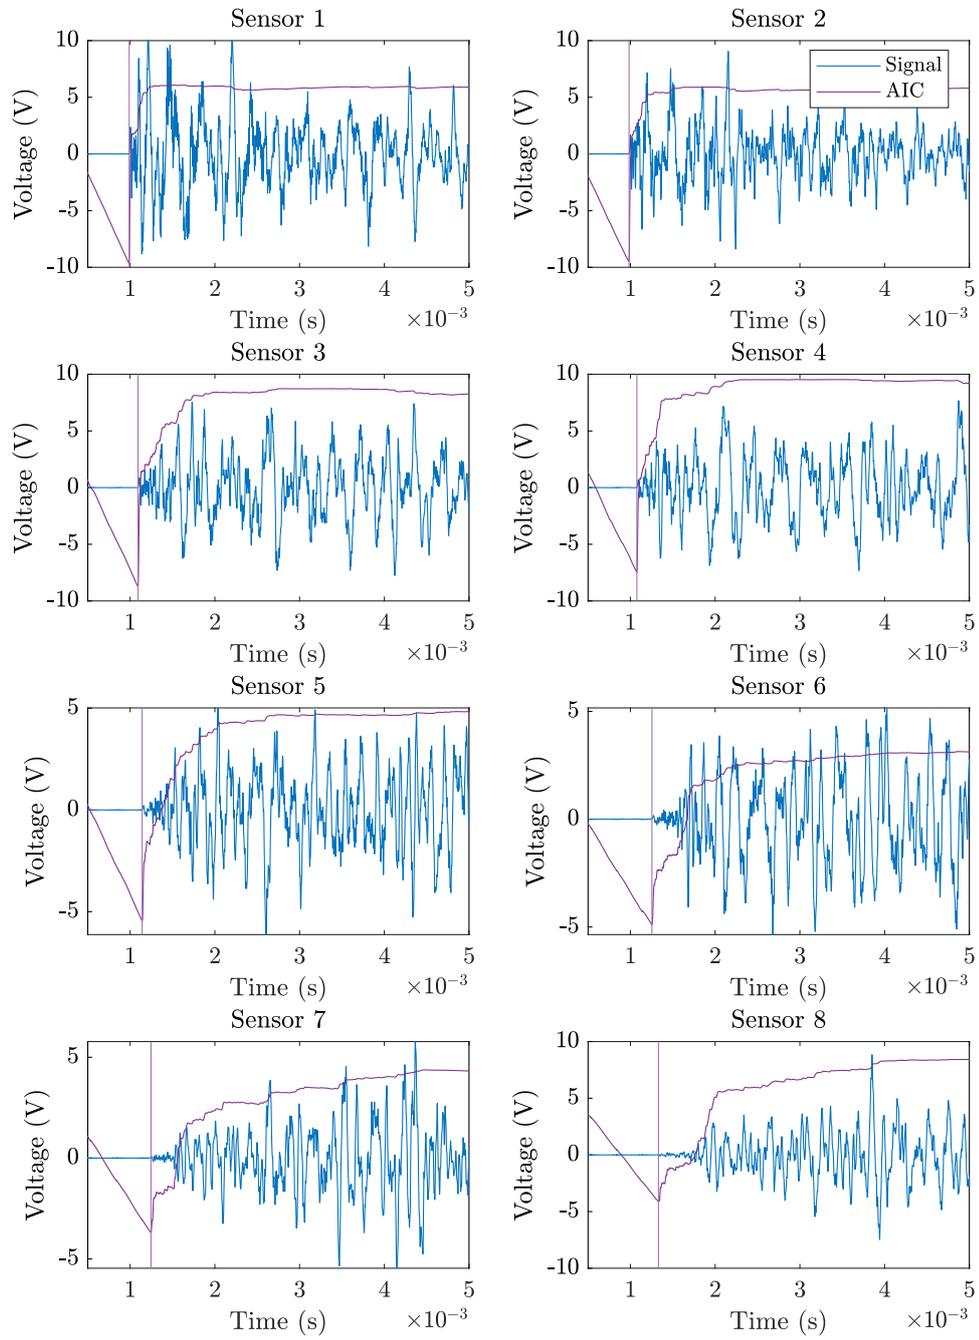


Figure 3.2: Arrival time identification of one event at each sensor using the AIC picker.

ability to perform well with <sup>2</sup>[engineering data] [15]. As the GP is being employed in a spatial modelling context that is asymmetrical about each dimension, it makes sense to switch to an anisotropic kernel where a unique length scale is defined for each input dimension. To transform equation (3.19) from an isotropic form, the kernel can be represented as,

$$k_{3/2}(\mathbf{x}_i, \mathbf{x}'_i) = \sigma_f^2 (1 + \sqrt{3}r) \exp(-\sqrt{3}r), \quad (3.24)$$

where

$$r = \sum_{d=1}^D \frac{\|x_{id} - x'_{id}\|}{l_d}, \quad (3.25)$$

with  $d$  denoting a specific dimension.

To visualise the process of learning a  $\Delta T$  map, Figure 3.3 plots all of the acquired data for sensor pair 2-8, with Figure 3.4 providing a denser mapping that has been learnt by the GP.

### 3.2.3 Source location likelihood

In a standard regression setting, one is interested in capturing a relationship between two sets of variables; those termed the independent or explanatory variables, and those for which the objective is to predict the values of (the dependent variables). Under a Gaussian process framework, the independent variables are generally assigned as the input points, whilst the quantities one wishes to predict are placed on the targets. In the context of source localisation, the desired objective is to estimate the location of an observed AE event, suggesting that a location coordinate be designated as a target with a  $\Delta T$  value (for example) assigned as an input. Although this direct means of modelling an inverse problem has demonstrated success [31], it is not reflective of the general approach to regression where the independent variables are assigned as model inputs. The method proposed in this chapter avoids such an issue

---

<sup>2</sup>Engineering data is used to indicate that the data may contain irregular fluctuations due to underlying noise or outliers, and that the overall behaviour of the data may be rougher than draws from other kernel choices that enforce stronger smoothing assumptions such as the squared exponential.

by considering a forward model, where the potential source location coordinates, which are indeed the explanatory ones, are used as model inputs, with the  $\Delta T$  values assigned as model targets.

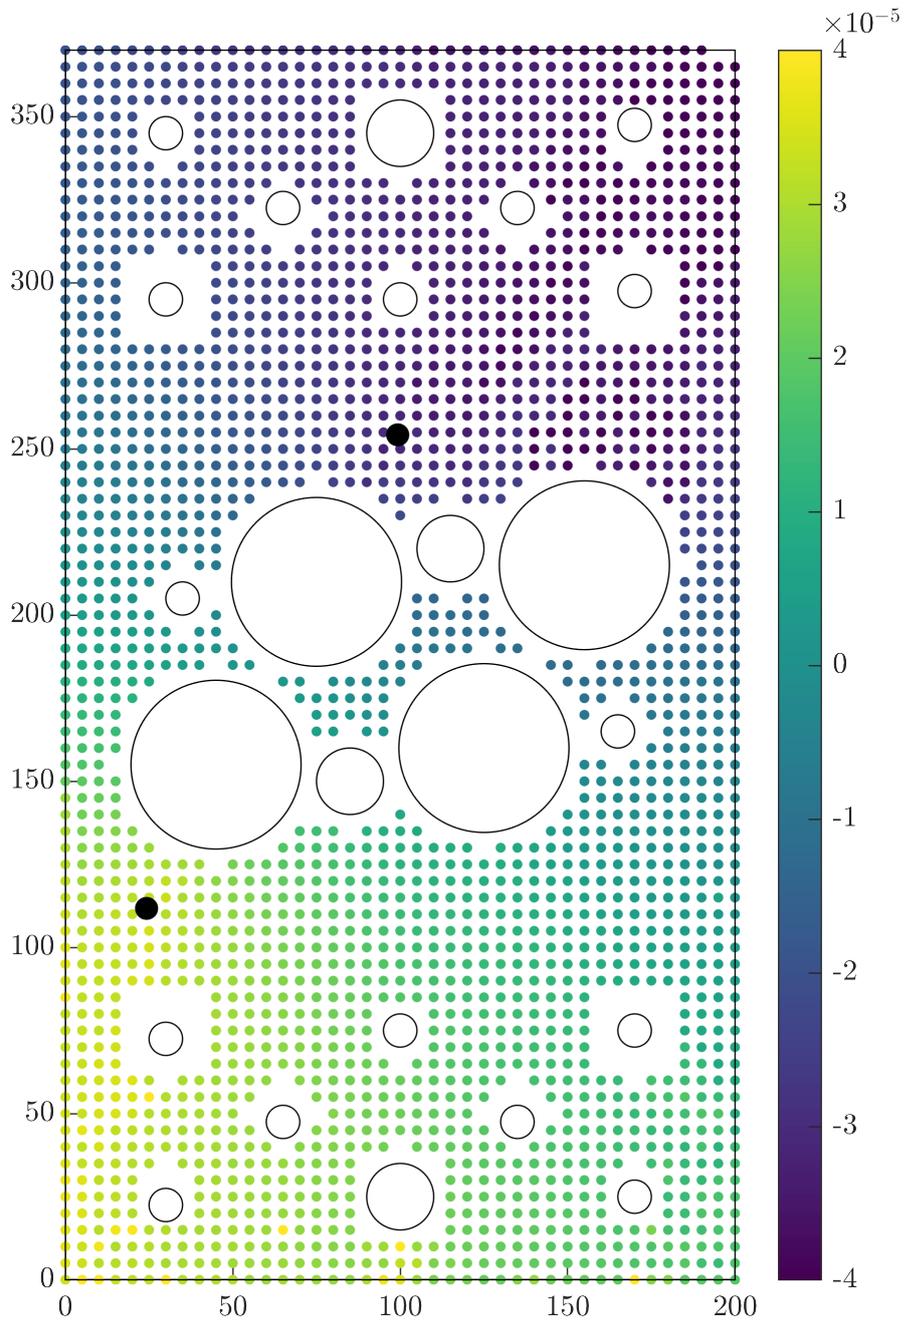


Figure 3.3:  $\Delta T$  values of artificial AE sources for sensor pairing 2-8, where each circle represents one event. Color bar indicates difference-in-time of arrival in seconds.

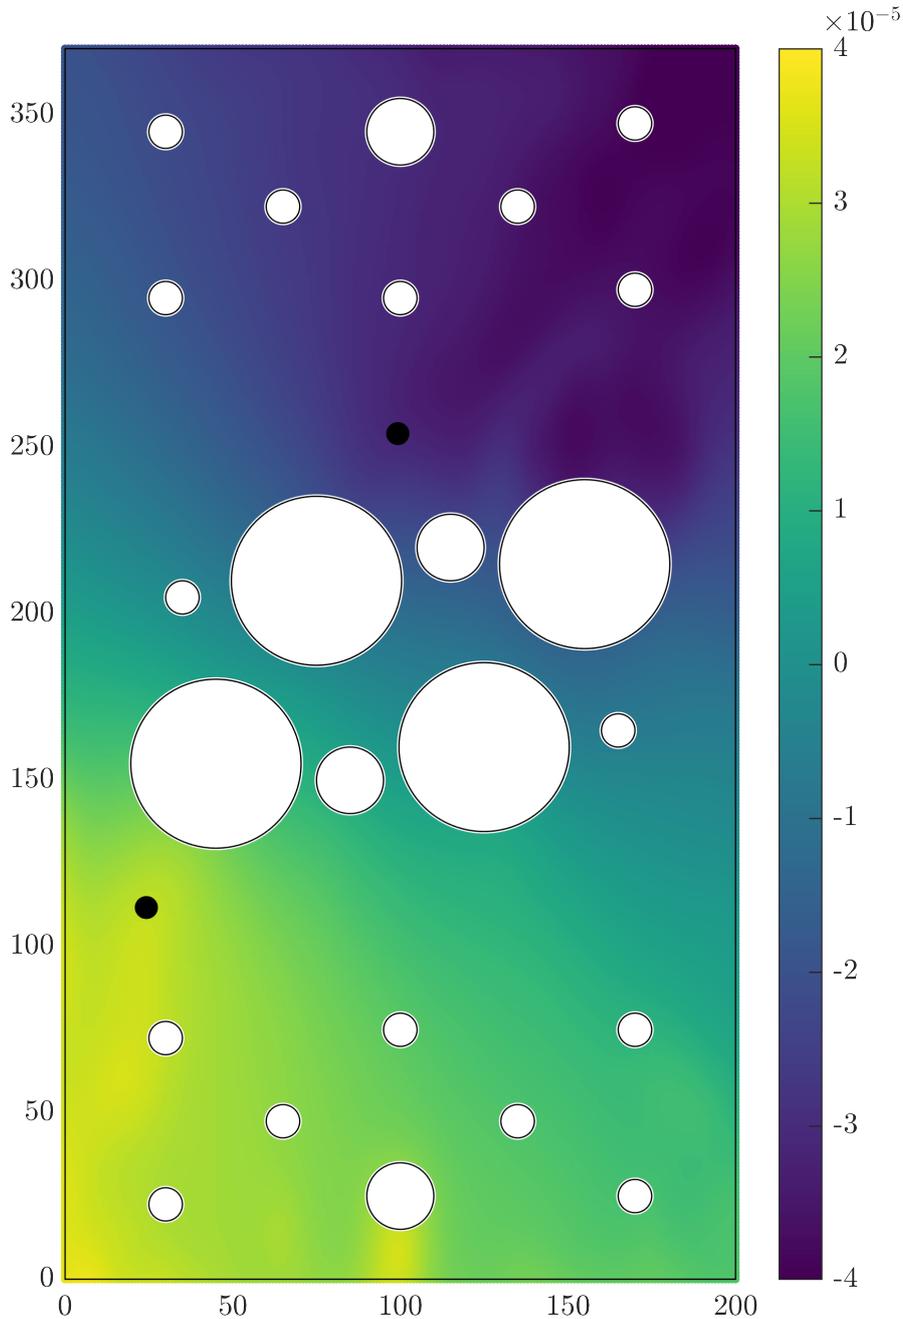


Figure 3.4:  $\Delta T$  map for sensor pairing 2-8 learnt by a Gaussian process. Color bar indicates difference-in-time of arrival in seconds.

An additional benefit of adopting a forward model is that the underlying assumption in standard GP regression of noise-free inputs [92] is more closely met than if taking an inverse model approach. Although the AIC method yields a reasonable prediction of the onset times, they are still only estimates of the true values. The  $\Delta T$  values should, therefore, be interpreted as being corrupted with some level of noise, and

therefore if placed on the inputs of the GP, do not uphold the required assumption. This will be particularly pertinent when operating outside of the ideal conditions of the laboratory, where increased levels of external noise will have more of a disruptive influence on the time of flight estimates. In the case of the AE source locations, the user has direct control of the position at which these events are generated. Therefore, providing that the experiment has been set up with good practice in mind (such as proper calibration of positioning devices), it is viable to assume that the noise present on the AE event locations is insignificant relative to that on the  $\Delta T$  measurements.

By placing grid coordinates on the input to the GP, the location of a future observed AE event can no longer be estimated solely by the predictive distribution. Instead, having mapped a grid of coordinates onto their corresponding  $\Delta T$  values via GP regression, it is possible to assess the likelihood of a new set of observed  $\Delta T$ s for each sensor pair  $\mathbf{y}_{obs} = \{y_{obs,1}, y_{obs,2}, \dots, y_{obs,J}\}$  given the learnt functional mapping from the  $(x, y)$  coordinates to a measured  $\Delta T$  in the training data, again for each sensor pair  $j = 1, 2, \dots, J$ . For the set of  $J$  models, this means assessing,

$$\log p(y_{obs,j} | \mathcal{D}, \mathbf{x}_*, \mathcal{M}_j) = -\frac{1}{2} \log \mathbb{V}[\mathbf{f}_{*,j}] - \frac{(y_{obs,j} - \mathbb{E}[\mathbf{f}_{*,j}])^2}{2\mathbb{V}[\mathbf{f}_{*,j}]} - \frac{1}{2} \log 2\pi, \quad (3.26)$$

where  $\mathcal{D} = \{\mathbf{X}_i, y_{i,j}\}_{i=1}^N$  is the set of  $N$  training pairs where, for each  $i$ ,  $\mathbf{X}_i$  is the  $(x, y)$  coordinate for that observation and  $y_{i,j}$  is the  $\Delta T$  for observation  $i$  at sensor pair  $j$ . Equation (3.26) gives the likelihood of a newly observed  $\Delta T$  from an unknown source location given a candidate location  $\mathbf{x}_*$  for a given sensor pair, i.e. model  $\mathcal{M}_j$ . Although this likelihood remains conditioned on the hyperparameters, it is assumed that these are now fixed given the training stage of the model and they are not shown to avoid clutter. One also may notice the equivalence between equation (3.26) and the likelihood term arising under the learnt model equation (3.23).

By calculating the exponential of equation (3.26) at a set of candidate points across the structure (i.e. for  $\mathbf{x}_* = X_*$ ), the likelihood that the event originated from each point can be computed. Defining this set of candidate locations as the predictive grid, it is clear that the accuracy of using equation (3.26) to identify the source origin will be dictated by the density of the grid. When implementing a GP, the main computational burden lies with the inversion of the covariance matrix, scaling as  $\mathcal{O}(N^3)$ . Given that this operation is exclusive to the training stage, defining a suitably dense predictive grid will generally not be the dominant factor in computing

time. What the size of a ‘suitably dense’ predictive grid will be, however, is dependent on factors such as the complexity of the wavefield, number of training points used, expected/desired detectable defect size and computational availability. As such, these variables dictate that the sizing of the predictive grid should be constructed on a case-by-case basis.

To demonstrate the method visually, the data is split into a training and testing set, where each data point consists of the source location and the corresponding  $\Delta T$  vector for an artificial event. For a single AE event randomly selected from the testing set, Figure 3.5 shows the likelihood that the corresponding  $\Delta T$  value originated from each location on the predictive grid. The results are shown for four sensor-pairings, where the red cross indicates the true location.

Investigating each of the maps, it can be seen that rather than providing a unique solution, there exist contours of highly probable locations for each sensor pair that will generally intersect at the true location. This is of course expected; acoustic emission theory states that when considering difference-in arrival times of the fastest wave mode to localise a source, providing the source does not sit on a direct path between two sensors, a minimum of three sensors are required to provide a unique solution. To account for this behaviour, it is necessary to marginalise over each of the individual models, which is equal to the following,

$$p(\mathbf{y}_{obs} \mid \mathcal{D}, \mathbf{x}_*) = \int p(\mathbf{y}_{obs} \mid \mathcal{D}, \mathbf{x}_*, \mathcal{M}) p(\mathcal{M}) d\mathcal{M} \quad (3.27)$$

By observing that the likelihood in each sensor pair depends only on the  $\Delta T$  between those sensors ( $y_{obs,j}$ ) and that the collection of sensor pairs forms a finite and discrete set of possible  $\mathcal{M} = \{\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_J\}$ , the integral in equation (3.27) can be transformed into a finite sum.

$$p(\mathbf{y}_{obs} \mid \mathcal{D}, \mathbf{x}_*) = \sum_{j=1}^J p(y_{obs,j} \mid \mathcal{D}, \mathbf{x}_*, \mathcal{M}_j) p(\mathcal{M}_j) \quad (3.28)$$

Placing an equal importance on all sensor pairs, it is possible to set  $p(\mathcal{M}_j) = 1/J \forall j \in \{1, 2, \dots, J\}$ . Therefore,

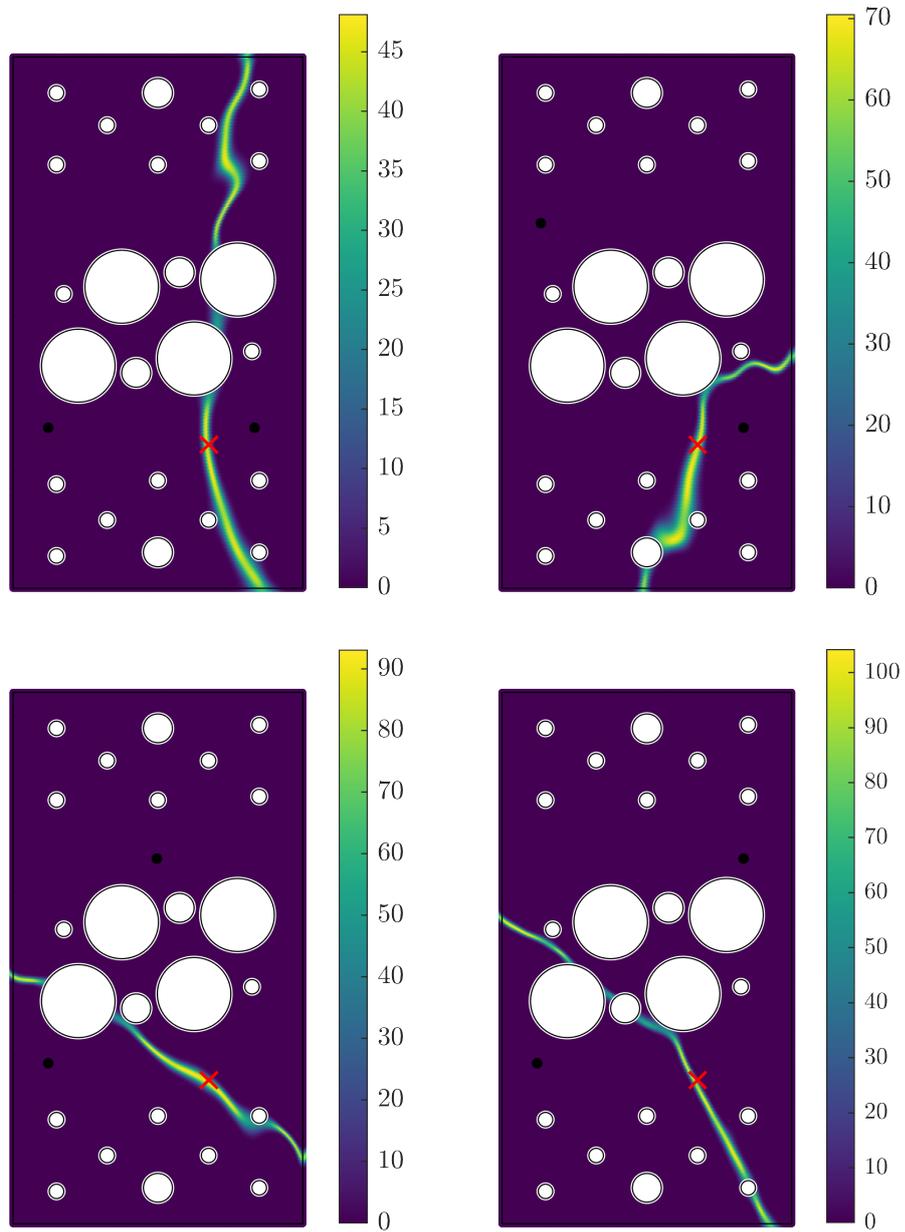


Figure 3.5: Source location likelihood maps corresponding to four individual sensor combinations, where the colour bar represents the likelihood value. The sensor pair used is represented by black circles, whilst the red cross marks the true source location.

$$p(\mathbf{y}_{obs} | \mathcal{D}, \mathbf{x}_*) = \sum_{j=1}^J p(y_{obs,j} | \mathcal{D}, \mathbf{x}_*, \mathcal{M}_j) \quad (3.29)$$

The task of predicting the true source location for the measured emission event can then be framed as maximising the likelihood of the observed  $\Delta T$ s given the candidate location  $\mathbf{x}_*$  and the previously observed training data  $\mathcal{D}$ . In other words to determine,

$$\hat{\mathbf{x}}_* = \underset{\mathbf{x}_*}{arg\ max} \{ \log p(\mathbf{y}_{obs} | \mathcal{D}, \mathbf{x}_*) \}$$

where the candidate AE source location  $\mathbf{x}_*$  with the highest likelihood of generating the observed  $\Delta T$ s from the observation  $\mathbf{y}_{obs}$  is determined to be the predicted source location,  $\hat{\mathbf{x}}_*$ . The advantage of marginalising across the possible models (sensor pairs) is that one is not required to choose a best subset of sensors with which to proceed, instead it is possible to capture the information from all possible measurements and combine them in a consistent manner. It also allows, if a-priori knowledge is available, weighting of different sensor pairs based on their efficacy in identifying source locations through modification of  $p(\mathcal{M})$ .

### 3.3 Results and discussion

Following the implementation of the localisation strategy detailed in the previous subsection, for the same test point used in Figure 3.5, Figure 3.6 represents the marginal likelihood across all sensor-pair models that the event originated from locations across the plate. Inspecting Figure 3.6, it can be seen that there exists a region of high likelihood around the true location, demonstrating that the origin of the AE event has been correctly located. If desired, a single estimate of the origin location can then be made by considering the location on the map that maximises the marginal likelihood.

The key advantage of the probabilistic maps is that the source location likelihood can be evaluated at any number of locations across the plate. The user is then flexible to decide on how to proceed. The simplest option is to use the algorithm in isolation, where the maximum likelihood solution can be used to identify the most likely location of the AE source. Alternatively, it is also possible to feed the location likelihoods into a wider risk-based SHM framework, where likelihood values are used

to inform the calculation of risk associated with some given event [6, 114, 115].

In addition to identifying the highest likelihood event location, the probabilistic source location maps also provide an intuitive visualisation of the associated confidence, which can be used to inform an operator how large a region of the structure may warrant closer inspection. For example, in the case that the algorithm identifies the most likely source location with high certainty, one would expect to see a single compact region of high likelihood, as observed in Figure 3.6, and therefore requires only a small area of the structure to be inspected. However, where a prediction is made with less certainty, this may be reflected in either a single region of high likelihood becoming more spread, or through the emergence of multiple highly probable locations.

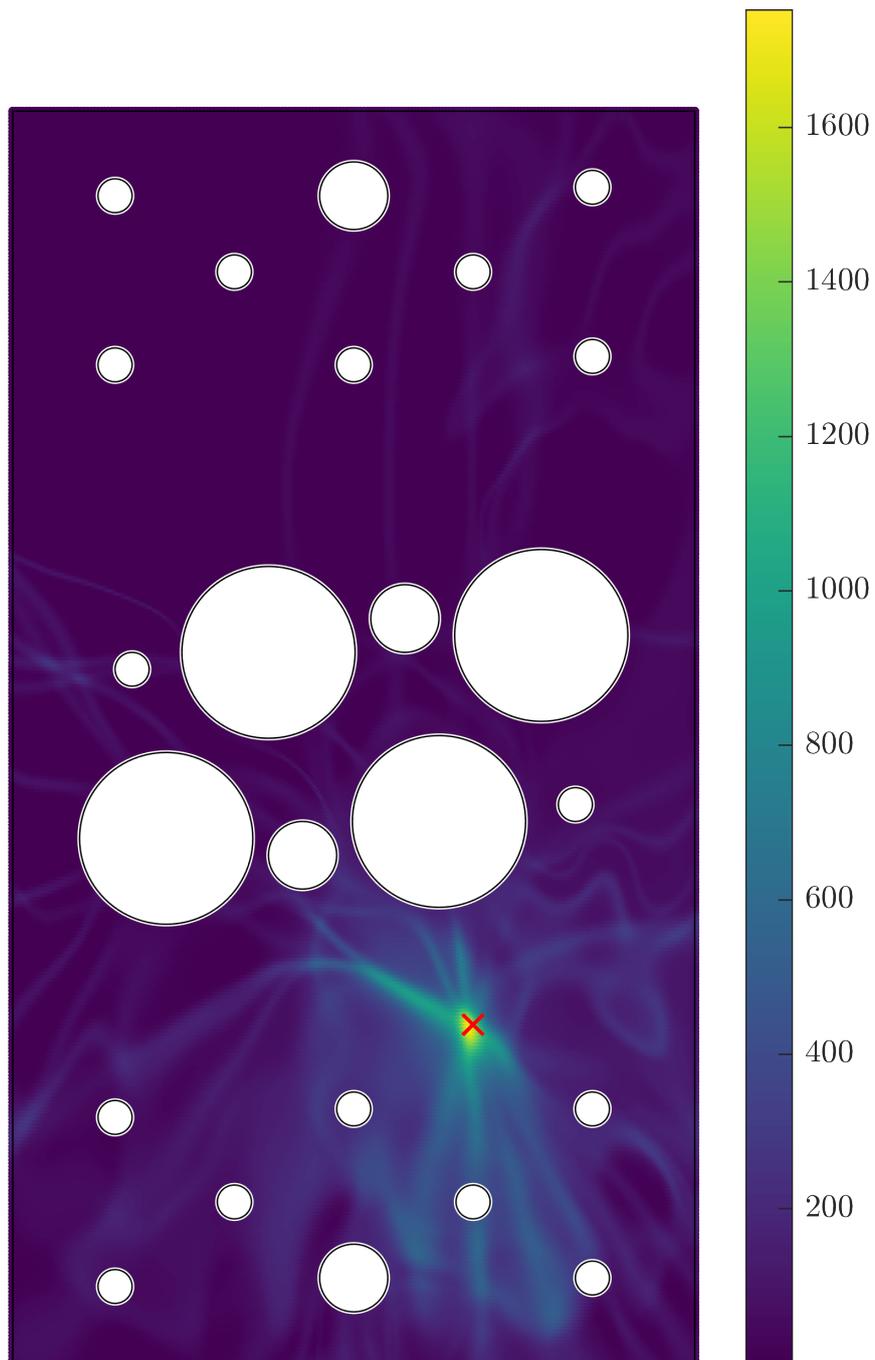


Figure 3.6: Mapping that quantifies the combined source location likelihood of an AE event. The true origin is given by a red cross and the colour bar represents the combined likelihood value.

To demonstrate a prediction with lower confidence, Figure 3.7 provides the source location likelihood map for a test point where the region of high likelihood is far more dispersed about the true location than that of the previous example, and indicates

that a larger region should be investigated further. In this case of low confidence, the ability to flag an area/number of potential source locations will prove more informative over considering single point estimates with an associated uncertainty bound.

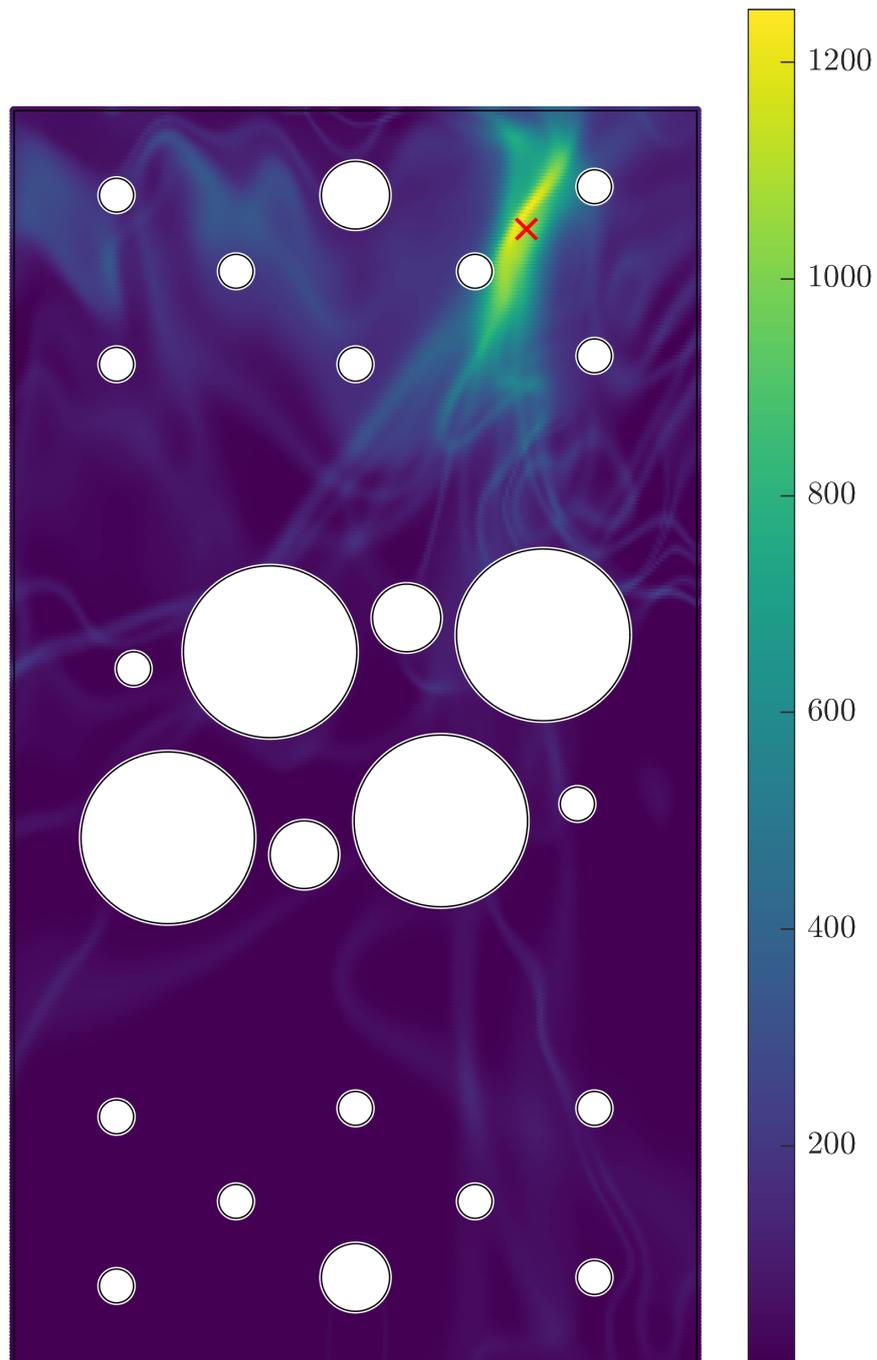


Figure 3.7: Mapping with an increased uncertainty around the location of the AE event. The true origin is given by a red cross and the colour bar represents the combined likelihood value.

### 3.3.1 Single source location prediction

Whilst single location predictions are not the sole focus of this work, it is still useful to consider them as a way to assess the general performance of the localisation algorithm. To do this, the predicted location (i.e. the most likely) of each observation in the test set is compared to the true value, with the results illustrated in Figure 3.8. Upon visual inspection, it can be seen that there is a high level of agreement between the predicted and true location for the majority of points within a test set of 100 events. Assessing the root-mean-squared-error (RMSE) of the test set according to equation (3.30) returns a value of 4.57 mm, which given the dimensions of the test structure, demonstrates the accuracy of the proposed approach.

$$RMSE = \sqrt{\frac{\sum_{x,y} ((x,y)_{pred} - (x,y)_{true})^2}{N_{test}}}. \quad (3.30)$$

### 3.3.2 Sensitivity to training grid size

One of the main costs associated with implementing a data-driven localisation strategy is the acquisition of training data, requiring artificial AE events to be generated across the test structure. Whilst autonomous data-collection solutions have been suggested [31], it is still beneficial for the number of training events to be reduced so that less time has to be spent collecting data. To investigate how sensitive the localisation algorithm is to the number of data points used in the training phase, the RMSE of single-point predictions using training sets with different grid spacing is investigated, where grid spacing is defined as the distance between each training point in both the  $x$  and  $y$  direction. For each training grid size, an averaged RMSE across ten testing sets that each consist of 100 randomly selected locations is then assessed, with the results plotted in Figure 3.9.

An obvious first conclusion to make is that as the spatial resolution is increased, which corresponds to reducing the grid spacing, the RMSE decreases. This is entirely expected; a training set representative of a denser grid will require less of the feature space to be interpolated over. A more notable observation is that as the grid spacing becomes larger, the corresponding increase in error is of a significantly smaller factor. For example, shifting the grid spacing from 5mm to 20mm - which is a 300% increase

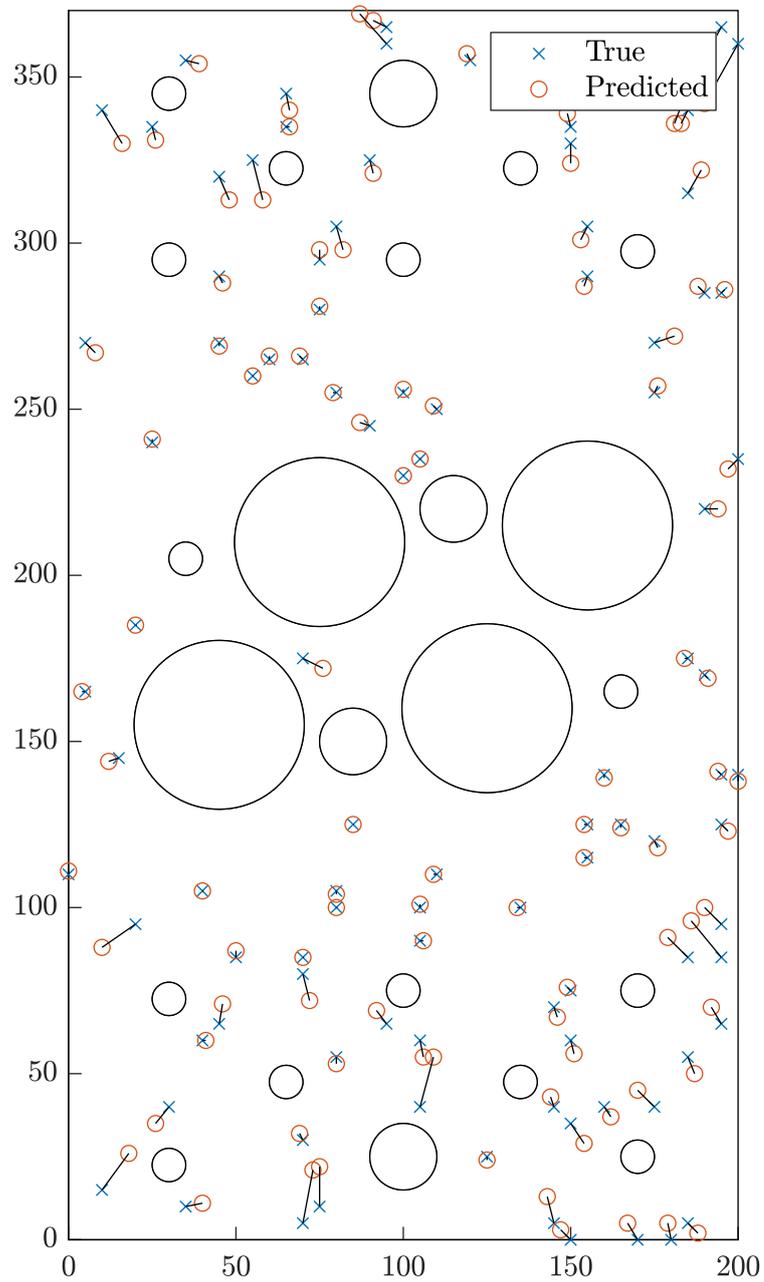


Figure 3.8: Comparison between the predicted and true AE event location for 100 testing events.

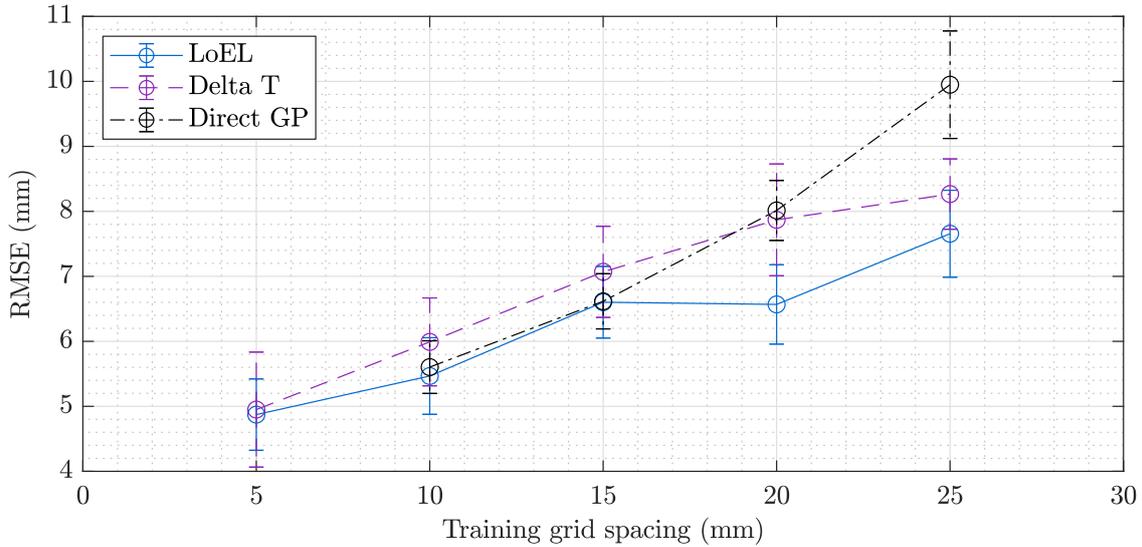


Figure 3.9: Mean and standard deviation of RMSE for 10 test sets as the spacing between the number of training set observations is varied. A grid spacing of 5/10/15/20/25mm corresponds to 2177/599/302/160/107 training points.

- returns only a 34.4% rise in error. Given that increasing the grid spacing from 5mm to 20mm would allow the required number of training events to be reduced from 2177 (at 5mm grid spacing) to 160 (at 20mm), in many circumstances, the relatively small increase in error could be justified by the considerable amount of time that would be saved during the acquisition stage. Additionally, as the data points are naively removed in a uniform manner, following a more informed training point selection process, such as removing less events around geometrical features, would likely see more favourable errors as the number of training measurements is reduced.

### 3.3.3 Comparison with existing methods

To compare the LoEL method with similar state-of-the-art approaches, the exercise completed in the above subsection is repeated using both the latest variant of the Delta  $T$  mapping technique (details in [86]) and also the method presented by Hensman et al. [31], which will be termed the direct GP approach. As this is the authors' implementation of both of these methods, the training and testing locations will differ to those presented in the results of [31]. Additionally, it should be noted that a grid spacing of 5mm was not computed for the direct GP due to excessive computational demands, which was estimated to be two weeks run time on a standard

workstation due to re-training the model for each test set.

Considering Figure 3.9, LoEL compares favourably with both Delta  $T$  and the direct GP approach, returning a lower RMSE at the majority of the training grid resolutions. An increased predictive performance is particularly evident at the larger grid spacings, where it can be seen that the difference in error between the proposed approach and those for comparison increases. This enhanced interpolative capability can be ascribed to two factors. Firstly, as the behaviour of the acoustic emission waves across the plate is inherently nonlinear, the accuracy of the interpolation is dependent on the ability to represent such complexity. In the proposed approach, the choice of kernel function within the GP prior allows nonlinear relationships to be captured, whilst for Delta  $T$ , the method is limited to fitting a linear regression between known data points. Therefore, as training data become more sparse, the assumption of a linear functional form will become increasingly restrictive, leading to a worsened interpolative performance. Secondly, as linear interpolation fits exactly through known targets, it is not possible to directly account for the process noise associated with the  $\Delta T$  values. Conversely, under the LoEL approach, the  $\Delta T$  measurements are considered to be the summation of a true underlying value, and some Gaussian noise term, which, for reasons discussed in Section 3.2.3, is more reflective of the underlying process being modelled.

For the direct GP, the radial basis function (RBF) kernel (equivalent to an SE) that is implemented is capable of modelling nonlinear behaviour and therefore overcomes the first limitation of the Delta  $T$  method. However, the infinitely differentiable property of the RBF kernel likely makes too strong a smoothing assumption on the data, which can be mitigated against by switching to a Matérn kernel as done here. Perhaps more crucially, by mapping from  $\Delta T$  measurements to a spatial location (direct approach), the assumption of noise-free inputs required by standard Gaussian process regression is not upheld, whilst the uncertainty associated with the TOA measurements is unable to be incorporated into the model as is achieved with the framework proposed here. This will clearly be detrimental to the localisation accuracy, particularly for the two least dense training sets where the error rises sharply in comparison to the method proposed in this work. An additional benefit of the proposed approach over the direct GP is a significant reduction in learning time. For the direct GP, a Gaussian process prior is learnt for each grouping of three or more sensors, which for eight individual sensors, corresponds to 438 GPs. In the case of the method proposed here, a GP is trained for each sensor pairing, reducing the number of GPs

to 28. Although at prediction time LoEL requires more function evaluations than the direct GP, as discussed earlier, the main computational bottleneck with GPs exists in the training stage, scaling  $\mathcal{O}(N^3)$  as opposed to at prediction time where the mean is computed at  $\mathcal{O}(N)$  and the variance  $\mathcal{O}(N^2)$ . Hensman et al. [31] do present some discussion around mitigating this elevated computational cost in their approach, for example, by using the same set of hyperparameter values for each array of the same dimensionality. However, the method proposed in this work remains more computationally efficient.

### 3.4 Conclusions

Within this chapter, a probabilistic framework for localising acoustic emission sources in complex structures has been established, offering a number of distinct advantages over the state-of-the-art:

- The likelihood that an observed AE event originated from a given location is quantified across the surface of the structure, where the location that returns the maximum likelihood identifies the most likely source origin.
- By using a forward model approach and placing the locations of the training events on the GP inputs, the model is representative of how regression problems are generally treated, as well as more closely upholding the noise-free input assumption than in the case of a direct strategy.
- The model allows measurement noise on the TOA estimates to be naturally included.
- There is the potential ability to flag multiple possible damage locations through quantification of source location likelihood over the whole structure.
- The probabilistic mapping offers an intuitive visualisation of the confidence associated with a location prediction, which can assist an operator in deciding how large of an area to inspect.
- As the number of data points used in the training phase is reduced, when applied to a complex-like structure, an enhanced interpolative performance is demonstrated in comparison to other similar localisation methods. This

behaviour is imperative when looking to implement localisation strategies in industrial settings.

One potential limitation of the proposed approach is the reliance on the retrieval of accurate time of arrival estimates. Although the Gaussian process framework allows some degree of uncertainty in the onset times, learning the  $\Delta T$  maps requires there to be a sufficient level of accuracy in the TOA values. In this work, an AIC picker was adopted to determine the onset times, which comes with the caveat that signals fed into the picker must only contain a single AE event. In scenarios where one is reliant on data containing multiple events, such as in a streaming context, then an additional step in the data processing stage would be required. The authors believe that an adaptive thresholding strategy would be suitable here [9], however, this specific aspect will not be revisited in the contributions made by the work of this thesis.

With regard to the nature of the onset time uncertainty, the assumption is made that the noise process can be modelled as an additive, fixed variance Gaussian, and is, therefore, insensitive to the location of the event relative to the sensor pair. For larger structures in particular, it is likely that adopting a noise process that is dependent on the position of the event would be more suitable, allowing for a noise that is dependent on the relative location between an event and receiving sensor. To this end, Chapter 4 will consider the use of heteroscedastic Gaussian processes - a more flexible Gaussian process model that allows an input-dependent noise to be learnt.

It would also be desirable to explore alternative approaches for combining the predictions across each of the sensor pairs. For instance, a weighting scheme that can reflect the uncertainty resulting from the location of a particular sensor pair will result in a more robust approach, allowing sensors that better cover a given area to make a greater contribution to the overall source location prediction. As the availability of training measurements is reduced, such a functionality will be particularly beneficial. Again, proposed solutions and discussions on this matter will be presented in Chapter 4.



# HETEROSCEDASTIC GAUSSIAN PROCESSES FOR ENHANCED UNCERTAINTY QUANTIFICATION

### Highlights:

- *A  $\Delta T$  forward model is developed that more appropriately models the uncertainty on the onset times.*
- *It is shown that the onset time features are heteroscedastic, which may be handled with a local (input-dependent) noise process in the observation model.*
- *Predictions from individual sensor pairs are weighted according to how well a particular area of the structure is covered by the pairing, improving the overall prediction accuracy and confidence.*

In the previous chapter, it was demonstrated how Gaussian processes may be used to form a probabilistic framework for localising acoustic emissions. Although effective, one fundamental assumption made by standard Gaussian process regression is that the noise present on the targets is constant across the input space. When considering the spatial modelling of  $\Delta T$  values, as the complexity of the propagation path between a measurement location and sensor pair grows, the variability in the onset

time features will increase, which can be modelled as an increased noise presence on the true function values. There are a number of mechanisms that can be attributed as the cause of this behaviour; as a wave travels further, the effects of phenomena such as scattering, attenuation and dispersion all become more prominent, resulting in a more challenging onset time identification. Similar effects may be seen when geometrical discontinuities in the propagation medium lie on the direct propagation path between an event and sensor location, as well as where an individual pairing of sensors offers poor spatial coverage of a structure (e.g. two sensors placed next to one another). As the arrival times form the basis of the  $\Delta T$  maps, this amplified noise then propagates forward into the location predictions. The combination of all of these effects results in the emergence of a spatially-dependent noise process; that is, the noise is a function of the location in the input space, rendering a uniform noise model suboptimal.

In this chapter, the use of heteroscedastic Gaussian processes will be presented as an approach for learning  $\Delta T$  maps in a way that allows this input-dependent noise to be captured. The additional source of uncertainty may then be fed into the localisation framework developed in the preceding chapter, returning an improved localisation performance than in the case of the standard model.

## 4.1 Spatially-dependent noise in $\Delta T$ values

As previously introduced, by placing a standard Gaussian process prior over the latent  $\Delta T$  functions, the assumption made is that the noise present on these function evaluations is uniform across the input domain. In other words, the variance of the set of targets is fixed. However, if the level of variability of the targets fluctuates, where the exact level of fluctuation is dependent on the location of the corresponding input, then this assumption breaks down. For the  $\Delta T$  features, their associated uncertainty is largely dominated by two factors, sensor noise and the estimation of the onset times. Whilst it is reasonable to assume that the epistemic uncertainty associated with a particular sensor remains constant - that is to say the accuracy of sensors can be considered fixed - the variability on the onset times should be a function of the complexity of the propagation path between the source and sensor, and as such, the uncertainty on the onset times will vary across the input space, resulting in a spatially/input-dependent process that is captured here as noise.

The existence of an input-dependent noise variance is not uncommon in regression tasks, and occurs regularly in many fields such as econometrics, financial forecasting and meteorology [116–119]. This type of noise is more generally referred to as *heteroscedasticity*, where a collection of random variables is heteroscedastic if their variance is not constant. To visualise the presence of heteroscedasticity on the arrival time differences, Figure 4.1 plots the  $\Delta T$  values collected for sensor pair 4 & 6. Investigating the figure, it can be seen that the onset times vary fairly smoothly across the plate, apart from in the lower left of Figure 4.1. In this area of the structure the sensor coverage is poor, with emissions originating from this region having to travel further to the closest receiving sensor than the rest of the plate, as well as being heavily obscured by the holes. This elevated propagation path complexity then results in onset times that, locally, shift further away from some smooth mean value. It can, therefore, be deduced that the  $\Delta T$  features do not follow a constant variance process, and instead, exhibit heteroscedastic behaviour.

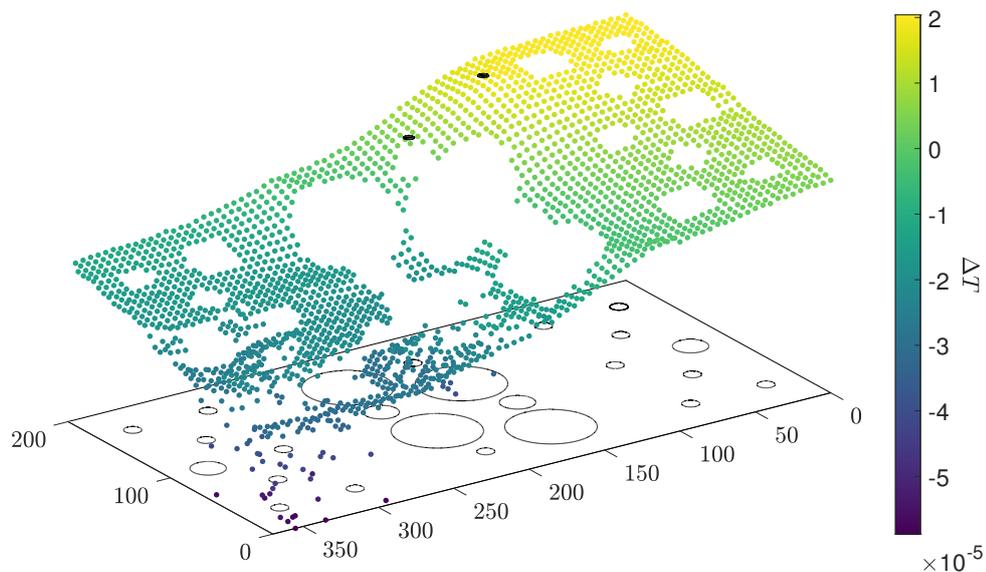


Figure 4.1: True  $\Delta T_{4,6}$  values, with sensor locations marked by the black circles. The targets are plotted on both the  $z$  and colour axis. The non-smooth fluctuation of the features in some areas of the plate indicate the presence of heteroscedasticity.

When fitting a standard Gaussian process regression model, by enforcing a fixed variance noise process, there is no mechanism for capturing this locally varying noise. This local variation can be more clearly visualised by considering the residual between predictions made by the GP and the true targets, and is plotted in Figure 4.2 for the same sensor pairing considered above. With a uniform training grid, for

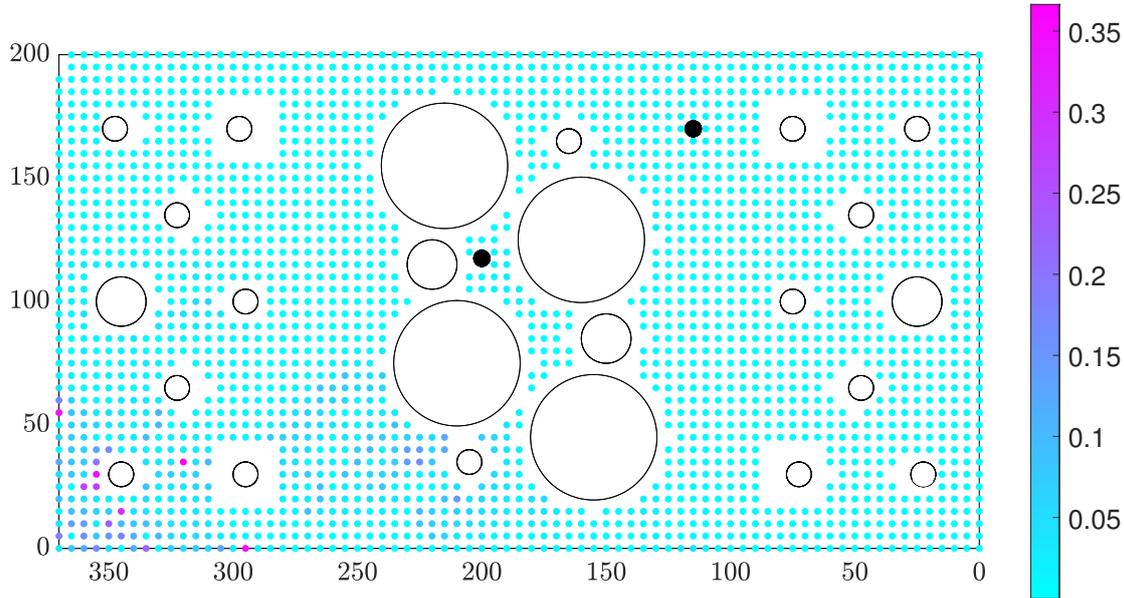


Figure 4.2: Residual between GP prediction and true target values for sensor pairing 4 & 6.

locations in the input space where a fixed noise model is suitable, this residual should appear almost constant. However, the figure clearly shows that in some areas of the plate, the residual value deviates significantly away from zero. This behaviour is most prominent in the bottom left corner of Figure 4.2, indicating that there is some spatially-correlated structure in the residual values that cannot be captured by a constant noise variance assumption. As expected, this region coincides directly with that of the onset times that vary the most about some smooth trend in Figure 4.1.

For sensors 4 & 6, their positioning allows for good spatial coverage of the upper middle and right side of the plate (referring to the orientation of Figure 4.2), but less so on the left, particularly the lower left. As discussed, this region of poor coverage is then where variation on the  $\Delta T$  features is largest.

Considering a sensor pairing with a more complete coverage of the entire plate, Figure 4.3 plots the residual between the GP predictions and target values for pairing 1 & 8. A first observation to make is that the presence of a locally-driven variation on the onset times still clearly exists, with there being uncaptured structure in the  $\Delta T$  times. However, the exact location and amplitude of the error is different to that of the previous sensor pairing. For example, the location in the input space at which the residual is greatest differs for both of the pairings, with sensor pair 1 & 8 seeing most error in the upper left and lower right of Figure 4.3. Additionally,

the amplitude of the error is lower for  $\Delta T_{1,8}$  than  $\Delta T_{4,6}$ , as shown by comparing the values of the residual in Figure 4.2 and 4.3 - note the difference in the color axis range between the two figures. As the variation in the onset times is dependent on the propagation path between a source location and receiving sensor, the error will change for different sensor pairs, dependent both on the individual positioning of the sensors relative to a particular location, and also the joint coverage offered over the entire structure. Generally, variation on the onset times will increase in positions that move further away from the sensors, as well as where sensor coverage is poor.

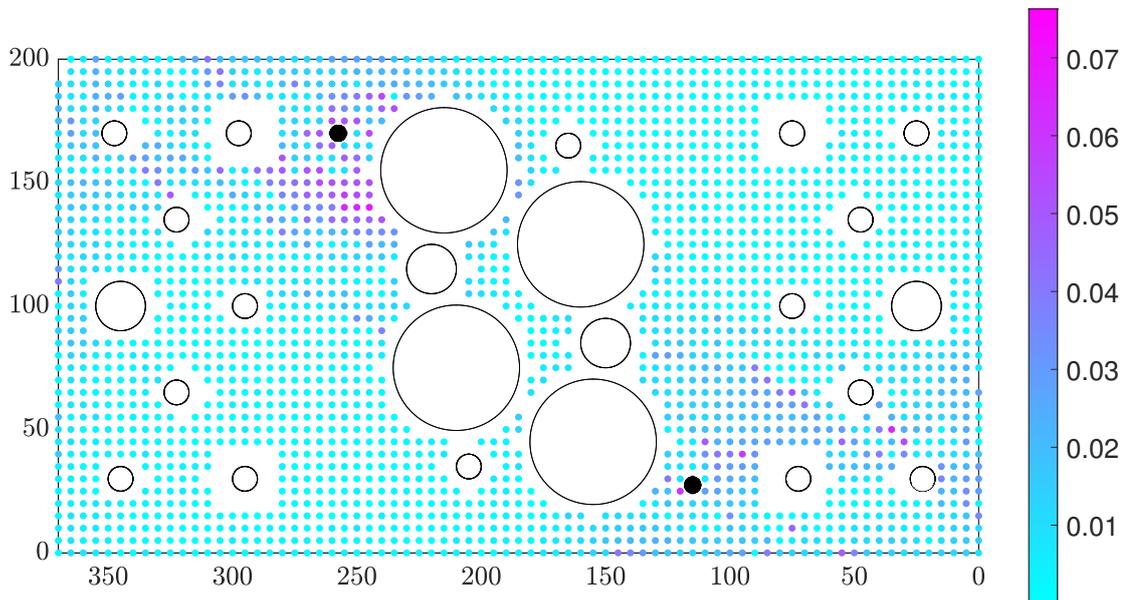


Figure 4.3: Residual between GP prediction and true target values for sensor pairing 1 & 8.

## 4.2 Heteroscedasticity in Gaussian process models

In order to transition to a heteroscedastic Gaussian process model, it is necessary to remove the assumption of a constant variance noise term. In a practical sense, this translates to allowing  $\sigma_n^2$  to vary with the model inputs. The question of how one should attempt to learn this new variance now arises. A common approach is to place a Gaussian process prior over this term [120–122], effectively treating the noise variance as a random variable instead of some fixed parameter. The model now contains two GPs; one that governs predictions of the true function evaluations, and another that predicts additive noise present on these evaluations. Such a model can be defined as a heteroscedastic Gaussian process (HGP). Before moving on, it

should be noted that switching to a non-Gaussian likelihood may also be an option for capturing the trends in the data discussed above, but that the heteroscedastic model naturally lends itself to modelling the  $\Delta T$  features. Continuing with the heteroscedastic GP model, mathematically, it can be seen that one is now concerned with learning latent functions of the form,

$$y = f(\mathbf{x}_i) + \epsilon_i \quad \epsilon_i \sim \mathcal{N}(0, r(\mathbf{x}_i)), \quad (4.1)$$

where,

$$r(\mathbf{x}_i) \sim \mathcal{GP}(m_r(\mathbf{x}_i), k_r(\mathbf{x}_i, \mathbf{x}'_i)). \quad (4.2)$$

Given that the variance of the additive noise is exclusively positive, this can be enforced by taking  $r$  to be the exponential of a GP draw [122],

$$r(\mathbf{x}_i) = e^{g(\mathbf{x}_i)} \quad g(\mathbf{x}_i) \sim \mathcal{GP}(\mu_0, k_g(\mathbf{x}_i, \mathbf{x}'_i)). \quad (4.3)$$

As with any GP prior, a covariance function  $k_g$  needs to be specified, whilst a constant mean of  $\mu_0$  is considered, controlling the scale of the noise. In this sense, the HGP can be viewed as a generalisation of the standard Gaussian process, where we transition to a more expressive and flexible noise model.

Although the use of a heteroscedastic noise model enhances the predictive power of the model, both the marginal likelihood and predictive distributions (i.e. the conditional mean and variance) are no longer analytically tractable. A number of approaches have, therefore, been suggested on how inference in a HGP may be approximated.

Generally speaking, intractable distributions may be approximated in one of two ways; generate samples from the target distributions so that an empirical function representative of the corresponding probability density may be formed, or approximate the posterior directly with some tractable density [123]. The former of these methods is often approached through Markov chain Monte Carlo methods, where Markov processes are chained together and iteratively sampled from to recover a target distribution. MCMC has then been employed to recover posterior estimates of

heteroscedastic Gaussian processes by Goldberg et al. [120] and Chan et al. [124]. Although guaranteed to converge to the true distribution with sufficient samples in accordance with the strong law of large numbers [125], MCMC methods can be computationally expensive, resulting in long convergence times. To speed up HGP inference, Kerstring et al. [121] suggest taking a maximum a posterior estimate (MAP) for learning the latent noise variance, which can be computed using expectation maximisation (EM). However, as MAP estimates are limited to a single point mass, such methods can be prone to overfitting. To remedy this, Lažara-Gredilla and Titsias [122] suggest integrating over the latent variables of the model and derive a variational framework in which HGP inference can be computed. Variational approaches look to approximate the target distribution with some simpler form of distribution, defined as the variational distribution. Some measure of similarity between a candidate and true posterior is then minimised such that the predicted distribution best approximates the true posterior. Variational approaches, therefore, fall into the category of the latter type of methods that were introduced earlier for approximating distributions, with EM an example of a variational approximation constrained to a point mass (e.g. the MAP estimate). In this sense, variational inference can be viewed as a generalisation of EM.

To implement this variational approximation, following [122], as with most variational approaches, the Kullback-Leibler (KL) divergence provides the measure of similarity between a candidate and the true distribution. Under the assumption of a Normally-distributed variational posterior, the evidence can then be lower-bounded as,

$$F(\boldsymbol{\mu}, \Sigma) = \log \mathcal{N}(\mathbf{y}|\mathbf{0}, K_f + R) - \frac{1}{4} \text{tr}(\Sigma) - \text{KL}(\mathcal{N}(\mathbf{g}|\boldsymbol{\mu}, \Sigma) || \mathcal{N}(\mathbf{g}|\mu_0\mathbb{I}), K_g), \quad (4.4)$$

where,

$$R = \text{diag}\{e^{(\mu_i - \frac{1}{2}\Sigma_{i,j})}\}, \quad i, j = 1, \dots, N. \quad (4.5)$$

Equation (4.4) serves as a surrogate objective function from which the marginal likelihood can be maximised, where  $K_f$  and  $K_g$  represent the covariance matrices of each of the GPs;  $\text{tr}(A)$  is the trace of the matrix  $A$ ,  $\mathbb{I}$  is the identity matrix,  $\text{KL}(a||b)$  represents the KL divergence between two distributions  $a$  and  $b$ , and  $\boldsymbol{\mu}$  and

$\Sigma$  are the variational parameters to be learnt, which are adjustable parameters that govern the variational distribution. It is possible to reparametrise the variational parameters in terms of some positive semidefinite diagonal matrix  $\Lambda$ , reducing the number of free parameters to be learnt from  $N + N(N + 1)/2$  to  $N$  [122]. From the perspective of computational complexity, reducing the number of hyperparameters will generally speed up the optimisation time.  $\boldsymbol{\mu}$  and  $\Sigma$  are then expressed by the following expressions:

$$\boldsymbol{\mu} = K_g(\Lambda - \frac{1}{2}\mathbb{I}) + \mu_0\mathbb{I} \quad (4.6)$$

$$\Sigma^{-1} = K_g^{-1} + \Lambda \quad (4.7)$$

Given that the lower bound  $F(\boldsymbol{\mu}, \Sigma)$  is equivalent to  $F(\Lambda)$ , equation (4.4) should be maximised with respect to the variational parameters  $\Lambda$  and model hyperparameters  $\boldsymbol{\theta}$ , which by construction, jointly minimises the KL divergence between the true and approximate posterior.

Finally, the posterior distribution  $p(\mathbf{y}_*|X_*, \mathcal{D})$  should be computed in order to calculate the predictive mean and variance. However, again, it is not available in closed form. Fortunately, the variational framework provides analytic expressions for the first two moments of the approximate variational posterior  $q(\mathbf{y}_*)$ , and therefore closed form expressions can be given for the mean and variance [122],

$$\mathbb{E}_q[\mathbf{y}_*] = K_{f, X_* X} (K_{f, X X} + R)^{-1} \mathbf{y}, \quad (4.8)$$

$$\mathbb{V}_q[\mathbf{y}_*] = K_{f, X_* X_*} - K_{f, X_* X} (K_{f, X X} + R)^{-1} K_{f, X X_*} + e^{\mu_* + \frac{1}{2}\sigma_*^2}, \quad (4.9)$$

where,

$$u_* = K_{g, X_* X} (\Lambda - \frac{1}{2}\mathbb{I}) + \mu_0\mathbb{I}, \quad (4.10)$$

$$\sigma_*^2 = K_{g, X_* X_*} - K_{g, X_* X} (K_{g, X X} + \Lambda^{-1})^{-1} K_{g, X X_*}. \quad (4.11)$$

Given the assumption that each prediction is Normally distributed, it is possible to place a predictive distribution over  $\mathbf{y}_*$  given test inputs  $X_*$ .

### 4.3 Learning $\Delta T$ maps with a heteroscedastic model

Having outlined the procedure for implementing the heteroscedastic Gaussian process, learning the  $\Delta T$  mappings as required for the localisation framework can begin, where a good localisation performance relies on accurately modelling the  $\Delta T$  features. The data set introduced in Chapter 2 will again be used to enable a comparison in performance between the heteroscedastic and homoscedastic model. Additionally, unless stated otherwise, the training set used throughout this chapter consists of 10mm spaced data points across the plate (other than where holes exist), totalling 567 observations. Such a spacing is equivalent to the grid spacing of 10mm in Section 3.3.2.

To investigate the performance of the HGP, first consider the  $\Delta T$  map that has been learnt for sensor pair 3 & 5. Figure 4.4 shows the events used in training, with Figure 4.5a a mapping of  $\Delta T$  predictions. The figure shows that the HGP is able to capture the behaviour of the  $\Delta T$  feature across the plate, particularly in the regions that exhibit sharp changes in the onset times, such as areas that are shielded by the holes. Comparing with the homoscedastic model, Figure 4.5b plots the equivalent mapping learnt by the standard GP, where it can be seen that both models return similar mappings.

To explore other sensor pairings, let us extend this analysis to consider the nMSE returned by the HGP and GP across all sensor pair mappings on a test set of the remaining AE events not used in the training stage. Plotting the result in Figure 4.6, there generally only exists small deviations between the mean error obtained by the two models for many of the sensor pairs, indicating that the mean predictions of the GP and HGP are largely similar over much of the structure. With that said, there are a number of sensor pairs where the HGP does return notable reductions in nMSE; for instance, sensor pair 1 & 2 (index 1 on Figure 4.6) and sensor pair 6 & 7 (index 26 on Figure 4.6) - see Table 5.2 for all sensor pair index labels. To explain, recall that the difference between the observation model of the GP and HGP is in the form of the additive noise term. As the HGP has the capacity to capture local shifts in variation, where the HGP model will prove most beneficial in regard to

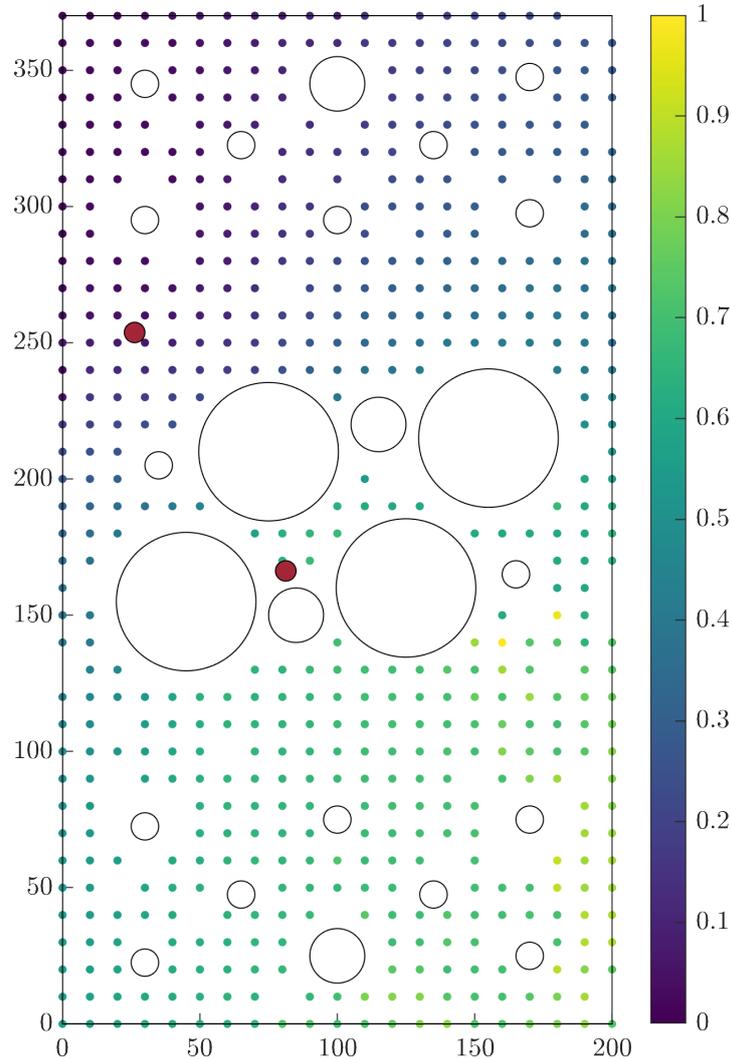


Figure 4.4:  $\Delta T$  mapping for sensor pair 3-5. Colour axis indicates the  $\Delta T$  value, with the red circles marking the sensor locations.

mean predictions will be in regions of data where the variance moves significantly away from the global average. In this scenario, the HGP can return mean functions that can better smooth through regions of local variation shift, instead of attempting to fit through individual points contained in the increased noise. For pairings that see an improved mean error, the corresponding sensors are positioned adjacently to one another, and so provide a poor coverage over much of the plate. With a poorer sensor coverage comes an increased variation on the onset time features (see Section 4.4, as well as Chapter 5 for further discussion), particularly as the event location moves further away from the receiving sensor, which is better smoothed through by the HGP.

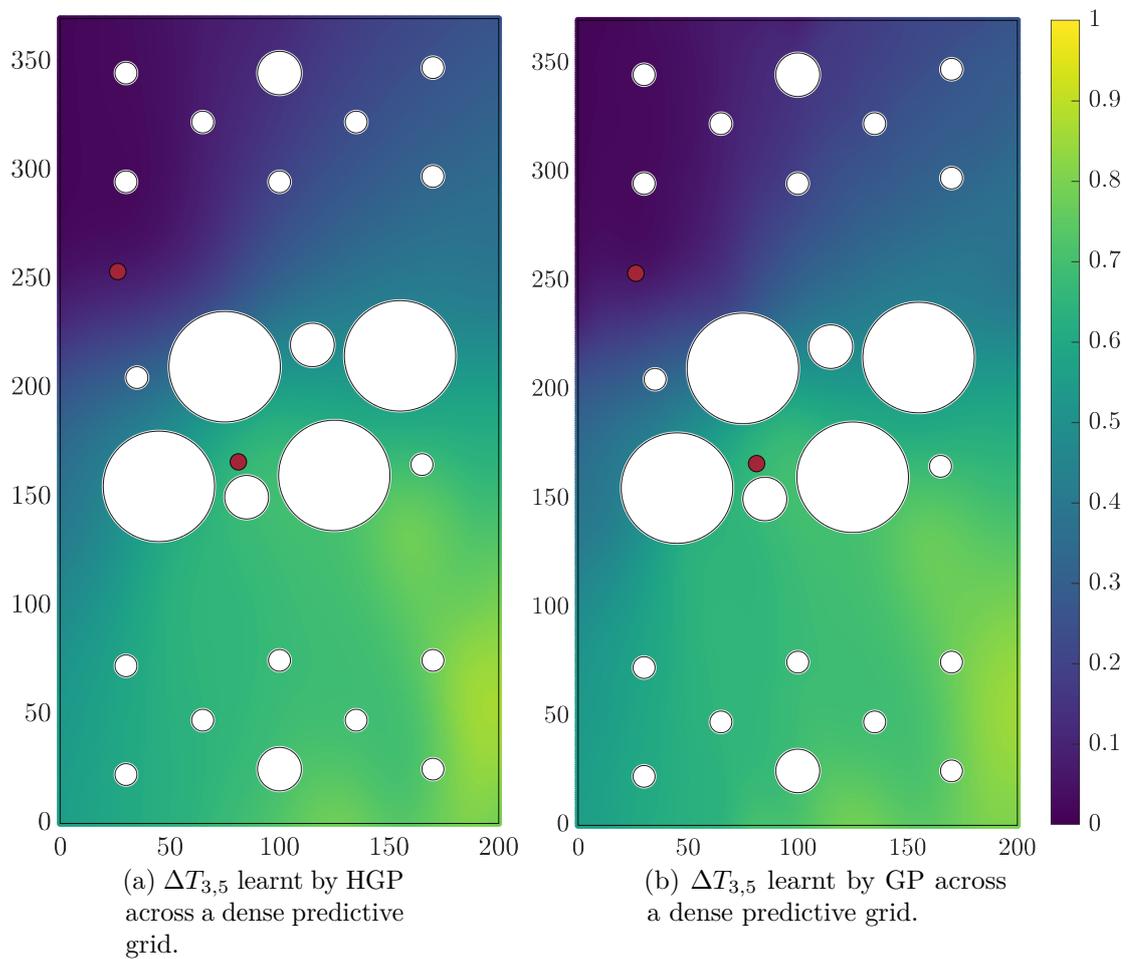


Figure 4.5:  $\Delta T$  mapping learnt for sensor pair 3 & 5 by both model forms.

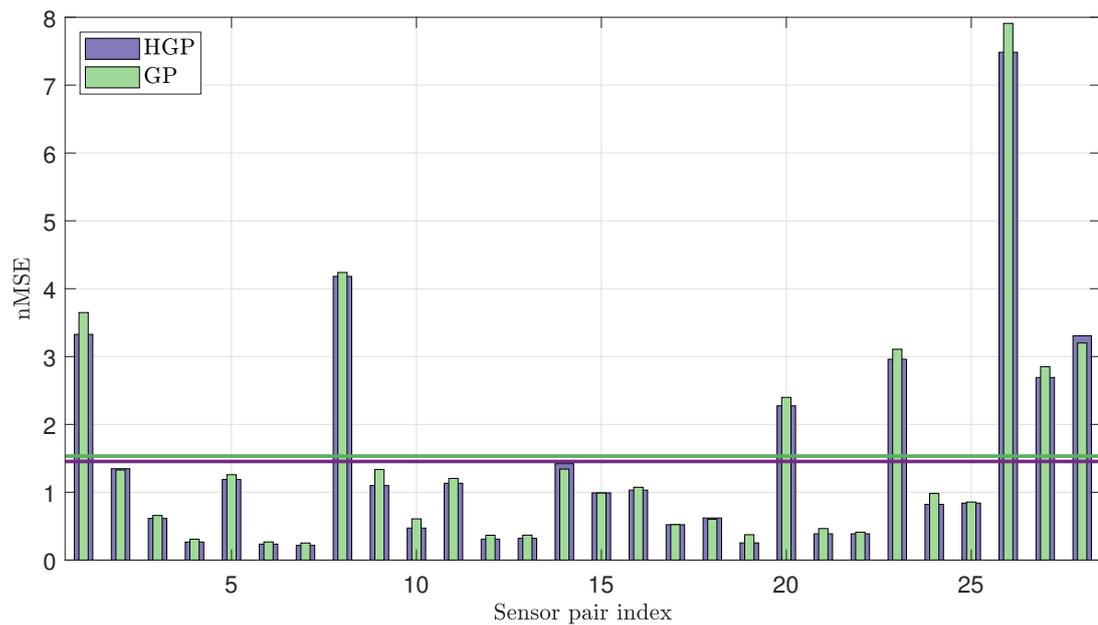


Figure 4.6: nMSE returned on test set across all sensor pairs for the GP and HGP. Mean score over all sensors given by horizontal line.

## 4.4 Uncertainty in $\Delta T$ maps with a heteroscedastic model

Having discussed the ability to learn the expected values of the  $\Delta T$  functions, let us turn our attention to capturing the variance on the  $\Delta T$  estimates. In the homoscedastic case, the variance outputs are solely a function of the distance between a given test point and the location of the training points. Where training points are uniformly sampled across the input space, with the same training locations used for every sensor pair, the predictive variance will remain constant across the domain. The only difference will be the overall scale of the noise (e.g.  $\sigma_n^2$ ), which will be dependent on the noise associated with that pair. Figure 4.7 plots the predictive variance of a sample of four sensor pairings, each offering coverage of different parts of the plate, where it is seen that the variances returned are largely uniform over the plate. The slight deviation in variance that does occur in parts of the plate is due to a lack of training measurements close to the boundaries, particularly in the centre of the plate.

The use of a global noise process will likely underestimate variation in some parts of the plate, whilst overestimating in others. To investigate this hypothesis, let us now examine the posterior variance learnt by the HGP. For the same four sensor pairings previously considered, Figure 4.8 plots this mapping. It can immediately be seen that there is significant difference in the predictive variance, with the increased expressiveness of equation (4.9) now allowing the confidence of a particular location to account for the relative positions of the sensors. For instance, taking the lower right plot of Figure 4.8 as an example, the top right area - particularly around the smaller hole - returns a high predictive variance. Inspecting the location of the two sensors, their coverage is very much limited to the lower left side of the plate. This positioning is particularly restrictive given the large number of holes in the centre of the plate, which will introduce increased complexity into the propagation path of the emissions. In the absence of sensor coverage at both sides of these geometrical features, estimating the onset times will become more challenging, resulting in an increased uncertainty that the HGP is able to capture. Considering the upper right (sensor pair 3 & 5) and the lower left plots of the figure (sensor pair 1 & 2), it can be seen that both of these sensor pairs too exhibit a region of high uncertainty away from the sensor locations. It is also evident that the HGP can account for the increased variation resulting from poor sensor coverage, with the maximum variance

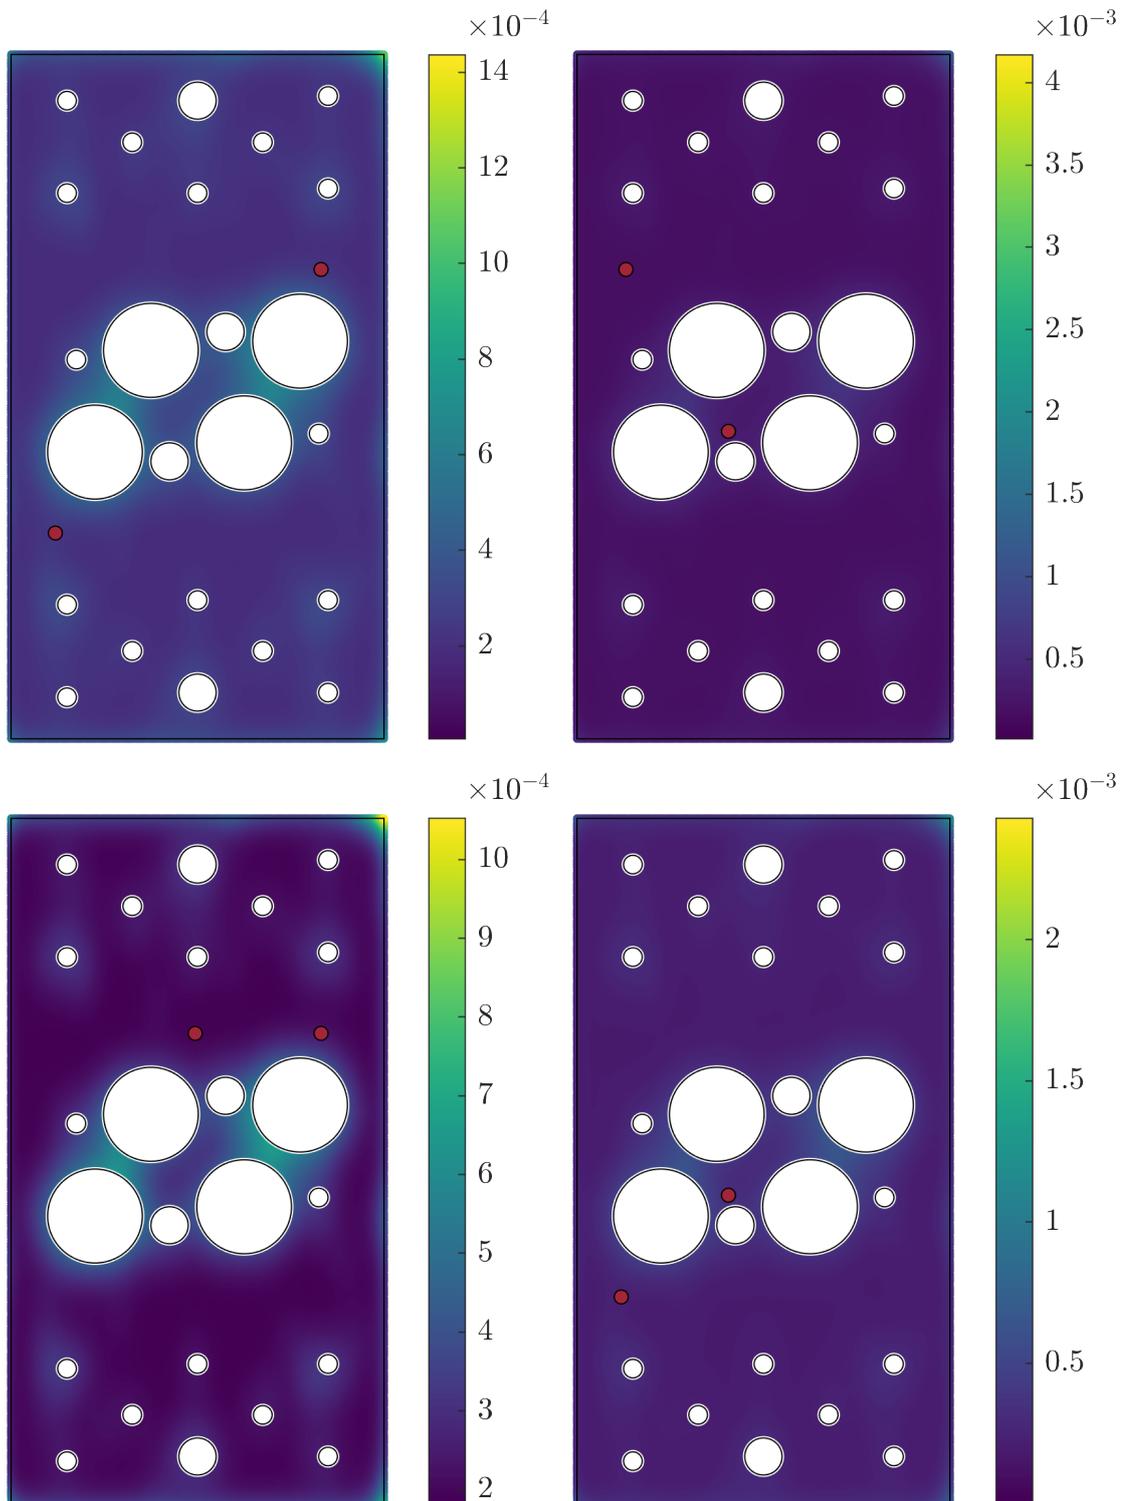


Figure 4.7: Predictive variance of  $\Delta T$  values learnt by the homoscedastic GP for four sensor pairings. Colour axis indicates the variance at a particular spatial location, with the red circles marking the sensor locations. The lower right of the plate here corresponds to the top right of Figure 4.2, having been rotated  $90^\circ$  clockwise.

in the mapping of sensor pair 1 & 2 an entire order of magnitude greater than what was predicted by the standard GP. The only sensor pair where the predictive variance remains fairly constant across much of the plate is that of pairing 1 & 8, where variance only begins to grow towards the lower left and upper right of the plate. Given that these two sensors are positioned such that their coverage extends over a significant portion of the plate, it would be expected to see a relatively fixed, low uncertainty.

To more formally quantify how well each model is able to capture distributions over the onset times, the MSLL can be considered. As the log loss is a probabilistic error metric, it also accounts for how well the variance on the test data has been captured by the predictions, thus serving as an ideal indicator for comparing the suitability of the two noise models. Figure 4.9 plots the MSLL obtained the test set used in the previous section for each of the sensors pairs. Given that the measure quantifies the loss obtained under the model, negative values indicate a better model fit.

For the standard GP, as the predictive density is Normally distributed, equation (3.23) is used to analytically evaluate the log loss. For the HGP, as there are no guarantees the predictive density is Gaussian, a Gauss-Hermite quadrature is used to calculate the MSLL. However, a good agreement was observed between the loss obtained through quadrature and that returned from approximating the HGP predictive density as a Gaussian, indicating the predictive densities were roughly Gaussian. Figure 4.9 shows that the HGP returns a much more favourable log loss, confirming that the HGP has been able to better capture the variation on the test set than the standard GP.

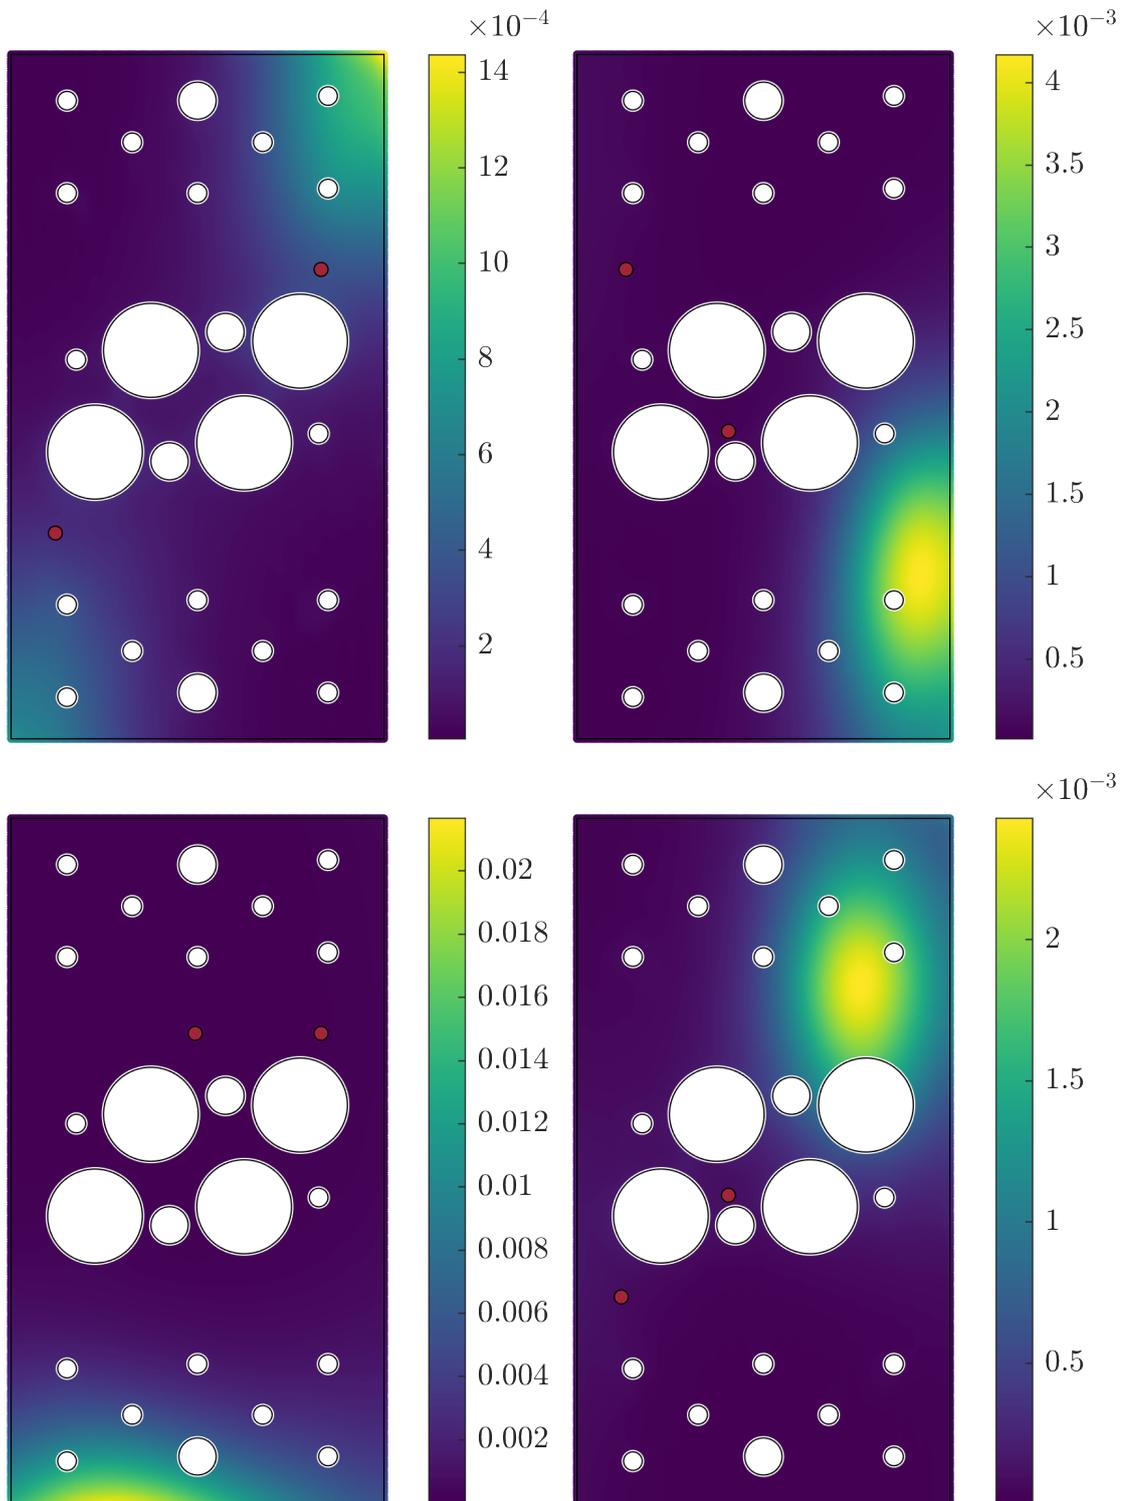


Figure 4.8: Predictive variance of  $\Delta T$  values learnt by the heteroscedastic GP for four sensor pairings. Colour axis indicates the variance at a particular spatial location, with the red circles marking the sensor locations. The lower right of the plate here corresponds to the top right of Figure 4.2, having been rotated  $90^\circ$  clockwise

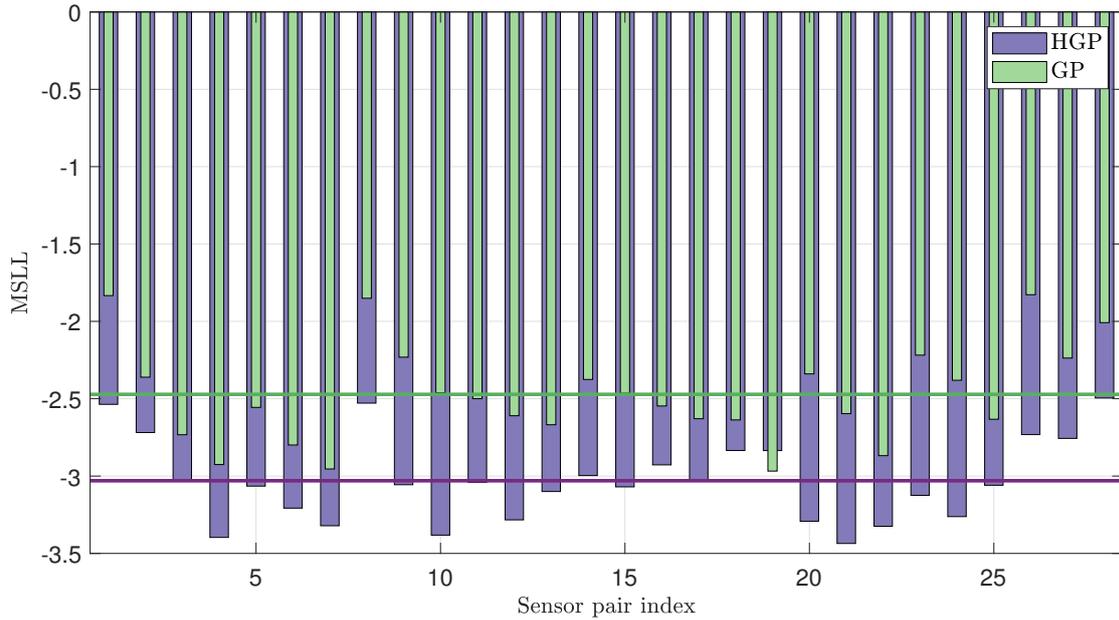


Figure 4.9: MSLL returned on test set across all sensor pairs for the GP and HGP. Mean score over all sensors given by horizontal line.

## 4.5 Localisation of AE events

In the previous section, it was seen that by transitioning to a heteroscedastic model, the predictive uncertainty on the predictions is now more reflective of the underlying behaviour of the  $\Delta T$  functions. When attempting to model some process, it is often important to attempt to suitably capture the associated uncertainty, and as such, this result in isolation is important. When considering the inverse problem, e.g. ‘given an observed  $\Delta T$  function value, where in the input space has it arisen from?’, the value in this improved noise model grows. This is because at a particular point, the assessment of the source location likelihood relies on the full predictive distribution. By accounting for the relative position between a given input and sensor pairing in the calculation of the variance, the resulting distribution is more informative than in the fixed noise case. Where sensor coverage is good, the corresponding source location likelihood can be scaled to reflect this increased confidence and vice versa. In this sense, the heteroscedastic model can essentially be viewed as a locally weighted model, depending both on a particular sensor-pair, and the relative position of the AE event to the sensors.

Before investigating the source location predictions made by the HGP, for the purpose of clarity, let us briefly restate the AE localisation problem as well as the proposed

framework developed in this thesis to address such a challenge. Generally, the task of AE localisation can be viewed as: having observed some signal or extracted feature, such as  $\Delta T$ s for multiple sensor pairs, predict the location at which the event originated from. Following the construction of a set of probabilistic mappings from grid coordinates  $(x, y)$  to  $\Delta T_j$ , following the previous chapter, it is possible to assess the likelihood of a new set of observed  $\Delta T$ s  $\mathbf{y}_\star = \{y_j\}_{j=1}^J$  that correspond to a new AE event. Mathematically, assessing this likelihood equates to,

$$\log p(y_j|\mathcal{D}, \mathbf{x}_\star, M_j) = \mathcal{N}(y_j|\mathbb{E}_q[y_j], \mathbb{V}_q[y_j]) \quad (4.12)$$

where  $E_q[y_j]$  and  $\mathbb{V}_q[y_j]$  are computed according to equations (4.8) and (4.9) with a training set  $\mathcal{D}$ . From a qualitative perspective, equation (4.12) provides the likelihood that an observed  $\Delta T$  from an unknown location originated from a candidate location  $\mathbf{x}_\star$  for a given sensor pair (model  $\mathcal{M}_j$ ). To predict the source location, the likelihood of the observed  $\Delta T$ s should then be maximised w.r.t. the input location  $\mathbf{x}_\star$ ,

$$\hat{\mathbf{x}}_\star = \arg \min_{\mathbf{x}_\star} \{\log p(\mathbf{y}_\star|\mathcal{D}, \mathbf{x}_\star)\} \quad (4.13)$$

where the most likely source location is given by  $\hat{\mathbf{x}}_\star$ . Notice that the model (sensor pair)  $M_j$  has been removed from the likelihood, having marginalised over all possible models,

$$p(\mathbf{y}_\star|\mathcal{D}, \mathbf{x}_\star) = \sum_{j=1}^J p(y_j|\mathcal{D}, \mathbf{x}_\star, \mathcal{M}_j) \quad (4.14)$$

To visualise the localisation process, a test event not included within the training set is selected. The likelihoods of the corresponding  $\Delta T$  features are then evaluated at a dense grid of test points across the structure, with the results plotted for a selection of sensor pairs in Figure 4.10. The true source location is marked with a red cross. Like with the homoscedastic GP, it can clearly be seen that on each of the sensor pair maps, there is a region of high likelihood that intersects the true source location. These regions are then analogous to the hyperbolae of potential source locations that would be returned from traditional TOA analysis.

One comment to make here is on the appearance on the likelihood surface for sensor

pairing 1 & 2, which appears as more of a cluster of high likelihood instead of some smooth curve as forms for the other three pairings. The reasoning for this behaviour is that the positioning of sensors 1 & 2 results in both sensors lying along similar propagation paths, causing changes in  $\Delta T$  as one moves over the plate to be much more gradual. Consequently, identification of possible emission locations for this pair of sensors becomes much more challenging.

To compare these single sensor pair likelihoods with those obtained under a homoscedastic noise, Figure 4.11 plots the equivalent result for the standard GP.

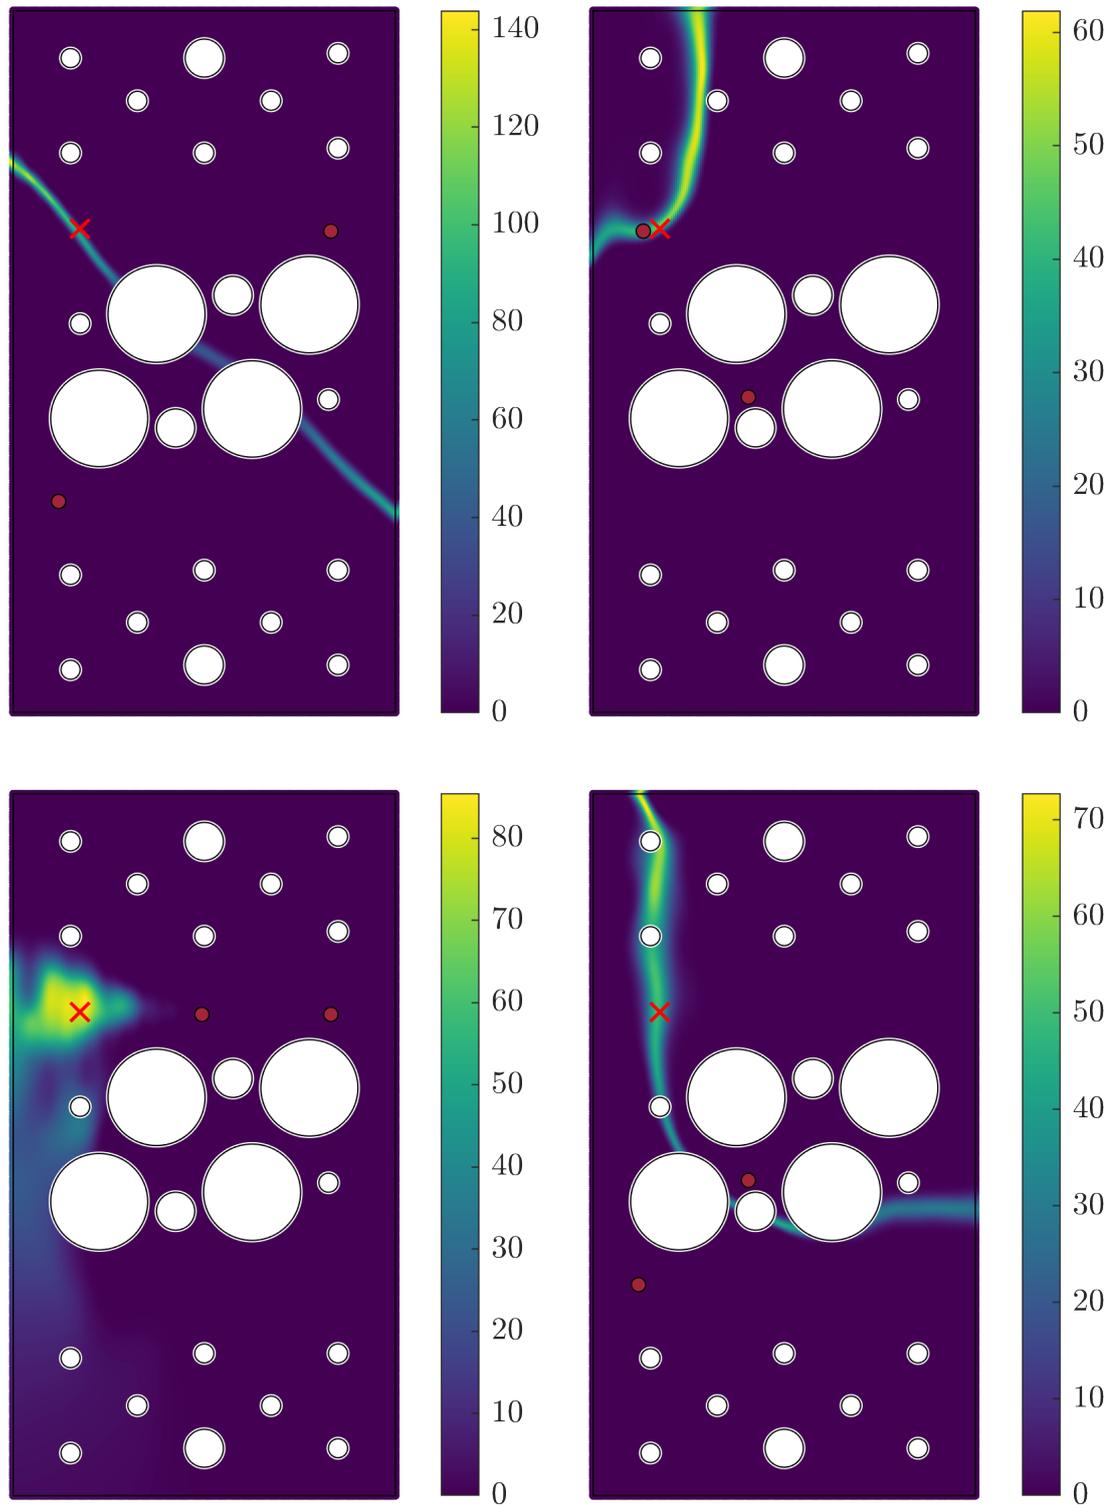


Figure 4.10: Conditional emission location likelihood of test event returned from heteroscedastic model for four sensor pairings. The true location is highlighted by the red cross, with the red circles marking the sensor locations.

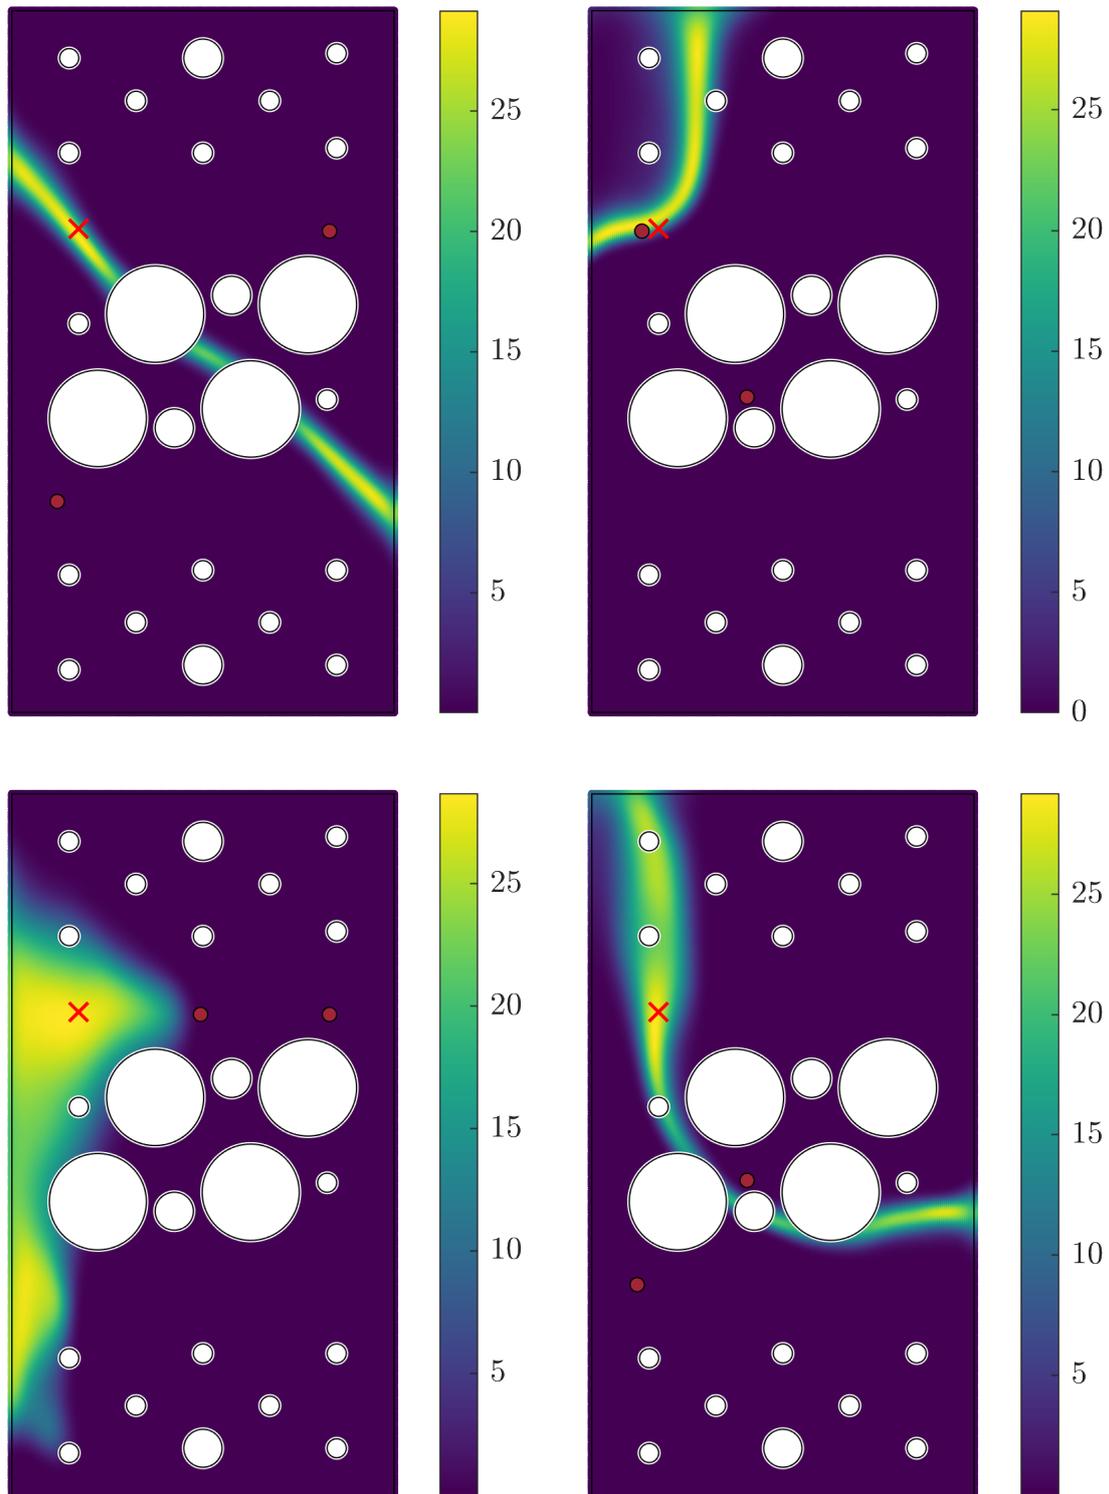


Figure 4.11: Conditional emission location likelihood of test event returned from standard model for four sensor pairings. The true location is highlighted by the red cross, with the red circles marking the sensor locations.

It can be seen that in the heteroscedastic model, the regions of high likelihood are narrower and more distinct, with likelihoods along parts of the contours higher than that returned by the standard GP model. To explore the reasoning for this behaviour, let us explicitly restate the equation for assessing the likelihood,

$$\log p(y_{obs,j} | \mathcal{D}, \mathbf{x}_*, M_j) = -\frac{1}{2} \log \mathbb{V}(\mathbf{f}_{*,j}) - \frac{(y_{obs,j} - \mathbb{E}(\mathbf{f}_{*,j}))^2}{2\mathbb{V}(\mathbf{f}_{*,j})} - \frac{1}{2} \log 2\pi. \quad (4.15)$$

Examining each term, the  $\frac{1}{2} \log 2\pi$  term can be ignored as it will remain constant across both models. There are then two remaining components that determine the likelihood of a test point; the squared difference between the observed  $\Delta T$  of an event and that predicted by the model, scaled by the predictive variance  $\frac{(y_{obs,j} - \mathbb{E}(\mathbf{f}_{*,j}))^2}{2\mathbb{V}(\mathbf{f}_{*,j})}$  and a logged variance term,  $\frac{1}{2} \log \mathbb{V}(\mathbf{f}_{*,j})$ . With the predictive mean returned by both models largely similar over much of the plate, changes in the likelihood between the GP and HGP will mainly be driven by differences in the predictive variance, which both scales the likelihood arising from the differences in observed and predicted  $\Delta T$ , as well as in the form of an additive term. As an example, let us revisit the predictive variance for sensor pair 1 & 8, for which the log values are plotted in Figure 4.12 to magnify differences between the two models. The figure shows that in the upper left of the plate, a lower predictive variance is returned by the HGP than the GP, indicating that in this area of the input space, locations are well covered by the sensors. Where the variance is lowered, confidence in regions of high likelihood are increased, resulting in the greater relative conditional likelihoods that are seen along the probable source location contour for sensors 1 & 8 in Figure 4.10 compared to Figure 4.11. Conversely, as may be somewhat obvious, where confidence is reduced then lowers the source location likelihoods, such as in the middle of the plate.

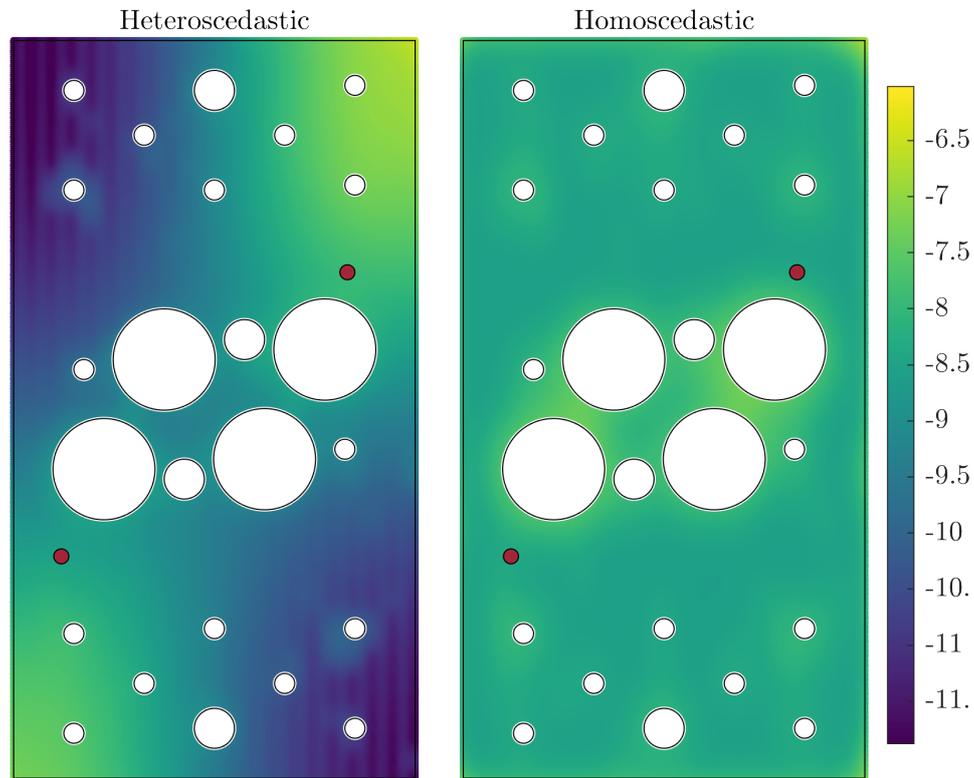


Figure 4.12: Predictive variance for sensors 1 & 8 for the left: heteroscedastic and right: homoscedastic GP. Colour axis indicates the variance, with the red circles marking the sensor locations.

The final step in the localisation framework is to combine the information from all the sensors pairs, with the result plotted in Figure 4.13. The predicted emission location from each model is indicated by a black cross. Comparing the plots, the prediction made by the heteroscedastic model is closer to the true origin of the emission than the standard GP. It is also evident that there is a significant increase in the confidence of the heteroscedastic prediction, indicated both by an increased likelihood at the predicted origin, as well as the region of high likelihood that encompasses the most likely estimate being less diffused. Whilst a reduction in location prediction error is an important result, an improved confidence is equally significant; should an operator use these maps to guide maintenance and repair action, the enhanced noise model that increases the confidence in the predictions is very useful, particularly where these maps would feed into a higher level of a health monitoring framework.

Considering a second test event, but this time taken to be from the lower part of the plate, here it can be seen, again, that the HGP performs much better, returning a more accurate location prediction which is also more certain.

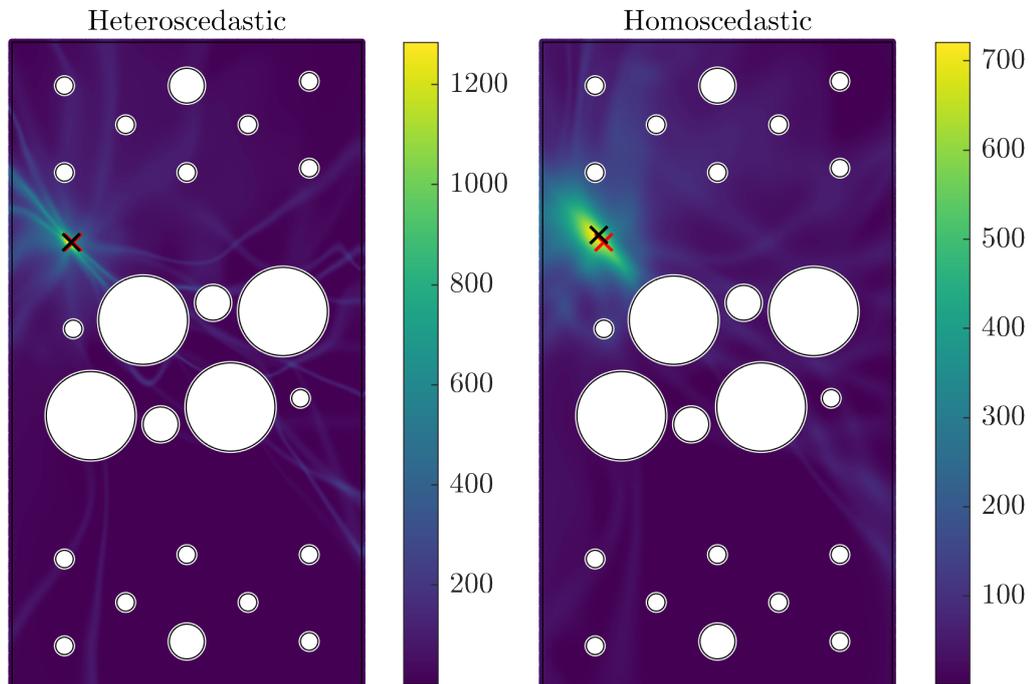


Figure 4.13: Marginal location likelihood map for test AE event. Red cross is true location, black is predicted.

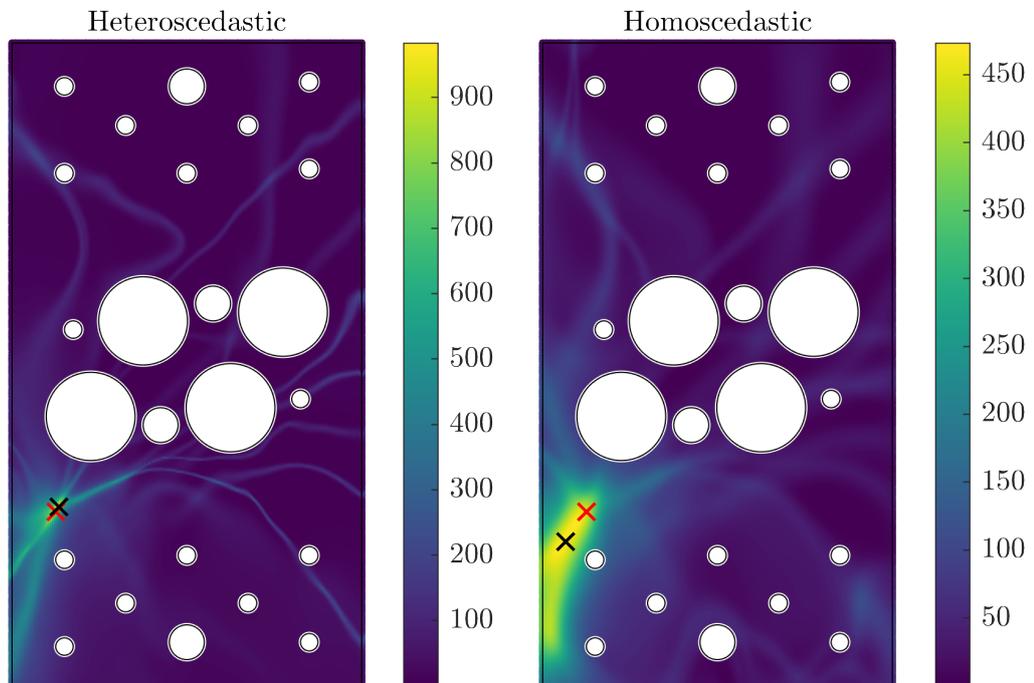


Figure 4.14: Marginal location likelihood map for a second test AE event. Red cross is true location, black is predicted.

As a final comparison, the maximum likelihood estimate of 100 test events is computed by both models. These predictions are then plotted in Figure 4.15, where it can be seen that the mean estimates returned by the HGP are, overall, closer to the true locations than the estimates of the GP model. This improved accuracy can be quantified by the RMSE of the predictions, with the HGP obtaining a score **3.52mm**, compared to the almost doubled value returned by the GP of **6.56mm**.

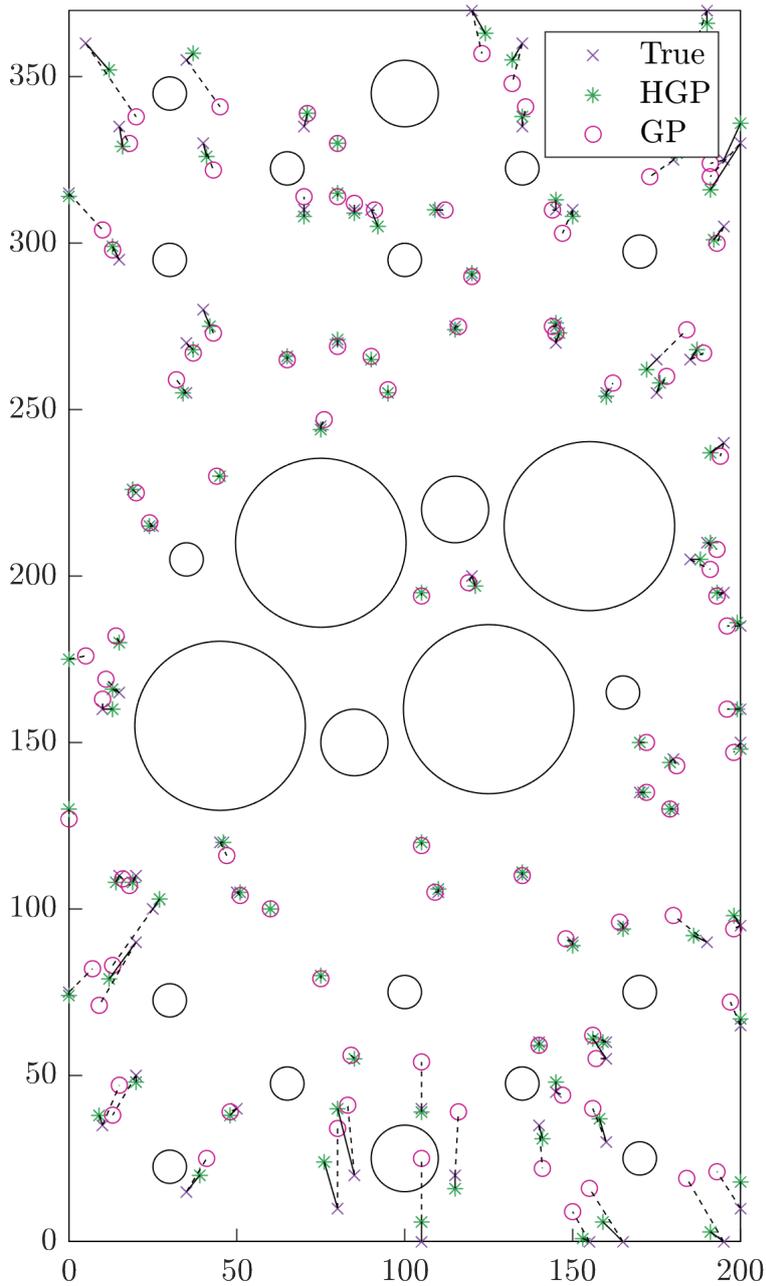


Figure 4.15: Comparison of predicted source location for 100 random points from the test set for both models. RMSE of HGP: 3.52, RMSE of GP: 6.56.

## 4.6 Conclusions

This chapter has presented the use of a heteroscedastic Gaussian process model for learning  $\Delta T$  maps. Previously, it was assumed that the uncertainty present on the onset times may be treated as draws from a constant variance distribution that is learnt globally, imposing a fixed noise process on the  $\Delta T$  values over the whole of the structure. In reality, when investigating the spatial behaviour of the onset times, the variation present on these features fluctuates according to how well the particular sensor pairing covers an area of interest. For instance, for locations whose line of sight to a sensor pairing is heavily obstructed, the local variation of onset times in that area increases, observed by considering the residual between the target value and that predicted by a standard (homoscedastic noise) Gaussian process. To capture this behaviour, this chapter proposes switching the constant variance assumption with some function that depends on the location in the input space, allowing for an input-dependent noise process. To learn this function, a second Gaussian process prior is placed over the noise, and is learnt jointly with the GP learning the predictive mean within a variational framework. The predictive equations, therefore, now comprise of two latent functions. A first that learns the expected (mean) value of the  $\Delta T$  times, and a second, that governs the variance of the additive noise process at a particular input location.

To investigate localisation performance using the heteroscedastic noise model, the  $\Delta T$  maps are learnt by the HGP for the complex plate, allowing for comparison between the two GP models. The first observation made is that the mean estimates of the onset times learnt by the HGP are similar to those of the standard, except in the case where the variance of the onset times significantly deviates from the global average, with the HGP having the capability to better smooth through this increased noise. Considering the predictive variance, placing a second GP prior over the noise variance allows for a more expressive noise model. For the HGP, regions of high uncertainty now emerge depending upon how well the sensor pair covers a particular area of the structure. Conversely, where some region of the input space is well covered by a sensor, confidence in predictions in this area are now increased. Propagating these enhanced uncertainty predictions into the localisation predictions, it is seen that tighter high likelihood regions are generally returned about the true source location, with maximum likelihood estimates that return a lower mean error. This behaviour can be attributed to the fact that this local noise model effectively

allows predictions to be weighted according to the quality of sensor coverage, allowing the predictions obtained from sensor pairs that are positioned closer to a prediction point to be more heavily weighted, and likewise, the poorer the coverage, the lower the weighting. Such a functionality will be extremely important where one may have partial or uneven sensor coverage of a structure.

In the previous two chapters, the localisation framework has been viewed from a purely-data driven perspective. In the following chapter, a methodology will be presented that combines both a data-driven learner and engineering knowledge as a means for reducing the dependency on training data in the learning phase of the  $\Delta T$  maps.

# PHYSICAL CONSTRAINTS FOR PHYSICS-INFORMED ONSET TIME MAPPING

### Highlights:

- *The burden of data collection for  $\Delta T$  maps is alleviated through the use of physics-informed Gaussian processes.*
- *Physical insight is included into the  $\Delta T$  learning process by constraining the covariance function to known boundary conditions.*
- *The constrained model improves predictive accuracy where the availability of training data are reduced.*

As seen throughout the thesis, acoustic emission time of arrival mapping plays a crucial role in the ability to localise acoustic emission as a means of damage localisation. Although the methodologies developed so far achieve a good performance in demanding environments, the need to collect a dense set of artificial acoustic emission measurements across the whole structure results in a lengthy data acquisition process, somewhat hindering the practicality of the method. In this chapter, it is considered how physical constraints may be incorporated into the learning process to alleviate this problem. In the approach, a Gaussian process model is constrained to the phys-

ical domain such that information relating to the geometry and boundary conditions of the structure are embedded directly into the learning process, returning a model that guarantees that any predictions made satisfy physically-consistent behaviour at the boundary. A number of scenarios that arise when training measurement acquisition is limited are investigated, including where training data are sparse, and also of limited coverage over the structure of interest. It is shown that the approach significantly reduces the burden of data collection, where it is seen that incorporation of boundary condition knowledge significantly improves predictive accuracy as training observations are reduced, particularly when training measurements are not available across all parts of the structure.

## 5.1 Introduction

For the proposed AE localisation framework, it is required that artificial acoustic emission measurements are collected over the structure of interest. Given that collecting such measurements is a very manual and laborious task, constructing these maps can become very costly from both a financial and temporal perspective, severely hampering the method. This cost is particularly pertinent as the size of the structures considered grows. It is, therefore, critical that the size of the training set is minimised. The challenge here is then building models that allow reduced training set sizes but can offer a good localisation performance in complex structures. As discussed throughout the thesis, geometrical features will disrupt the propagation path of the ultrasonic surface waves, inducing sharp discontinuities in the AE arrival time functions. Such behaviour will require a standard data-driven learner to have a dense training set in regions containing these structural complexities so that the onset time behaviour can be suitably captured, conflicting with the desire to reduce the amount of data that needs to be acquired. An additional challenge is faced from the action of retrofitting a monitoring system, where gaining access to the complete coverage of all parts of a structure is often challenging, particularly at locations that contain joints and interfaces. As such, it is important to obtain good model generalisation in areas where measurements may be more sparse, which proves difficult when solely reliant on a data-driven learner. These problems are not unique to the application of acoustic emission localisation however, and instead arise consistently wherever black-box learners are deployed. Whilst their flexible nature enable a high level of performance in the presence of an abundance of data,

when presented with sparse training sets or forced to extrapolate, there is often no guarantee on how the predictions will behave.

This chapter proposing a means of addressing the above limitations of  $\Delta T$  mapping through incorporating physical knowledge into the learning process. Methodologies that consider such an approach can broadly be grouped under the term physics-informed machine learning, where physical insight is embedded into data-driven algorithms - for a general overview of this field, see [126]. Also referred to as grey-box modelling [127], the objective when adopting such a model architecture is to combine the expressive power of machine learning tools with physically derived laws or constraints. Predictions that are drawn from these models are then guaranteed to be representative of some underlying physical laws that govern the dynamics of the system under consideration.

One particular way to embed physically-derived insight is by constraining the learning algorithm with physical constraints such that subsequent predictions comply with these assumptions. In the context of Gaussian process regression, there are numerous ways constraints can be applied, as extensively discussed in [128]. For example, if one has knowledge of the shape of the latent function, then monotonicity or convexity constraints can be applied [129–133]. General bounds such as nonnegativity can also be imposed on the GP prior [134], as well as constrained to satisfy linear operators [135]. Where more specific insight is available, for example, the underlying equation of motion, derivation of an exact autocovariance is possible [136]. Constraints also exist in the form of vector-output Gaussian processes, where relevant physical relationships between multivariate targets can be embedded into the cross-covariance terms. Under a multivariate output framework, the inclusion of both ordinary differential equations [137, 138] and boundary conditions have been explored [139].

The focus of the work presented here will be on the application of constrained Gaussian processes for learning difference-in-time of arrival/ $\Delta T$  acoustic emission maps, with application and subsequent analysis on the complex plate structure introduced in Chapter 2. The nature of the constraints that are considered are those of physical boundary conditions, embedded into the model by firstly rewriting the covariance of the GP prior as a Laplacian eigenfunction expansion. Given that the eigenfunctions are unique to a user-specified domain, the GP can then be naturally constrained to some boundary conditions along a physical domain. This results in a model that retains the flexibility of a traditional machine learner, whilst adhering to known physical conditions that exist at boundary locations. To the authors' best

knowledge, up until now, all previous works that consider a data-driven approach to AE localisation have been purely black-box in nature. This chapter demonstrates that by ensuring that the model is constrained to known physical laws, the process of generating  $\Delta T$  maps is made feasible, improving the predictive performance in cases of sparse/few training observations and where training measurements only partially cover the full structure spatially, and therefore, only partial coverage of the input space for the data-driven learner. General discussions surrounding where one may implement the constrained Gaussian process are also considered.

## 5.2 Constrained Gaussian processes

Although Gaussian processes present a powerful tool, as with all machine learners, they are still black-box in the sense that their performance is entirely reliant on the data that the model is trained on. In cases where sufficient training data are scarce, then the resulting model may struggle to adequately learn the underlying behaviour of the features. As introduced above, one way to circumvent such issues is to incorporate physical insight into the machine learner. In this chapter, focus is directed on embedding boundary condition knowledge into a Gaussian process prior through constraining the covariance function. This presents one view of a constrained process. From the perspective of the nature of the physics considered, the advantage here is that the level of insight into the governing mechanistic laws can be relatively shallow; knowledge of boundary conditions are generally easier to come by than an exact governing differential equation. Additionally, through directly constraining the form of the prior, the approach does not rely on the addition of artificial observations at the boundary. In fact, the method employed here is a sparse approximation, and so is computationally cheaper than the standard implementation of a Gaussian process, both in terms of complexity and storage demands [140].

To constrain a Gaussian process in this manner, it is first necessary to make use of the following covariance function approximation [140]:

$$k(\mathbf{x}, \mathbf{x}') \approx \sum_j^m S(\sqrt{\lambda_j}) \phi_j(\mathbf{x}) \phi_j(\mathbf{x}'). \quad (5.1)$$

Under this representation, the covariance function is defined as a basis function expansion across  $m$  Laplacian eigenfunctions  $\phi$  of a user-selected domain, projected

onto the spectral density  $S$  of the covariance that has been evaluated in a point-wise manner at the corresponding Laplacian eigenvalues  $\lambda$ . To calculate the Laplacian eigenpairs, one is required to solve an eigenvalue problem of the form,

$$-\nabla^2 \phi_j(\mathbf{x}) = \lambda_j^2 \phi_j(\mathbf{x}), \quad x \in \Omega, \quad (5.2)$$

where  $\nabla^2$  is the Laplacian operator,

$$\nabla = \left( \frac{\partial}{\partial x^1}, \dots, \frac{\partial}{\partial x^D} \right) \quad (5.3)$$

with  $D$  the total dimensionality of the vector  $\mathbf{x}$ , and  $\Omega$  represents the domain of interest. As for any differential equation, its solution may be sought given boundary conditions described generally by,

$$\Psi[\phi_j(\mathbf{x})] = H(\mathbf{x}), \quad x \in \delta\Omega, \quad (5.4)$$

where  $\Psi$  denotes some operator and  $H$  is an arbitrary function that maps  $\mathbf{x}$  to a known value which exists on the boundary of a domain  $\delta\Omega$ . The solutions for  $\lambda$  and  $\phi$  are, therefore, bound to the chosen domain  $\Omega$ , and, consequently, are unique to the boundary conditions specified in equation (5.4). Upon substitution into equation (5.1), each draw from the prior is then guaranteed to abide by these constraints, returning an expression for the covariance that is dependent on the boundary conditions of the feature space, and embedding physical boundary conditions into the model.

To visually demonstrate, Figure 5.1 shows the plot of a sample from both an unconstrained (standard) and constrained Gaussian Process from a function defined across the range  $[0, 1]$ , with boundary condition  $y = 0$  at  $(0, 1)$ . Both GPs are conditioned on the same input points. As can be seen in the figure, only the constrained GP is able to correctly predict the boundary conditions, with the standard GP over or undershooting the true values.

To specify a suitable kernel spectral density, it is possible to employ Bochner's theorem which states that the covariance of a stationary function can be represented by the Fourier transform of a positive, finite measure [105]. If this measure has a corresponding density  $S$ , then the spectral density and the covariance are Fourier

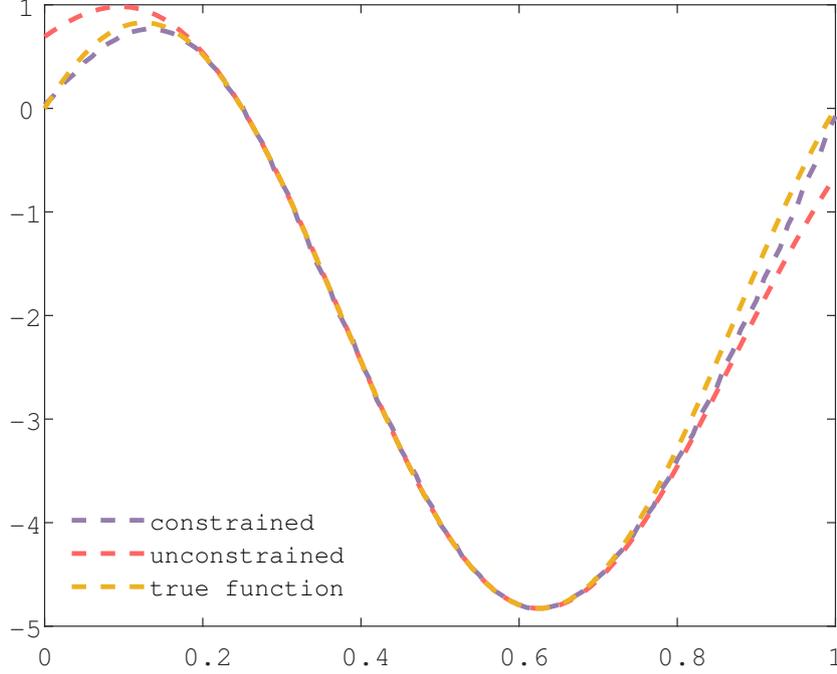


Figure 5.1: GP draw from an unconstrained and constrained GP across domain defined on  $[0, 100]$ , with boundary conditions  $y = 0$  at  $x = 0, 1$ .

duals of one another [141]. For example, the spectral density of the Matérn 3/2 kernel takes the following functional form,

$$S(\omega) = \sigma_f^2 \frac{4\pi(2\nu)^\nu \Gamma(\nu + 1)}{l^{2\nu} \Gamma(\nu)} \left( \frac{2\nu}{l^2} + \omega^2 \right)^{-(\nu+1)}, \quad (5.5)$$

where  $\nu = 3/2$  and  $\Gamma$  denotes the Gamma function. For a test point  $\mathbf{x}_*$ , the predictive posterior is then defined as:

$$\mathbb{E}[y_*] = \mathbf{\Phi}_* (\mathbf{\Phi}' \mathbf{\Phi} + \sigma_n^2 \mathbf{\Lambda}^{-1})^{-1} \mathbf{\Phi}' \mathbf{y} \quad (5.6)$$

$$\mathbb{V}[y_*] = \sigma_n^2 \mathbf{\Phi}_* (\mathbf{\Phi}' \mathbf{\Phi} + \sigma_n^2 \mathbf{\Lambda}^{-1})^{-1} \mathbf{\Phi}'_*, \quad (5.7)$$

with,

$$\mathbf{\Phi} = (\phi_1(X), \phi_2(X), \phi_3(X), \dots, \phi_m(X)), \quad (5.8)$$

$$\mathbf{\Lambda} = \text{diag}(S(\lambda_1), S(\lambda_2), S(\lambda_3), \dots, S(\lambda_m)), \quad (5.9)$$

The mean function in the predictive equations is now set to 0. The reasoning for this is not due to the mean being uninteresting; that's quite the opposite in a grey-box context. However, specifying a physics-based basis function often requires a significant level of knowledge regarding the dynamics of the system being modelled. The intention here is to present a general method that can be applied with known boundary condition information, which, in most cases, is readily available. For interesting examples of applying a physics-based mean function, the reader is referred to the following works [142, 143].

Finally, as shown in equation (5.5), there are a number of hyperparameters to learn. Again, a type-II maximum likelihood approach as detailed in Chapter 3 can be followed, where the modified NLML for the constrained GP is written,

$$\begin{aligned} \boldsymbol{\theta} = \arg \min_{\boldsymbol{\theta}} \{ & \frac{1}{2}(N - m) \log \sigma_n^2 + \frac{1}{2} \sum_{j=1}^m \Lambda_{j,j} + \frac{1}{2} \log |\sigma_n^2 \mathbf{\Lambda}^{-1} + \mathbf{\Phi}^T \mathbf{\Phi}| + \frac{n}{2} \log(2\pi) \\ & + \frac{1}{2\sigma_n^2} (\mathbf{y}^T \mathbf{y} - \mathbf{y}^T \mathbf{\Phi} (\sigma_n^2 \mathbf{\Lambda}^{-1} + \mathbf{\Phi}^T \mathbf{\Phi})^{-1} \mathbf{\Phi} \mathbf{y}) \}. \end{aligned} \quad (5.10)$$

### 5.3 Physical constraints for acoustic emission time of flight mapping

To implement the constrained GP, it is first required that the Laplacian eigenfunctions, as shown in equations (5.2) and (5.4), are solved. These expressions are unique for a given physical domain, and therefore enable boundary condition knowledge to be directly encoded into the model. From the perspective of a dynamicist, this process can be seen as analogous to finding the wavenumbers and approximate mode shapes of the structure. For simple geometries, it is possible to arrive at closed form solutions for the Laplacian eigenvalues. However, due to the geometrical complexity

of the plate, in this work the eigendecomposition is computed numerically. This is calculated by approximating the operator with a finite difference equation that is solved alongside the boundary conditions, with each boundary condition giving an equation that can be solved simultaneously with equation (5.2). To implement this numerical approximation, the Laplacian operator is first converted into its discrete counterpart by transforming the domain into a *grid mask*,  $u$ , which exists as a binary matrix where ones denote locations inside the domain, whilst zeros indicate the opposite. A discrete representation of the Laplacian can then be formed as a stencil matrix by applying a finite difference approximation of the Laplace operation on  $u$ ,

$$-\nabla^2 u(i, j) \approx \frac{1}{h^2}(-4u_{i,j} + u_{i-1,j} + u_{i,j-1} + u_{i+1,j} + u_{i,j+1}), \quad (5.11)$$

where  $(i, j)$  index the rows and columns of the grid mask, and  $h$  represents the step size between adjacent nodes within the grid. Equation (5.11) can then be manipulated to reflect known information solution information in the form of boundaries conditions where  $(i, j)$  lie on the boundary of the grid mask. For example, in the case of Dirichlet conditions (process value is specified), the solution can be fixed at the boundary positions. For the onset time functions, the associated boundary condition is that of a first order spatial derivative equal to zero (Neumann boundary conditions). In the model, this enforces that at the boundary of the structure the acoustic emission waves propagate in, the derivative of the arrival times is equal to zero. Equation (5.4) can therefore be rewritten as,

$$d\phi_j(\mathbf{x}) = 0 \quad x \in \delta\Omega. \quad (5.12)$$

To solve for first order derivative boundary conditions is slightly more involved than Dirichlet conditions and requires the use of a first order approximation at the boundary location. However, before incorporating this boundary type, a series of “ghost point” must be added to every location directly adjacent to a boundary. These ghosts points are purely artificial, and exist to help implement the finite difference approximation at endpoints of the domain. As an example, take the arbitrary 2D domain shown in Figure 5.2a, where each circle represents a node in the grid mask. Consideration of the boundary on the left hand side of the domain along the  $j$  direction requires the addition of ghost points at  $j = 0$ , as shown in Figure 5.2b.

Incorporation of a first order derivative at  $j = 1$  can then be achieved by applying a

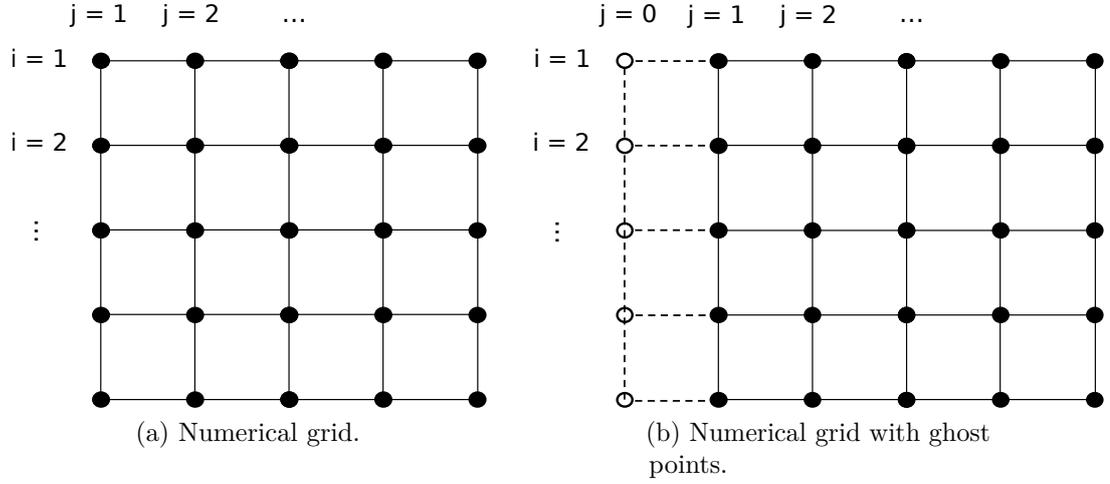


Figure 5.2: Numerical grid for arbitrary domain to depict how ghost points are introduced to include derivative boundary conditions.

backward difference approximation

$$u'_j = \frac{u_{i,j} - u_{i,j-1}}{h}, \quad (5.13)$$

which for zero derivative boundary conditions, can be simplified to,

$$u_{i,j} = u_{i,j-1}. \quad (5.14)$$

At this point, it is worth noting that although a higher order approximation would be obtained through the use of a central difference method, the use of a backward (and forward, as discussed shortly) difference scheme ensures that the resultant stencil matrix is symmetric [144], and therefore the corresponding eigenvectors are real. Returning to equation (5.11), at  $j = 1$ , the ghost points that appear at  $u_{i,0}$  can be removed by substituting in the above expression, yielding,

$$-\nabla^2 u(i, 1) \approx \frac{1}{h^2} (-3u_{i,1} + u_{i-1,1} + u_{i,2}), \quad (5.15)$$

and incorporating the boundaries into the stencil matrix. Where a ghost point lies at an index of +1 to the boundary, such as the right hand side of the domain in Figure 5.2a, then it is necessary for a forward difference approximation to be applied

$$u'_j = \frac{u_{i,j+1} - u_{i,j}}{h}, \quad (5.16)$$

replacing equation (5.14) with,

$$u_{i,j+1} = u_{i,j} \quad (5.17)$$

The process detailed above is then identical for any given direction, providing the finite difference equation applied is consistent with the target direction. For example, in the  $i$  direction, equations (5.14) and (5.17) become

$$u_{i,j} = u_{i-1,j} \quad (5.18)$$

$$u_{i,j} = u_{i+1,j}. \quad (5.19)$$

Iterating through each element of the grid mask, a stencil matrix can be computed that corresponds to the negative Laplacian of  $\Omega$ . A visualisation of the corresponding sparse stencil matrix can be seen in Figure 5.3. The leading  $m$  eigenvalues and eigenfunctions can then be calculated through a chosen numerical solver. For the work conducted in this chapter, following [145] for large-scale spatial mapping problems, 256 eigenbases are used. The first 16 of these basis functions that incorporate all physical boundaries present on the plate are shown in Figure 5.4. Note that although only homogenous first order derivative boundaries are considered, it is possible to include other forms of boundary condition through proper treatment of equation (5.11) at boundary locations.

## 5.4 Results and discussion

The feasibility of acoustic emission localisation for large and complex systems is severely hampered by the need to collect artificial AE events at locations on a dense grid across the structure. To explore how the constrained GP may help to alleviate this burden, in this section, results are shown that investigate scenarios where the number and location of training points are limited. To mimic the likely availability

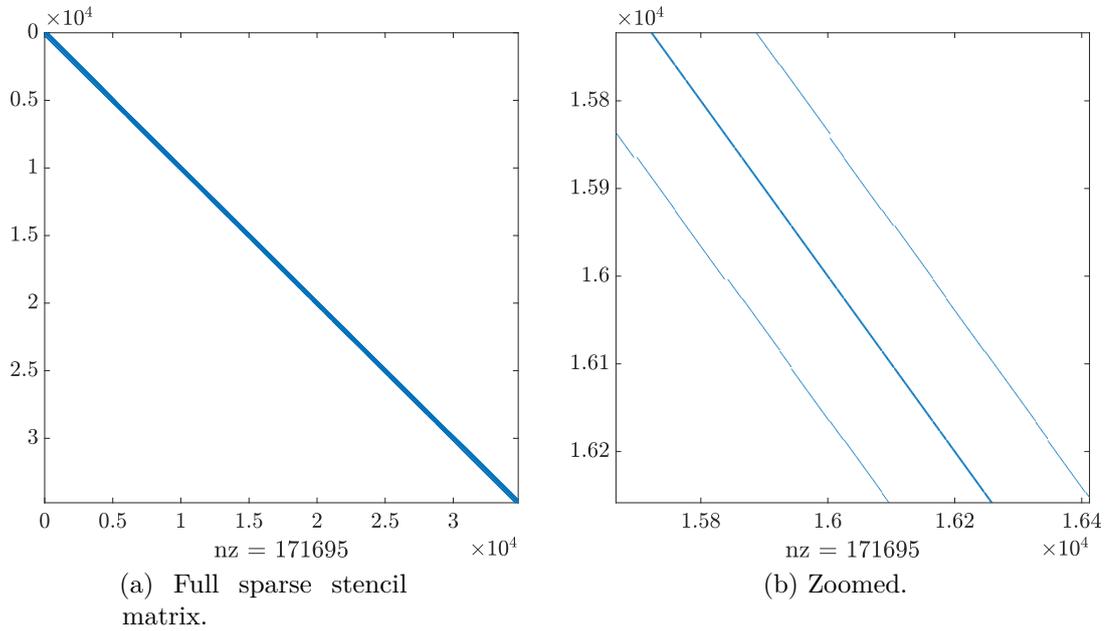


Figure 5.3: Sparse stencil matrix for the plate structure.

of training data from a measurement campaign on a real structure, the case where training measurements are available across the structure but with limited grid density is considered first. For most structures, however, particularly those with multiple components, it is unlikely that access to the whole structure would be available to establish a training dataset (e.g. between closely space components). The second scenario investigated therefore limits the training dataset to a single part of the plate.

The investigation will compare the performance of the standard and constrained GPs. Naturally, the availability of measurements themselves from the boundary for GP conditioning will affect the performance of both methods. This is quantified by explicitly considering additional measurements at the available boundary for both scenarios. In the first case, where measurements are available across the full structure, it is expected that the standard GP will outperform the constrained model when training grids are dense - the constrained GP is after all an approximation. It is also expected that the benefit of the constrained model will be seen where training data are sparse and when data are only available from limited locations across the structure (as will likely be the case in operation).

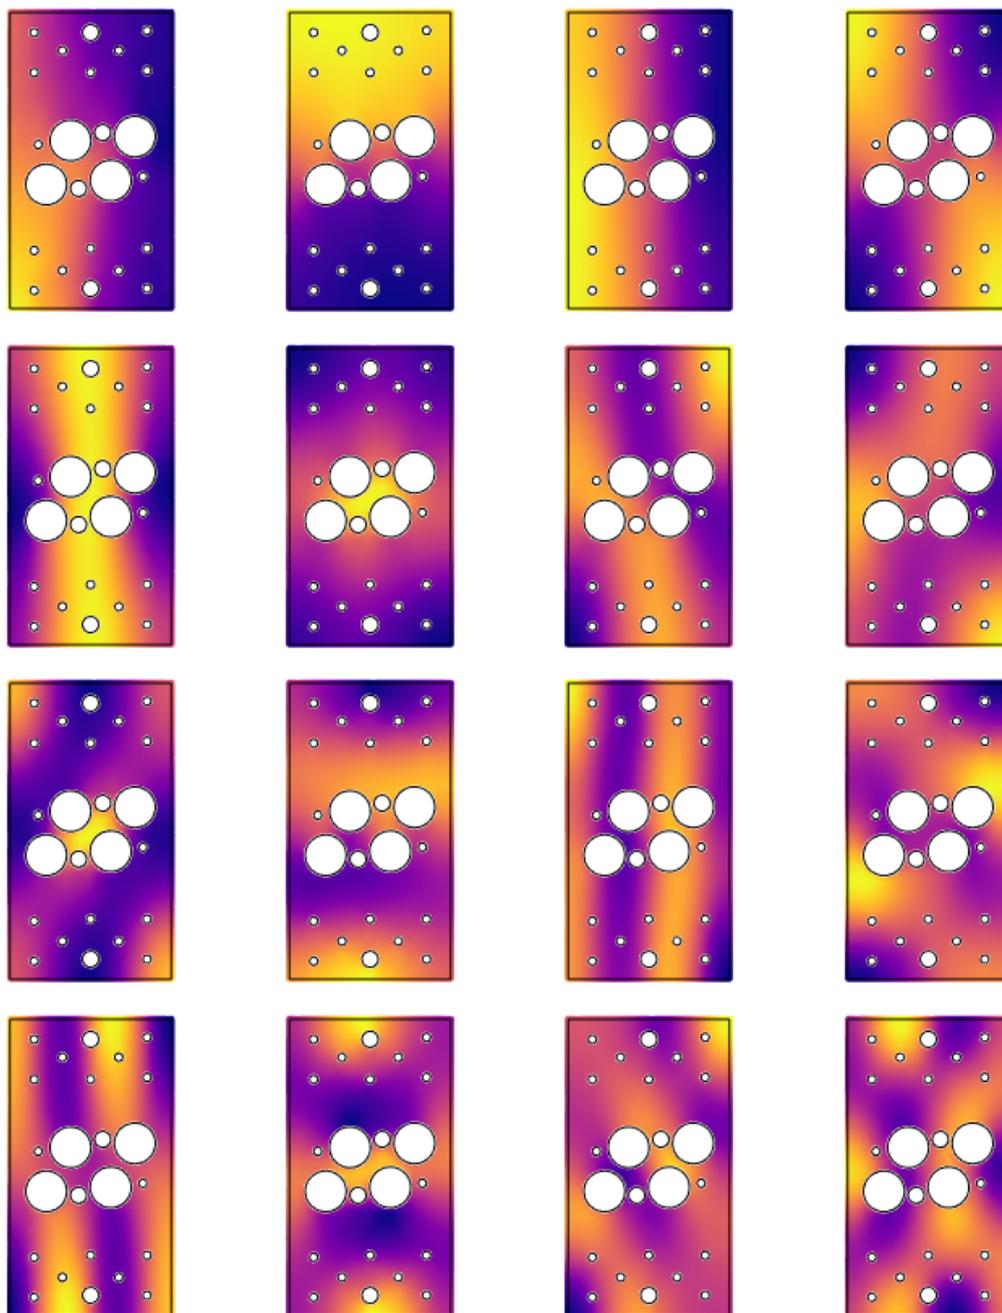


Figure 5.4: First 16 Laplacian eigenfunctions (left to right) of the plate.

### 5.4.1 Sparse training data available across the whole structure

When undergoing an AE data collection campaign, it is often not practical to collect a dense grid of training observations, particularly for larger structures or when setting up multiple monitoring regions. To consider how the constrained GP performs where the training data availability is reduced, a number of training sets containing varying numbers of observations are formed, with measurements available across the full spatial limit of the structure. For an individual set, the spacing between training points is uniform (excluding where the holes are located), with the value that the spacing is fixed at varying across the sets. The total range of training sets considered is outlined in Table 5.1.

The nMSE returned on the test set for each training set, averaged over all available sensor pairs, is plotted on Figure 5.5 respectively for both the constrained and standard GP.

Spacing between training points (mm)	10	15	20	25	30	40	50	60	70
Total number of observations	600	302	160	107	71	50	32	14	11

Table 5.1: Spacing between observations and corresponding total number of datum points for each training set used.

The figure demonstrates that as the number of training points reduces to the larger grid spacings, the constrained GP begins to offer a slightly greater accuracy. At the lower end of the grid spacings where the training set is denser, it can be seen that the standard GP is marginally favourable, which as introduced above, is expected. Where training data coverage and availability is good, sufficient boundary condition insight can be obtained from training measurements. That is not to say, however, that one should always seek to retain the full covariance where data are easily assessable. When computing predictions with the constrained GP, computational complexity reduces from  $\mathcal{O}(N^3)$  to  $\mathcal{O}(Nm^2)$ , with storage requirements moving from  $\mathcal{O}(N^2)$  to  $\mathcal{O}(Nm)$ . A benefit here, therefore, is that when the number of training points exceed several thousand, the use of the constrained GP presents a practical solution without turning to more complex computing techniques such as parallelisation and/or the use of graphical processing units.

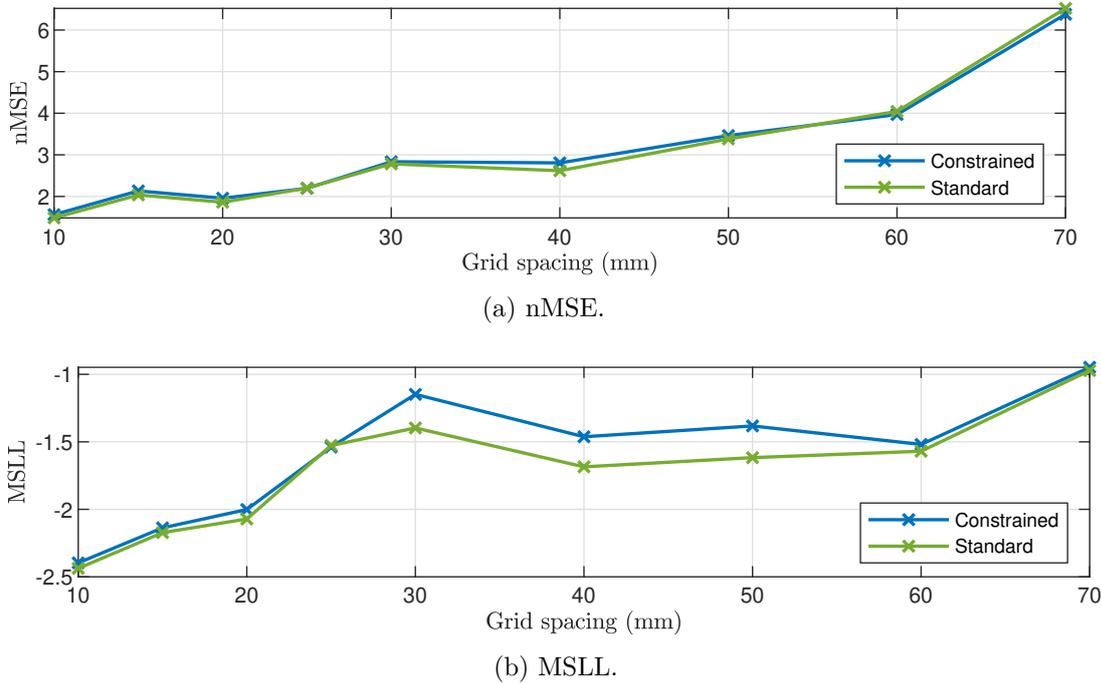


Figure 5.5: nMSE and MSLL of the predictions made by the constrained and standard Gaussian process as the number of training points are reduced.

In the above example, it was assumed that boundary measurements were available inline with the overall training grid density. For example, at a spacing of 20mm, boundary measurements were available every 20mm. It may arise, however, that one would wish to gather more insight in regions on or around boundaries. For instance, when mapping an acoustic emission wavefield, it is likely that one would want more insight around boundaries of the domain, where sharp discontinuities will be introduced into the propagation pattern. The first option available in this scenario would be to collect more measurements at the boundary location, which is considered by repeating the above experiment, but ensure that boundary measurements are available every 10mm<sup>1</sup>. This results in a scenario where one may take a fairly sparse grid of training measurements, but adopt a more fine resolution at boundary locations. Clearly, it is also possible to combine more training measurements with physical constraints. Figure 5.6 plots the results for both of these cases.

It can be seen that both models obtain similar error scores, with some improvement

<sup>1</sup>For the outside boundaries, measurements were taken directly on the boundary. For the inside boundaries, measurements were not available directly on the boundary. Therefore, inside boundary measurements were defined as data points directly adjacent to a given hole in line with a 10mm grid spacing.

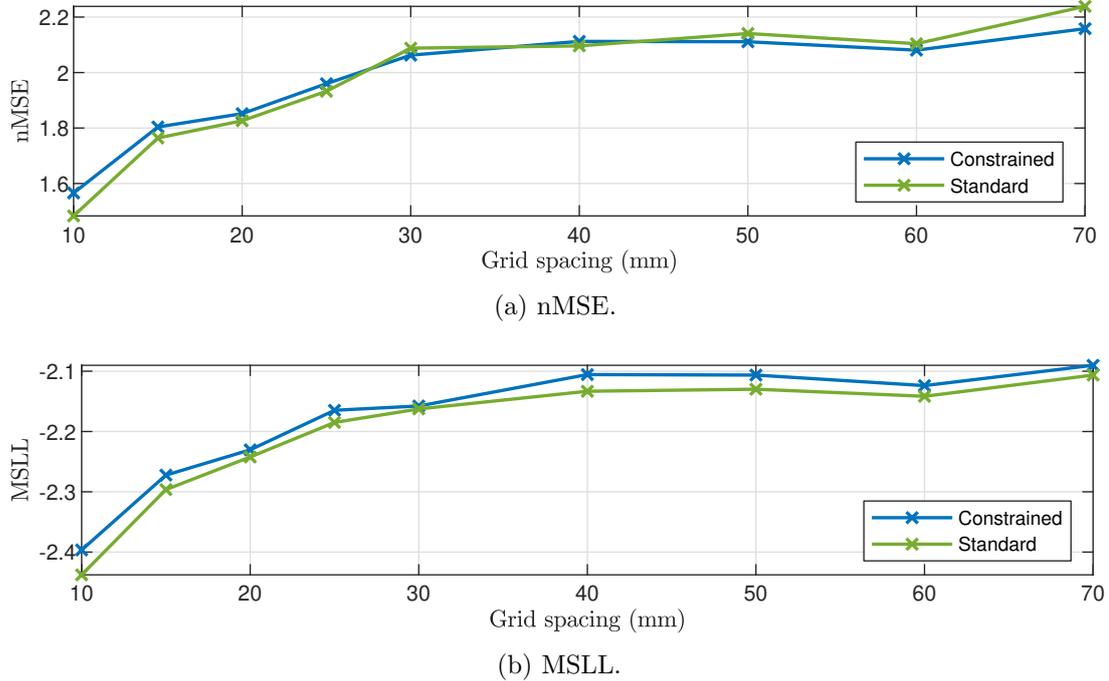


Figure 5.6: nMSE and MSLL of the predictions made by the constrained and standard Gaussian process as training points are reduced whilst retaining full boundary measurements.

returned by the constraints where training data becomes more sparse (note the difference in the range of  $y$  axis values between Figure 5.5 and Figure 5.6). In the case that a measurement campaign has been specifically conducted with both extra measurements at the boundary and a dense training grid, then the benefit of using the constrained GP from the perspective of mean predictive accuracy is negligible. There are, however, a number of disadvantages to simply adding more training points at the boundary. As the structures we wish to represent become more complex, particularly when in two or three dimensions, the number of datum points required to sufficiently capture a continuous domain will quickly grow. Given the cubic and quadratic scaling of complexity and storage respectively for standard GPs, prohibitive computational demands can quickly be reached, and so for large or complex structures that will demand many boundary measurements, the constrained GP will be a more feasible solution. An additional limitation when collecting data is that many engineering structures simply prevent acquisition at boundaries, for example, at joints and connections that physically obscure generating an artificial signal at that location. We, therefore, now examine a second scenario, where one has access to no boundary measurements, repeating the procedure in the preceding paragraph, but removing all boundary locations from the training set. The predictive performance on the

test set for both forms of model is plotted in Figure 5.7. The figure demonstrates an improved performance from the constrained GP as training data become fewer. The significance here however, is that performance gain of the constrained GP at larger grid spacings is higher than that obtained in Figure 5.5 and 5.6, with the error returned by both models for the denser training sets comparable. As the standard GP now has no knowledge of boundary conditions, it is forced to predict at boundary locations with little information if the training grid is not dense. In the case of the constrained GP, the constraints provide the kernel function additional information to the training measurements, incorporating physically relevant structure into the covariance that can then be used when predicting at locations on and adjacent to boundary locations.

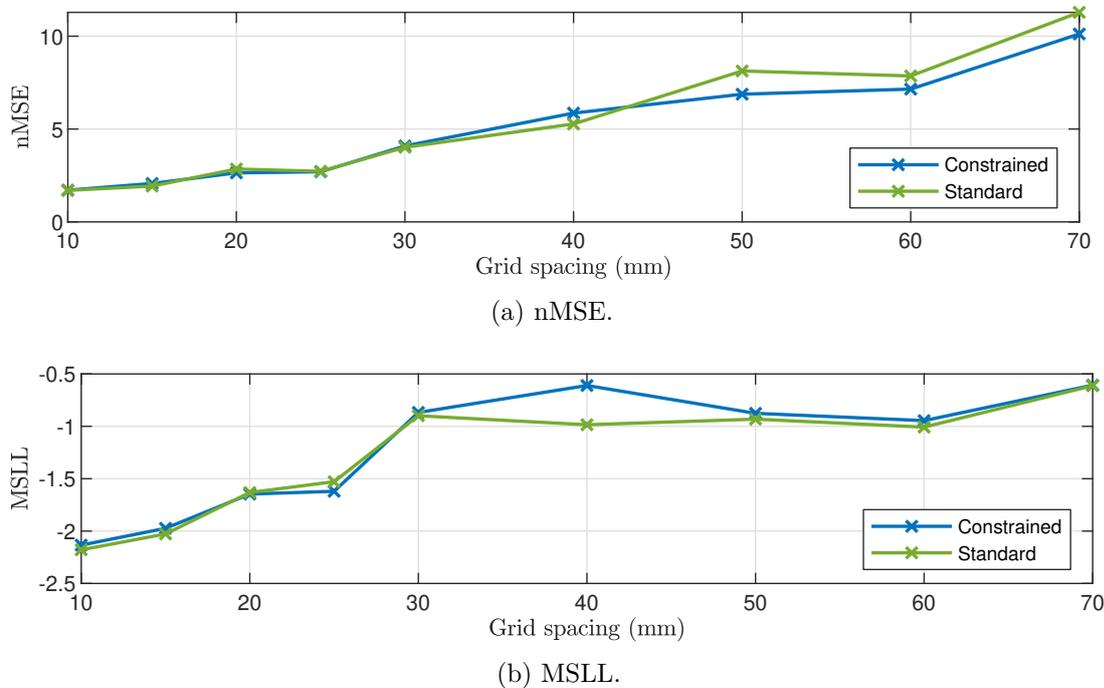


Figure 5.7: nMSE and MSLL of the predictions made by the constrained and standard Gaussian process as training points are reduced with no boundary measurements.

Overall, what can be deduced from the results in this section is that when few training points are available, in the absence of a dense collection of boundary measurements, the inclusion of boundary constraints improve the predictive capabilities of the model. Whilst there is no guarantee that an improved predictive accuracy will be obtained where training data is in more of an abundance, particularly at the boundaries, the use of the constrained GP may still be a consideration when one is limited by the computational demands of computing predictions in closed form through the

standard GP implementation.

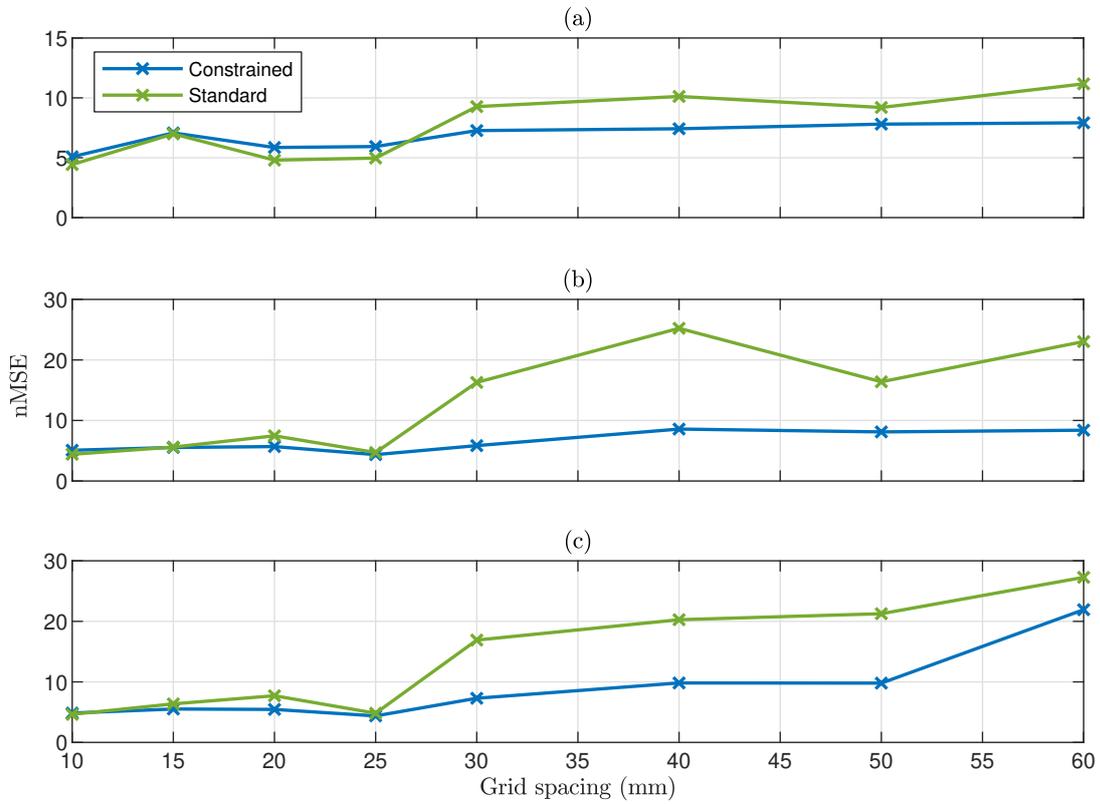
### 5.4.2 Training data from partial structure coverage

When undergoing an SHM data collection campaign, as previously discussed, it is often not possible to collect data across the entire input space. In a spatial mapping context, this limitation generally arises through being unable to collect data that fully covers the entire structure of interest. For example, when acquiring the artificial AE events that are used to learn the  $\Delta T$  mapping, it may not be possible to gain full access of the structure, particularly when a health monitoring system is being retrofitted. Such a scenario may arise if trying to collect data from the drive train of a wind turbine gearbox, for example, where the assembly of various interlocking gears and shafts will obscure access to many of the individual components that one may be interested in developing an AE mapping for. The fuselage of an aircraft is another example where full access is prevented without disassembly, particularly in the areas where the root of the wings are mounted.

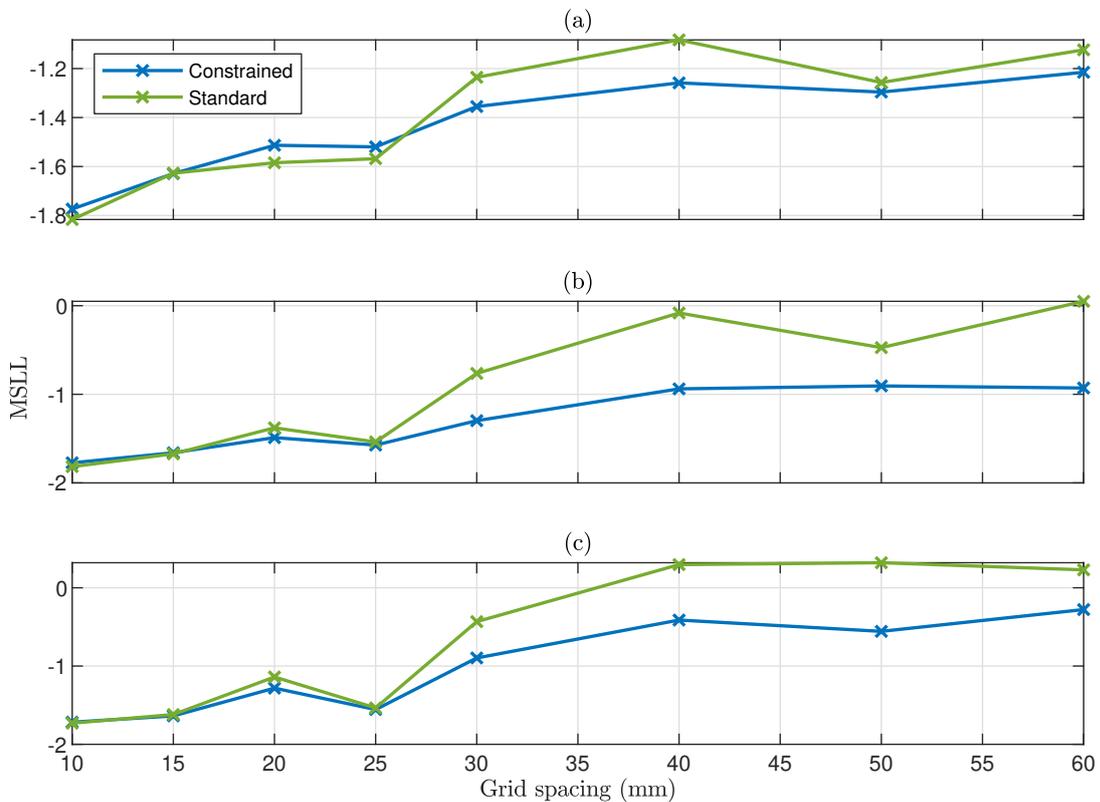
To explore how the constrained GP can mitigate against a lack of training data coverage, the data points used to train the models are first restricted to the middle section of the plate. The three training conditions with respect to the inclusion of boundary points in the training set are again considered; these are full boundary coverage, partial boundary coverage (i.e. inline with the overall training grid density), and finally no boundary location measurements. Predictive performance on the test set for each of these three conditions for both the constrained and standard GP is shown in Figure 5.8.

Across all three cases, the constrained GP offers either an improved or comparable predictive performance. For the wider grid spacings, particularly where partial or no boundary measurements are available, a large performance gain is obtained by constraining the GP, regularly exceeding an nMSE reduction of greater than 10%.

To further investigate the performance of the constrained GP, a single grid spacing of 30mm where no boundary measurements are available is now considered (case c) in Figure 5.8). At a particular grid spacing, as there are a total of 28 sensor pair combinations, there are thus 28 total feature maps to learn. Figure 5.9 plots the nMSE obtained on the test set for each individual sensor pairing. For reference, Table 5.2 lists the sensor pair numbers with the corresponding index used here.



(a) nMSE for training sets with a) full boundary location measurements b) partial boundary measurements c) no boundary measurements.



(b) MSL for training sets with a) full boundary location measurements b) partial boundary measurements c) no boundary measurements.

Figure 5.8: Predictive performance with restricted training coverage.

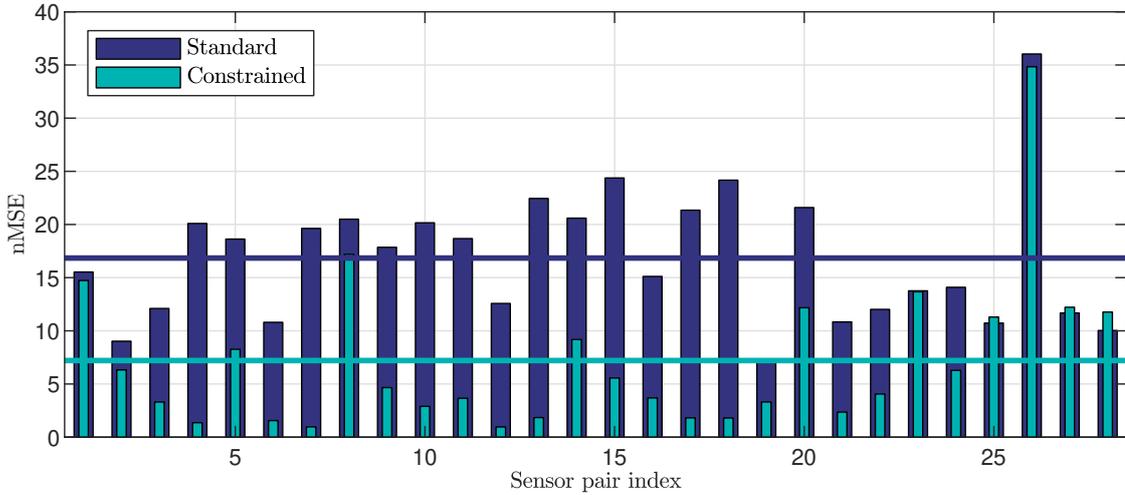


Figure 5.9: nMSE returned by each sensor pair model on a training set with limited coverage (case c). Horizontal lines correspond to the average of the standard and constrained nMSE across all sensor pairs.

Sensor pair index	Sensors	Sensor pair index	Sensors
1	1-2	15	3-5
2	1-3	16	3-6
3	1-4	17	3-7
4	1-5	18	3-8
5	1-6	19	4-5
6	1-7	20	4-6
7	1-8	21	4-7
8	2-3	22	4-8
9	2-4	23	5-6
10	2-5	24	5-7
11	2-6	25	5-8
12	2-7	26	6-7
13	2-8	27	6-8
14	3-4	28	7-8

Table 5.2: Sensor pair index and corresponding individual sensors.

As an initial observation, it can be seen that for the majority of the sensor pairs, the predictions returned by the constrained GP are more accurate. This can be quantified formally by considering the averaged nMSE across all of the sensor pairs for both models, where the constrained GP yields an averaged error of 7.30 in comparison to 16.91 from the standard GP. Considering examples of the constrained GP performing significantly better, sensor pairing 15, which corresponds to sensors 3 & 5, displays a large difference in error between the two models. Figure 5.10 maps the mean predictions of the test set for each of the GP models. Comparing both of the plots, it

can be seen that most of the variation between the  $\Delta T$  predictions exists in lower and upper part of the plate, with the boundary constraints constricting the magnitude of the  $\Delta T$  value predicted by the model.

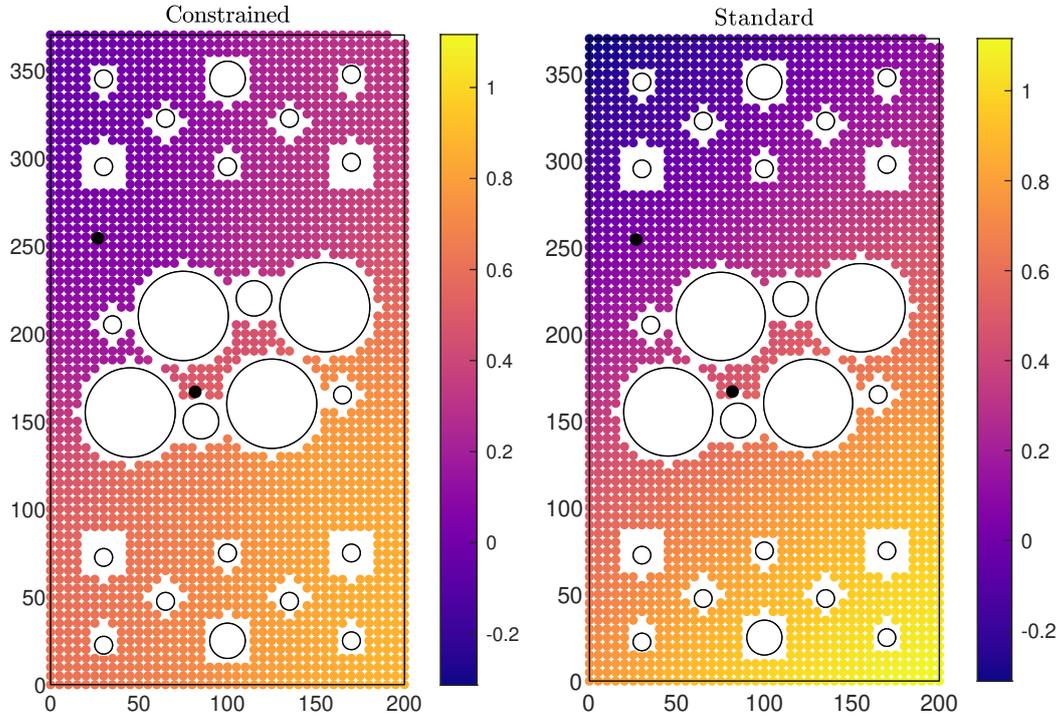


Figure 5.10:  $\Delta T$  predictive mean of both GP models on the test set for sensor pair 3 & 5. Sensor locations are highlighted by the black circles.

Continuing with the analysis of sensor pair 3 & 5, Figure 5.11 shows the predictive variance returned by both the standard and constrained GP on the test set. What the results show here is that the constraints embedded into the GP reduce the uncertainty of the prediction made across the testing set in comparison to the standard GP, particularly in the upper and lower region of the plate. It is interesting to note that in the centre of the plate, the constraints do slightly increase the variance on predictions made in this area. A possible explanation here is that the standard implementation of the GP does not sufficiently represent the uncertainty in this area, following the predictive variance being solely a function of distance to training locations. What this measure does not provide is any adjustment for how the presence of boundary conditions may alter the behaviour of the AE features. For example, in this middle section of the plate, the onset times will not vary smoothly due to the discontinuities imposed by the boundaries. Coupled with a lack of training data in this zone, with the awareness of the physical domain embedded into the model, the predictive variance returned by the constrained GP is more representative of the

underlying uncertainty. A further consideration to make here is the influence that the type of the boundary condition would have on the uncertainty.

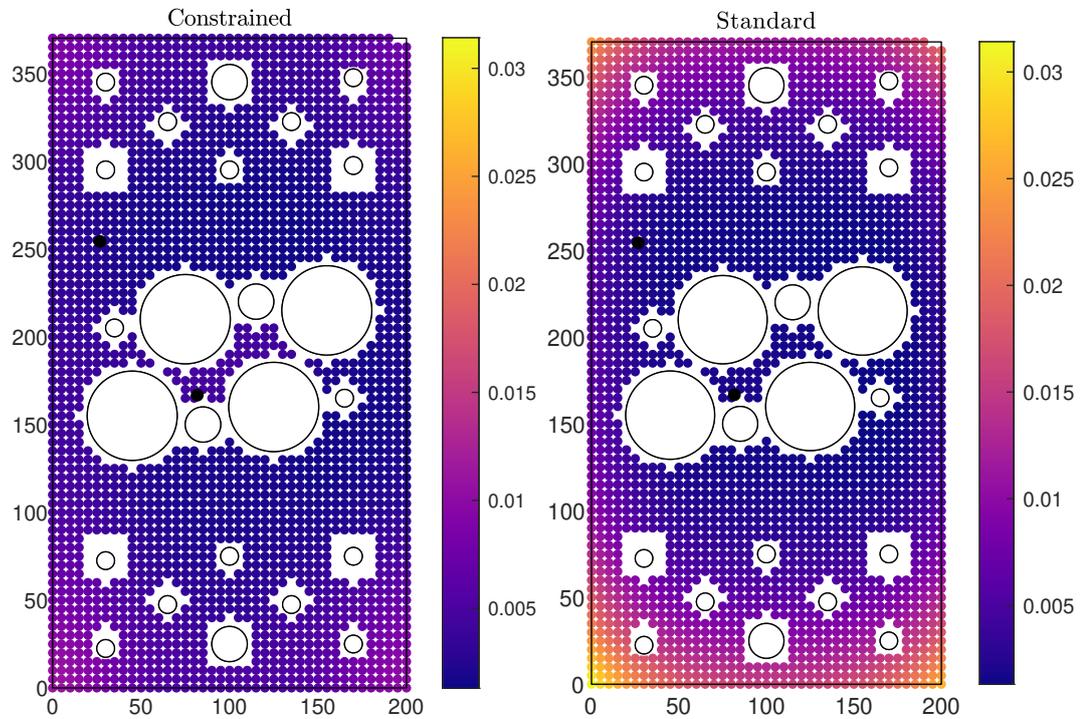


Figure 5.11: Predictive variance of both GP models for sensor pairing 3 & 5. Sensor locations are highlighted by the black circles.

Whilst these plots illustrate the learnt distribution over the AE features, they do not provide a direct measure of how well the predictions reflect the true target values. To analyse the discrepancy between the predicted and true targets for both model forms, a mapping of the squared error of the standard GP subtracted from the squared error of the constrained GP is considered. Under this metric, a positive value indicates that the error is larger in the standard GP, whilst negative values express a larger error in the constrained GP. Figure 5.12 maps this error metric across the test set for sensor pair 3 & 5.

Figure 5.12 clearly demonstrates that in the upper and lower segments of the plate, which are the regions away from the training points, the accuracy of the  $\Delta T$  predictions are improved by the physics-informed GP. As the prediction locations move further away from where the training data points are placed, the level of improvement offered by the constraints generally grows, particularly towards the boundaries at the upper and lower edges of the plate. For regions where training data coverage is good, then the use of boundary constraints will not significantly affect the mean predictions, explaining the similar error score obtained for both models in the centre of the plate.

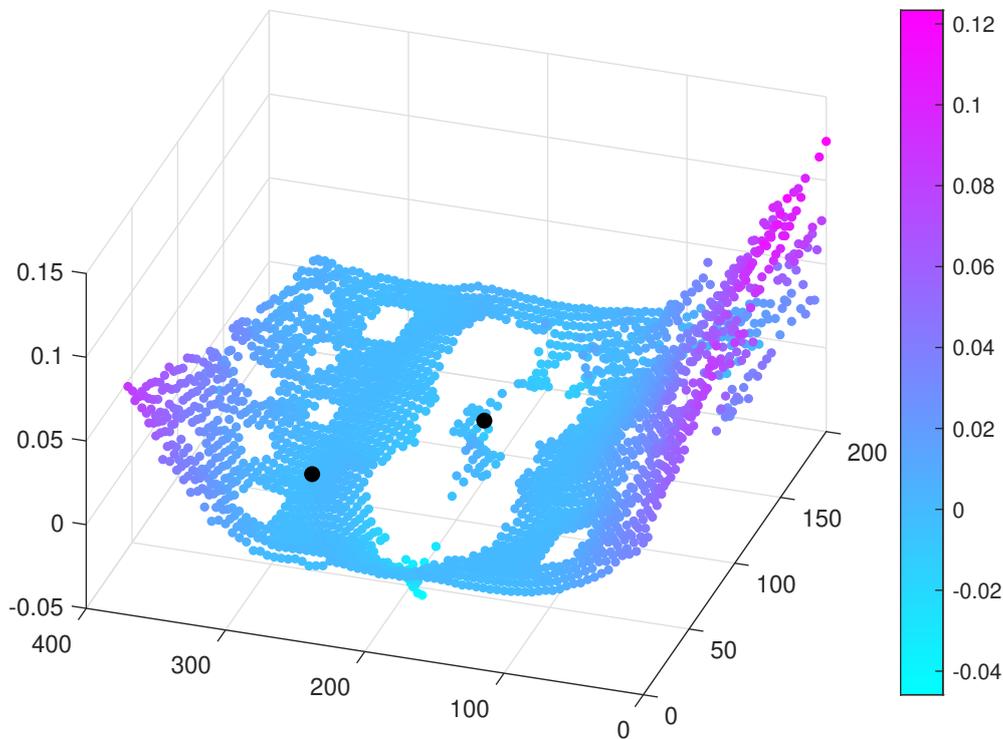


Figure 5.12: Squared error difference between standard and constrained GP for sensor pair 3 & 5. Training observations taken from the middle section of the plate.

If the testing set contained data points on (or closer to) the inner hole boundaries, then it is likely that an error decrease would have been observed at these positions for the constrained model. However, testing measurements were only collected at a minimum distance of around 5mm from the inner holes. A final observation that can be made is that the error reduction obtained in the upper left of the plate ( $x = 0$ ,  $y = 370$ ) is greater than that returned in the upper right ( $x = 200$ ,  $y = 370$ ), despite these two areas being geometrically symmetrical. To explain this behaviour, recall from the previous chapter that the true  $\Delta T$ 's in a particular part of the plate will

contain a level of variability that is dependent on the complexity of the propagation path between the locations of the sources and the receiving sensor. As the propagation path becomes more complex and the variability of the onset times increases, the feature map becomes more challenging to learn. For sensors 3 & 5, the sensor pairing is positioned closer to the upper left of the plate than the right, requiring AE sources from the upper right region to propagate further to the receiving sensors, leading to a significant increase in the level of variability in the onset times in the upper right region than the upper left. As the constraints implemented here act only in relation to boundary conditions, it should not be expected that they provide a means of capturing this variability that arises from sensor positioning, and as such, both models perform similarly in this region.

Examining a different sensor pair mapping, Figure 5.13 plots the difference in square error for sensor pair 4 (sensors 1 & 5). Again, it can be seen that away from the training data, the constraints significantly reduce the error on the  $\Delta T$  predictions, with the maximum improvement occurring on the upper and lower boundaries of the plate. However, unlike Figure 5.12, there is now a significant error reduction on the upper right of the plate, with the previously seen improvement on the upper left now absent. Again, the position of the sensor pair explains this behaviour, with sensor coverage improved in the top right of the plate, but reduced towards the upper left region.

The dependence on the quality of sensor coverage also explains the reasoning for a number of sensor pairs returning either a comparable or worse error metric for the constrained GP. For example, sensors 1 & 2 (index 1) and sensors 6 & 7 (index 26). In these cases, the sensor coverage is poor, with both sensors generally lying adjacent to one another. This positioning therefore results in large portions of the plate requiring waves to propagate further across the structure before being received, resulting in a more complex propagation path and therefore more variable  $\Delta T$  features. If the interest is in improving predictions in scenarios where sensor coverage is poor, then implementing a constrained machine learner in isolation is not suitable. This is because constrained learners are still reliant on some baseline level of training data; for the constrained GP, the covariance structure of the features still needs to be learnt from input data, which is then used exclusively to make predictions. In this case, one may incorporate a mean function derived from physical knowledge, and will be considered as part of future work. As seen in the previous chapter, switching to a heteroscedastic noise model also offers some assistance when sensor coverage is

poor, with the work in this chapter considering only the fixed noise variance case.

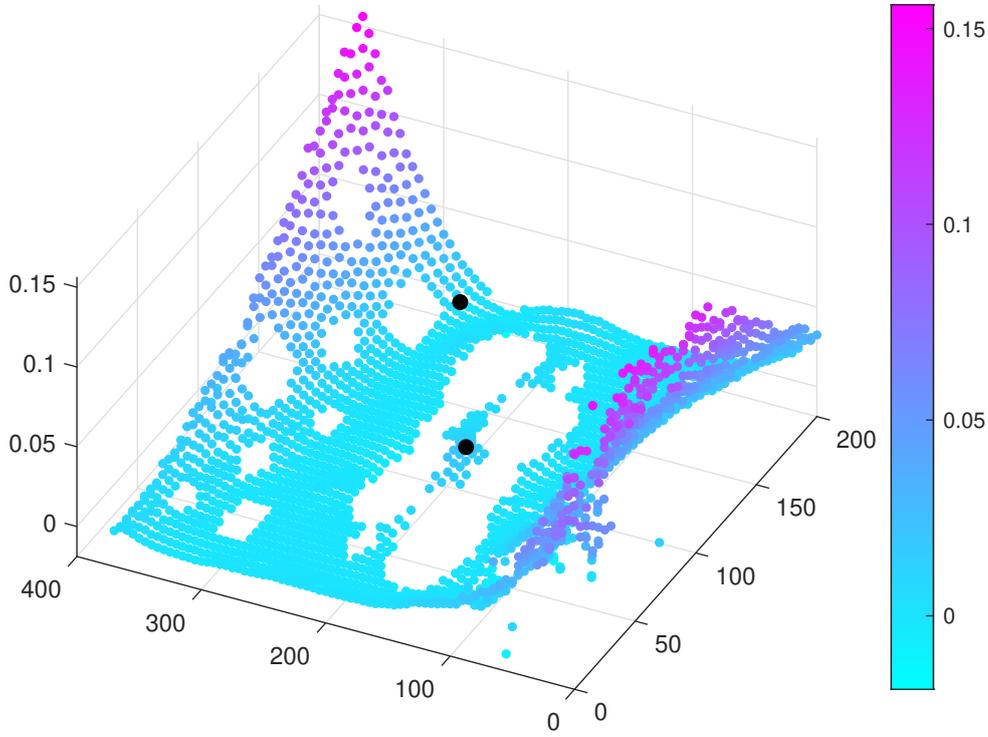


Figure 5.13: Squared error difference between standard and constrained GP for sensor pair 1 & 5. Training observations taken from the middle section of the plate.

## 5.5 Conclusions

In this chapter, combining physical insight with data-driven models has been presented from the perspective of learning  $\Delta T$  maps for acoustic emission localisation. Specifically, it has been shown how a Gaussian process regression model may be constrained such that it satisfies physical boundary conditions that arise from AE onset times.

In the work that has been presented in this thesis until now, only models that are exclusively data-driven have been considered. A limitation when adopting such an approach is that one is solely reliant on the training data during the learning process. For spatial modelling problems such as  $\Delta T$  mapping where there may exist restricted coverage of the full input space or a lack of boundary measurements, the ability to learn an underlying model in a purely data-driven capacity will, therefore, be hampered.

To aid the model learning where there exists reduced availability of training data, the work in this chapter includes boundary condition knowledge into the Gaussian process prior, where it is demonstrated that the predictions made are enhanced, both in mean predictions and an improved confidence, where access to training data is restricted. To embed the physical insight into the model, the covariance function of the Gaussian process kernel is constrained to known boundary conditions, where the resulting draws from the prior are guaranteed to uphold these conditions. Due to the cost of acquiring training data for engineering structures, these are scenarios that consistently arise in SHM, and therefore illustrate the benefit of incorporating physical insight into the Gaussian process through boundary condition constraints.

In this work, physical knowledge was included from the perspective of constraining the kernel function. Broadly speaking, there are two possible avenues of future work. Firstly, one may consider other means of incorporating physical behaviour into the model. For instance, a kernel that better represents the variance of the process driving the behaviour of the  $\Delta T$  functions, or a mean function derived from engineering understanding that reflects the expected onset time behaviour. Secondly, the noise process that is modelled in this chapter is input-independent, with the level fixed globally across the input domain. As such, the limitations of a global noise process discussed in Chapter 4 apply here, with the variance being underestimated in parts of the structure where sensor coverage is poorer. To unite the benefits of including physical insight into the GP with enhanced uncertainty quantification, future work should combine the constraints with a heteroscedastic noise. More will be discussed on the topic of future work in Chapter 7.



# AE LOCALISATION FOR WIND TURBINE BEARINGS

### Highlights:

- *The localisation framework is extended to non-Euclidean spaces, allowing localisation in spherical structures such as bearings.*
- *Acoustic emission localisation is investigated in a scaled-down wind turbine gear box bearing.*
- *Good localisation performance is achieved in the presence of noisy  $\Delta T$  features.*

So far in the thesis, attention has only been given to the development of localisation approaches in structures that are characterised in Euclidean space. Many high value assets that engineers are interested in monitoring, however, can be more appropriately represented by non-Euclidean measures. For instance, many components that are found in rotating machinery, such as bearings, are spherical, and exist more naturally in other topological spaces.

In this chapter, the application of the localisation framework developed so far in a non-Euclidean space is considered, allowing for the application of probabilistic acoustic emission localisation in spherical components like bearings. In the initial framework, it was assumed that the domain may be represented in Euclidean space,

embedded through the use of kernels that express the covariance between inputs as a function of a Euclidean distance. For the geometry of a bearing, a Euclidean distance will not adhere to the topological constraints of this spherical space, and is therefore not the most appropriate measure. To bypass this issue, the work in this chapter expresses the covariance as an eigendecomposition, where the inputs are projected onto some eigenbasis. The resulting kernel satisfies both the conditions of a valid covariance function, as well as the topological constraints imposed by the geometry of the bearing. The proposed method is then applied to localise AE on a spherical roller bearing that is designed to replicate planetary support bearings that are found in wind turbine gearboxes.

## 6.1 AE localisation within bearings

Bearings are an integral part of many mechanical systems in industries such as power generation, machining and automotive. It is, therefore, very important that these components are properly maintained in a way that reduces cost whilst maximising output, particularly in the offshore wind energy sector where bearing failures are the leading mechanical cause of wind turbine down-time [146]. As has been discussed, SHM is the field of study concerned with developing such maintenance strategies. A distinction is made when working in the context of rotating and reciprocating machinery, where the development of monitoring tools for these systems is defined as condition monitoring (CM). SHM and CM are entirely similar in principle, both requiring that monitoring data is acquired and fed into some statistical or physics-based model to infer a health state. Where the two fields differ is that CM deals only with components that are periodically excited in a continuous manner by the system itself; for instance, a driving shaft in a gear box.

Whilst AE localisation has received a lot of interest in a structural health monitoring context, the literature on applying similar methods in condition monitoring is much more sparse. Rogers [147] was the first to consider AE localisation in roller bearings, localising events along a radial line between two sensors (linear localisation) that occurred above a pre-defined threshold from TOA measurements. However, as the study only considers the use of two sensors that are positioned diametrically, the approach will be prone to returning non-unique location predictions, with the assumption made that the source lies along a radial line that passes through the

two receiving sensors. Additionally, no attempt is made to quantify the error of the localisation. Elforjani and Mba [45] implemented a similar strategy for localisation in slow speed roller element bearings, but instead incorporated the arrival times at four different sensors. The source location was then calculated analytically, taking the wave propagation speed to be that of the velocity of the zeroth-order symmetrical Lamb wave mode for the frequency and material thickness being investigated.

Although for smaller bearings it can be more cost-effective to simply replace them as signs of a fault begins to emerge, given that the size of bearings found in wind turbines has been increasing rapidly over the last decade, replacing old-for-new can be both a logistically and economically less favourable action than repairing individual defects at a known location. For instance, if an individual ball/roller was damaged in a ball/roller bearing, then it would be possible to replace the defective component at a much lower cost than the entire bearing. Were the defect located on the raceway, the surface may be polished or even ground down, and if needed, larger balls/rollers installed. Whilst this option may involve repairing or replacing a number of components, it could significantly reduce downtime where the full turbine does not have to be disassembled. From a sustainability perspective, reusing as much of the original bearing as possible is clearly desirable. For these reasons, there has been a renewed interest in localising defects in wind turbines. Examples of recent work include that of Naumann [58] and Martinez [148], who both consider the multilayer perceptron (MLP) variant of feedforward neural networks to learn the relationship between  $\Delta T$  measurements and source locations for a planetary wind turbine gearbox bearing, and like [31], directly model the inverse problem. A set of artificial acoustic emission events is first used to train the MLP to map from  $\Delta T$  to source location, allowing events with an unknown origin to be predicted based on observed arrival times. As seen previously in the thesis, the advantage of these types of methods is that no assumptions regarding the behaviour of the wave propagation are required to be made, which can instead be implicitly learnt from measured data.

With that said, neural network architectures typically contain many learnable parameters and so require large training data sets. As argued in the previous chapter, it is often the case that training data is limited when developing  $\Delta T$  maps, and so methods capable of handling small training sets are desirable in the pursuit of practical solutions for localising damage in our structures. Additionally, both approaches are purely deterministic, and do not attempt to model uncertainty on measurement observations, with no subsequent measure of uncertainty on localisation predictions.

In the context of condition monitoring, there has been little-to-no attention given to probabilistic AE localisation and thus there is progress to be made in this domain. In particular, for larger scale applications such as for bearings found in new generation wind turbines, operational noise will have a significant impact. With this in mind, it is essential that CM localisation methods are able to account for uncertain measurements, which thus far, has received minimal attention in the current literature.

To this end, the work presented in this chapter aims to develop probabilistic AE localisation in non-Euclidean spaces, allowing the extension of the methods developed in this thesis to bearings and other spherical components.

## 6.2 Generalising to non-Euclidean domains

In standard GP analysis, the covariance between two points is commonly defined as some function of their Euclidean distance. This measure is used in many popular kernels such as the squared exponential and the Matérn family of functions, for example, the Matérn 3/2, which is used extensively in this thesis. The standard GP formulation is, therefore, often not readily applicable to non-Euclidean spaces. The task is, then, to adjust the kernel selection to be more reflective of the geometric space of concern. In the case of a bearing, the spherical topology of this component may be more suitably represented through a geodesic distance. However, naively substituting out the Euclidean distance for a geodesic measure in the aforementioned kernels, as well as many others, will lead to an ill-defined kernel [149], with the resulting covariance matrix not satisfying the condition of positive (semi)definiteness, and so violating the requirements of a valid kernel.

One way to circumvent the issue of defining a valid kernel on a non-Euclidean space is through the use of spectral methods, where the kernel function is reconstructed as an eigendecomposition, as considered in Chapter 5. That is to say,

$$k(\mathbf{x}, \mathbf{x}') = \sum_{i=1}^{\infty} \lambda_i \phi_i(\mathbf{x}) \phi_i(\mathbf{x}'). \quad (6.1)$$

where  $\lambda_i$  is a vector of eigenvalues and  $\phi_i$  the eigenfunctions of a given kernel [150]. This is the approach that will be adopted in this chapter, where the eigendecom-

position is formed of the Laplacian eigenfunctions of the input domain [140], as considered previously in the thesis. In this sense, equation (6.1) can be rewritten as,

$$k(\mathbf{x}, \mathbf{x}') \approx \sum_i^m S(\sqrt{\lambda_{L,i}}) \phi_{L,i}(\mathbf{x}) \phi_{L,i}(\mathbf{x}'), \quad (6.2)$$

where the spectral density,  $S$  is set to be that of a Matérn 3/2, with  $\{\lambda_{L,i}, \phi_{L,i}\}_{i=1}^m$  representing a set of  $m$  Laplacian eigenfunctions and eigenvalues associated to the bounded domain. As the spatial derivative of the onset time functions will always be equal to 0 at the boundary of the domain, the eigendecomposition can be solved according to,

$$-\nabla^2 \phi_{L,i}(\mathbf{x}) = \lambda_{L,i}^2 \phi_{L,i}(\mathbf{x}), \quad x \in \Omega, \quad (6.3)$$

$$d\phi_{L,i}(\mathbf{x}) = 0, \quad x \in \delta\Omega, \quad (6.4)$$

Implementing the eigendecomposition for the geometry of a bearing then proceeds similarly to that of the plate structure, forming a finite difference approximation of the Laplace operator. Solving in parallel with the known boundary conditions, the resulting matrix represents the discretised Laplacian of the domain, for which the eigenpairs can be identified.

## 6.3 Case study

To demonstrate the methodology developed in this chapter, data collected from a spherical element roller bearing by Martinez [148] are used. An image of the set-up is shown in Figure 6.1. The test rig, which was initially designed to investigate fatigue life of wind turbine roller bearings, contains an SKF NU1010 ECP spherical roller bearing which has an inner raceway diameter of 57.5mm and an outer raceway diameter of 80mm. This particular bearing has been chosen as it is representative of the contact conditions experienced by commercial wind turbine bearings, but at a geometrically scaled-down level [151]. It, therefore, serves as an ideal initial investigative case study for the application of the probabilistic AE localisation framework to a bearing element.

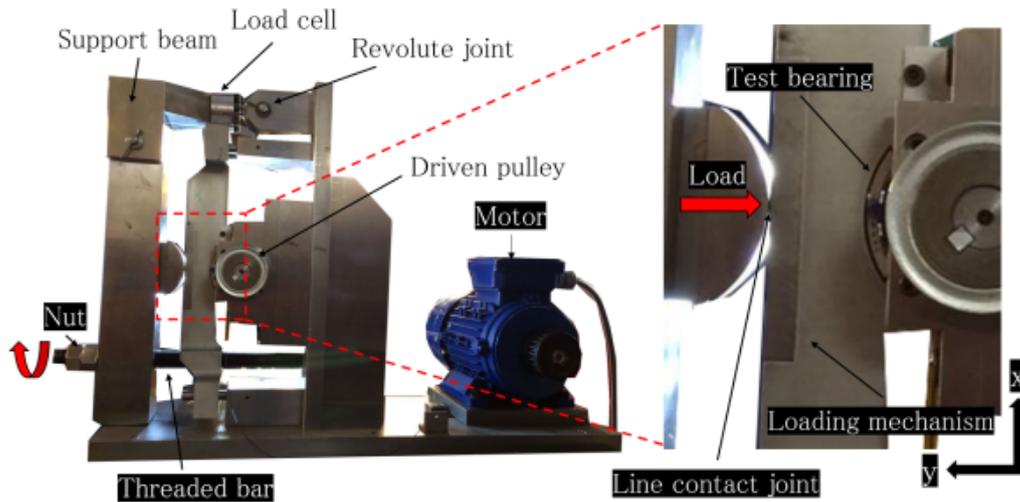


Figure 6.1: Full test rig set up - recreated from [148].

To acquire the artificial acoustic emission measurements as required to generate the  $\Delta T$  maps, the upper support of the rig was dismantled so that the outer race of the bearing was exposed. Three Mistras NANO30D sensors were then mounted onto the support bracket, each with a resonance at around 125-175kHz. To simulate acoustic emission propagation, 30 H-N sources were excited at each of the positions on the outer raceway marked on Figure 6.2, leading to a total of 300 AE events. The angular position of each event relative to the y axis is provided in Table 6.1. An associated  $\Delta T$  value for each sensor pair is then obtained through the use of the AIC onset picker (see Chapter 2.3).

By focusing on the outer raceway, the damage scenario more likely to occur is investigated [148]. This is due to the fact that unlike the inner raceway, where stress is periodically distributed as the race rotates, stress concentrations will manifest at fixed locations. Defects on the outer raceway will, therefore, accelerate at a greater rate than at the inner.

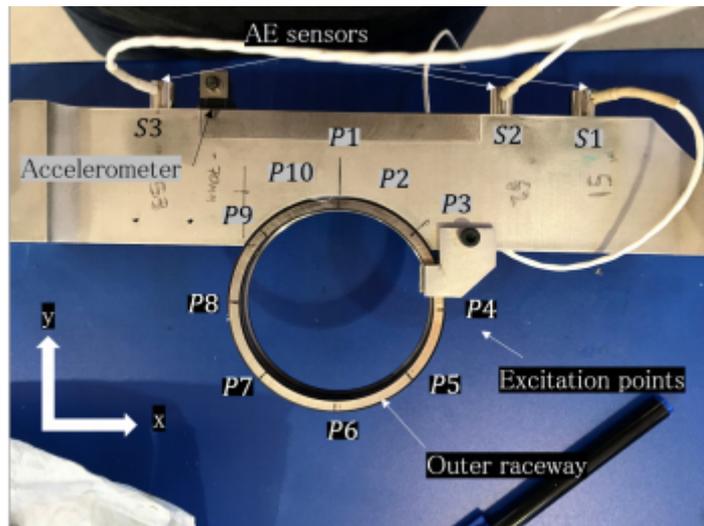


Figure 6.2: Isolation of outer raceway, with sensors and artificial AE event locations marked - recreated from [148].

Position	1	2	3	4	5	6	7	8	9	10
Angle ( $^{\circ}$ )	0	11	45	90	135	180	225	270	315	326

Table 6.1: Angular positions of events. Angle defined relative to the y axis in Figure 6.2

### 6.3.1 Exploration of the data

To examine the data collected for the bearing test rig, the mean and standard deviation of the true  $\Delta T$ s at each position are calculated for every sensor pair. These values are then plotted in Figure 6.3. Overall, the  $\Delta T$  functions vary smoothly about the circumference of the raceway, and will be continuous in the range  $0^{\circ} \leq \theta \leq n360^{\circ}$ , where  $n$  is an integer.

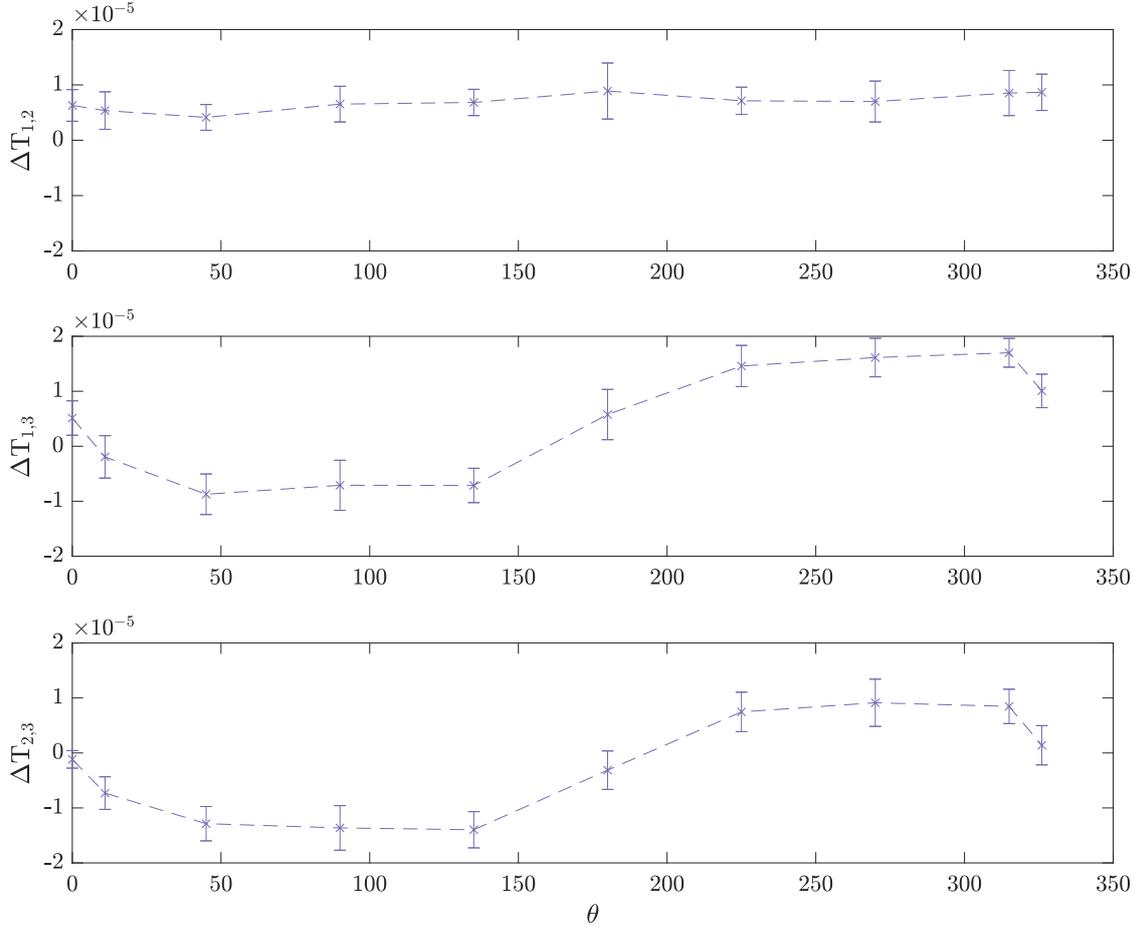


Figure 6.3: True  $\Delta T$  values for all the sensor pairs across the bearing.

Comparing the behaviour of the mean onset times across the pairings, the range of  $\Delta T$  is much lower for sensors 1 & 2 compared to 1 & 3 and 2 & 3. Given that sensors 1 and 2 are positioned close together, the propagation path of a given event to the two sensors will be similar, resulting in little discrepancy between their arrival times. Where two sensors in a pairing are spaced further apart, there is a greater potential for larger differences in arrival times, with the magnitude of the difference depending on the exact location of the event relative to the sensors. For example, for pairing 1 & 3, P3 - P5 return larger differences in onset times, with these positions located closer to sensor 1 than sensor 3. Similar behaviour is observed at P7 - P9, except now these positions are located closer to sensor 3 than sensor 1. Conversely, P6 lies somewhat equally between the two sensors, resulting in a smaller  $\Delta T$  value.

Considering the variation in the arrival times, sensor pair 1 & 3 has the highest combined standard deviation, with the second highest 2 & 3, followed by 1 & 2. As the

mean onset times for pairing 1 & 2 have a much lower overall range than that of the other two pairs, it makes sense that this pairing also produces the smallest standard deviation. However, having a measure of variation that is dependent on the range of the function values may mask the true underlying behaviour. For instance, what may actually be relatively small variations between readings at repeating positions can be amplified by an overall large range of  $\Delta T$ s, and vice versa. We, therefore, want a measure of variation that is relative to the range of the corresponding values, which can be acquired by normalising the  $\Delta T$  values to between 0 and 1.

Plotting the mean and standard deviation of the transformed features for each pairing in Figure 6.4, a very different picture is seen to earlier, with the  $3\sigma$  interval of sensors 1 & 2 now much larger than the other two pairings, indicating that this sensor pair has a much greater degree of fluctuation over repeated readings than the other two pairs. As sources traveling to both of these sensors will propagate along similar paths, small changes in the arrival times of repeated events will result in a bigger relative shift in  $\Delta T$  times, causing this pairing to observe greater variations. Additionally, where both sensors are placed similarly spatially, coverage over the full physical domain will generally be poor. Where this is the case, in parts of the structure that are not well covered by the sensors, extracting arrival times of the raw waveforms may become more of a challenge.

What is consistent between both the standard and normalised measure of standard deviation is that the positions at and around  $180^\circ$  return onset times with the most variation. This observation is inline with the discussions held in Chapter 4, where it was explained how an increasing distance between a source and receiver likely leads to a more complex propagation path, manifesting in greater variations in arrival time estimates, and thus  $\Delta T$  vectors, in events similarly positioned. More will be discussed on positions with greater ranges of onset times, and the effect this has on the localisation, towards the end of the chapter.

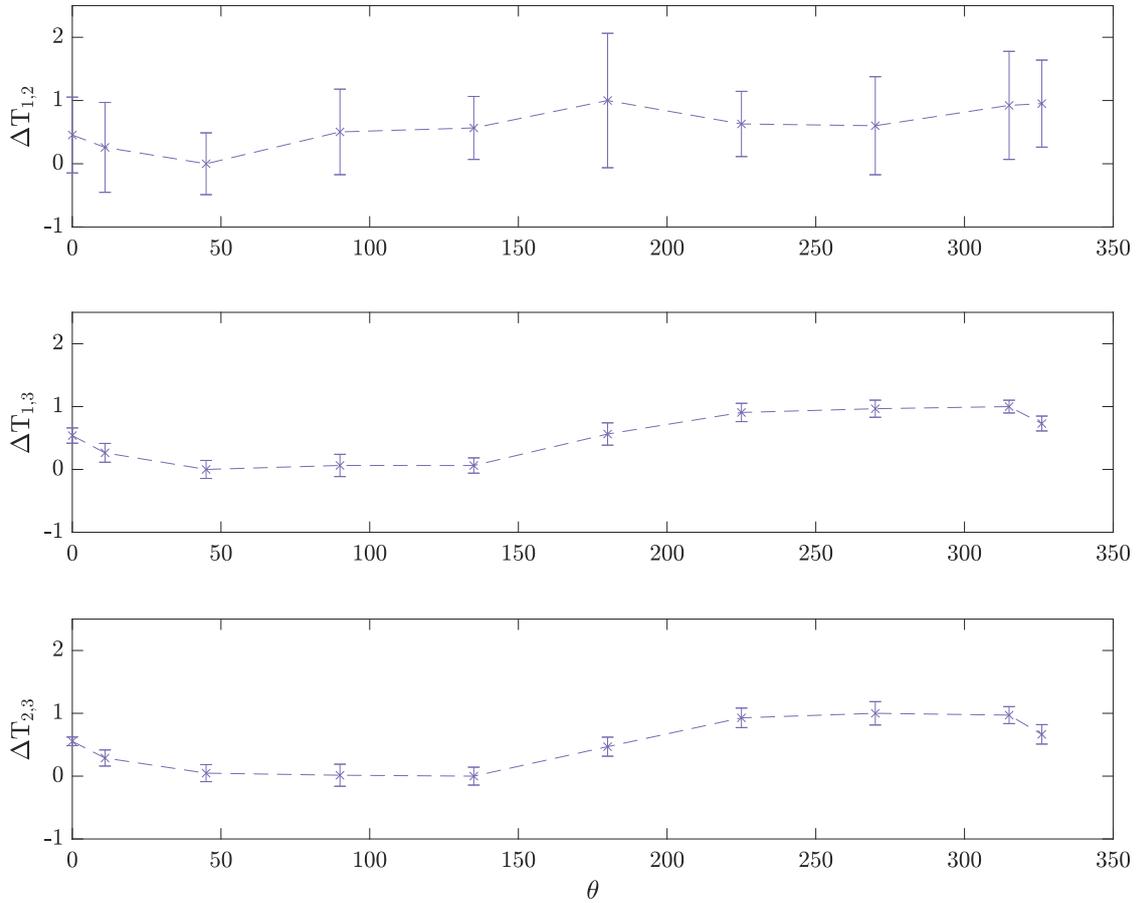


Figure 6.4: Normalised  $\Delta T$  values for all the sensor pairs across the bearing.

## 6.4 Localising AE events in a bearing

In order to investigate the performance of the proposed localisation strategy, the onset time data that was introduced in Chapter 6.3 was broken up into a training and testing set. For the training set, the first 20 repeats recorded at each location were included, with the testing set encompassing the remaining 10 repeats. Utilising the entire training set, a manifold GP representation was learnt for each of the sensor pairs, with a total of 3 GPs trained. Each observation in the test set was then used to quantify the accuracy of localising AE events with an origin unknown to the machine learner, with a total of 100 test events (10 at each location).

### 6.4.1 Learning $\Delta T$ maps

To investigate how well the GP has learnt the  $\Delta T$  functions, Figures 6.5-6.7 show the predicted  $\Delta T$  values on the test set for each of the sensor pairs, with the known values (which the GP has not seen in the training period) also marked on the plots.  $3\sigma$  confidence intervals are also included, having been derived from the predictive variance of the GP.

Considering sensor pair 1 & 2, following the discussion in the preceding section, there exists a large amount of variation in these  $\Delta T$  values over each of the repeats. As discussed extensively in previous chapters, scenarios where there exists a wide variation in measured targets at a particular input location are where probabilistic methods such as the GP prove extremely useful. As each output of the GP is represented as a distribution instead of a single point, it follows that the uncertainty on a prediction can be encoded through the predictive variance. As such, the GP is able to learn a confidence bound around the mean predictions, which as shown in Figure 6.5, encompasses the vast majority of the true  $\Delta T$  values. Therefore, although the mean output may not exactly match that of the target  $\Delta T$ , the GP retains sensitivity to this through the predictive variance, that later is fed into the source location prediction.

Interestingly, the uncertainty bounds do not contract as much as one may expect around the training points. In a GP regression, one typically observes a high certainty around training points, where an observation from the latent function is available to condition on. However, in this case, the large variation that is present on  $\Delta T_{1,2}$  at each measured position blows up the overall amplitude of the noise learnt by the GP, even at training locations. As such, the confidence bounds across the input space are relatively wide, indicating that the GP is not particularly confident in how well this feature has been learnt. As discussed in Chapter 4, it is likely there would be significant benefit in adopting a local noise process in the regression model. However, incorporating heteroscedasticity into the spectral approach will be the topic of further work, and discussion regarding this matter will be left to Chapter 7.

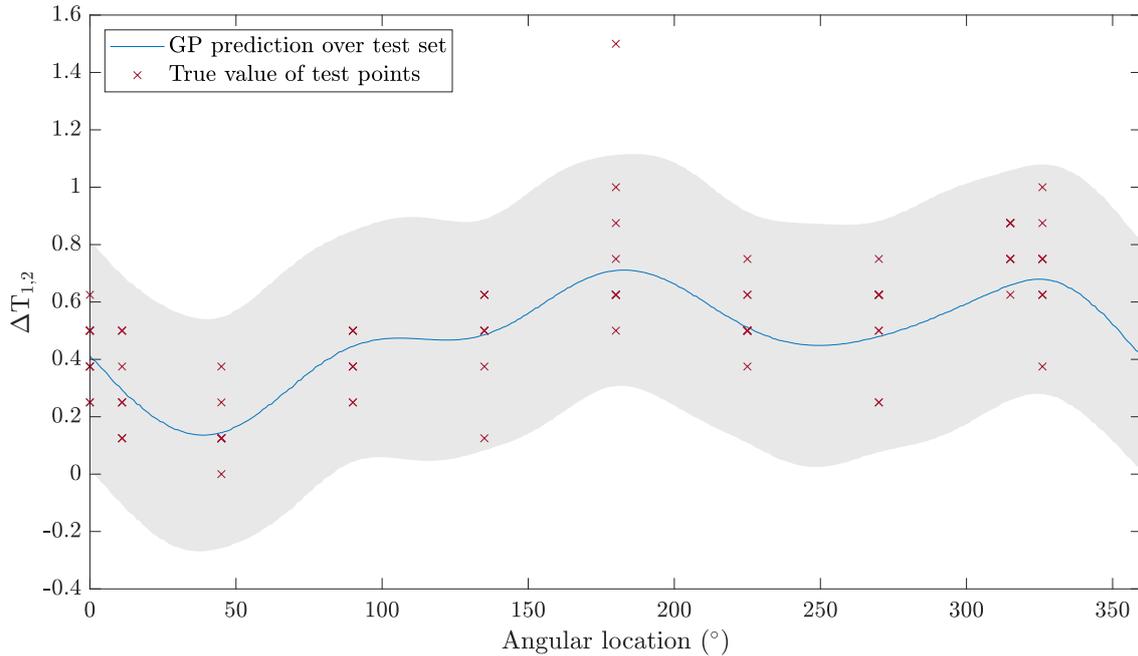


Figure 6.5:  $\Delta T$  predictions (normalised) across the test set for sensor pair 1-2. The true test targets are shown by the red cross.

Moving on to sensor pair 1 & 3 and 2 & 3, Figures 6.6 and 6.7 show that the target values from these two pairs fluctuate much less than for sensor pair 1 & 2, allowing for a mean fit that is in closer agreement with the test target values. This behaviour is then reflected in the uncertainty bounds, which are now tighter about the mean. With less variation in the  $\Delta T$  values, the predictions are also able to be more confident around the training locations. Additionally, it can be seen that the learnt  $\Delta T$  functions for both of these pairs is similar, and arises from the close positioning of sensors 1 and 2. This results in closely aligned paths of propagation, and so similar onset time differences.

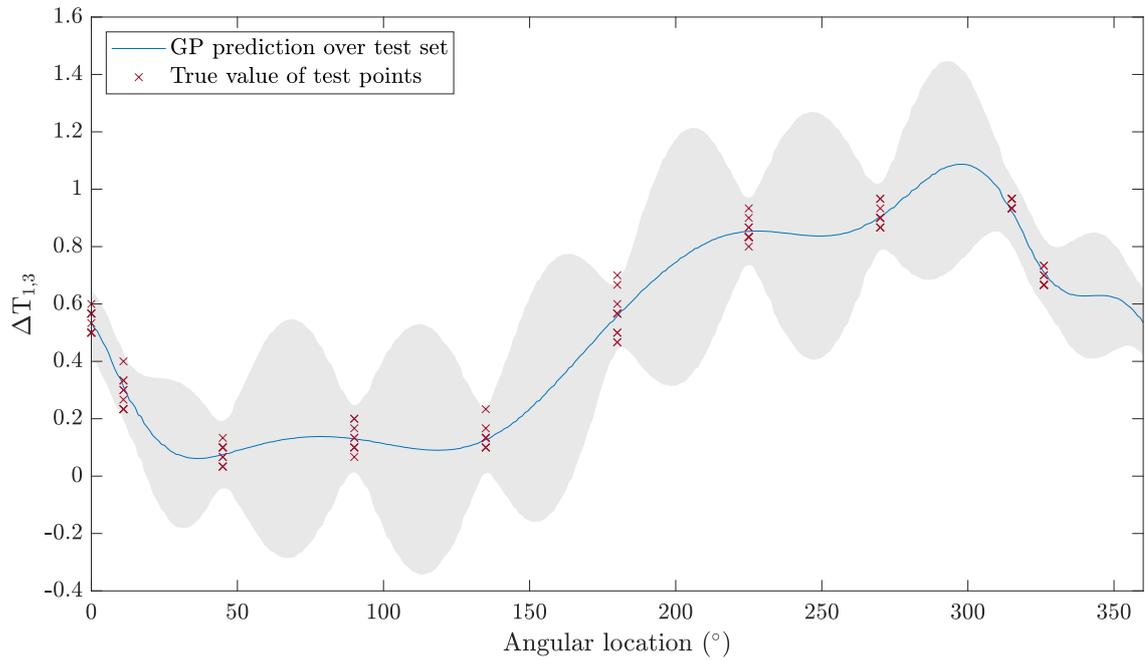


Figure 6.6:  $\Delta T$  predictions (normalised) across the test set for sensor pair 1-3. The true test targets are shown by the red cross.

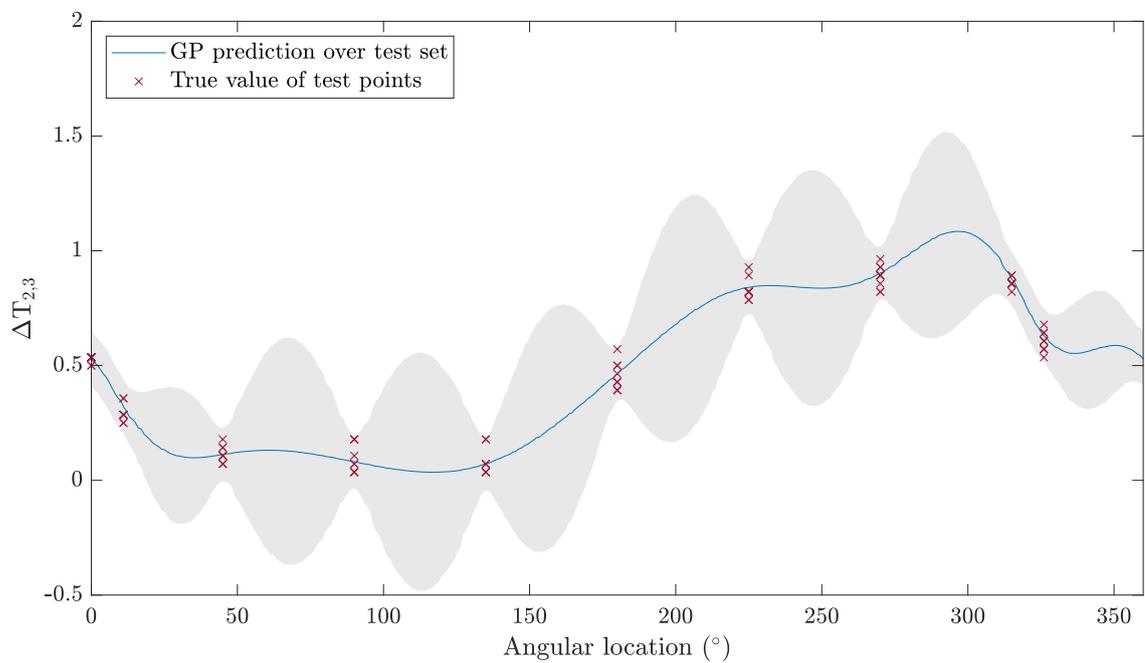


Figure 6.7:  $\Delta T$  predictions (normalised) across the test set for sensor pair 2-3. The true test targets are shown by the red cross.

A final point to note regarding these results is that the  $\Delta T$  values appear to be fairly

evenly spaced along the  $y$  direction of the plots, suggesting that they are drawn from a discrete distribution instead of a continuous distribution. Although we know that the  $\Delta T$  values should be continuous given they represent time, the reason for this behaviour is that the floor of the sampling rate is being hit, and ideally, the experimental set up should have considered an increased acquisition rate. Where data is known to be discrete, it makes sense to switch to a discrete likelihood model. However, given that it is expected that the  $\Delta T$  values be continuous, the work in this chapter will continue with a Gaussian likelihood.

### 6.4.2 Source localisation

In order to visualise the localisation process for the bearing, Figure 6.8 plots the conditional likelihood for each sensor pairing, evaluated for an event taken from the test set. The true location of this source is marked with a red cross. A first observation that can be made is that there exists quite a wide spread of regions of high likelihood in pairing 1 & 2, with the locations of high likelihood more diffused over the bearing than for the other two pairs. The difficulty in recovering the underlying source location contour for this pairing results from the high variance of the  $\Delta T_{1,2}$  predictive distributions.

With a wider distribution, the total probability mass under the distribution will be spread over an increased range, resulting in lower likelihoods. For the two other sensors, the regions of high likelihood are more distinct, with a larger maximum likelihood value. When calculating the marginal likelihood, these two models will, therefore, contribute more towards the total source location likelihood.

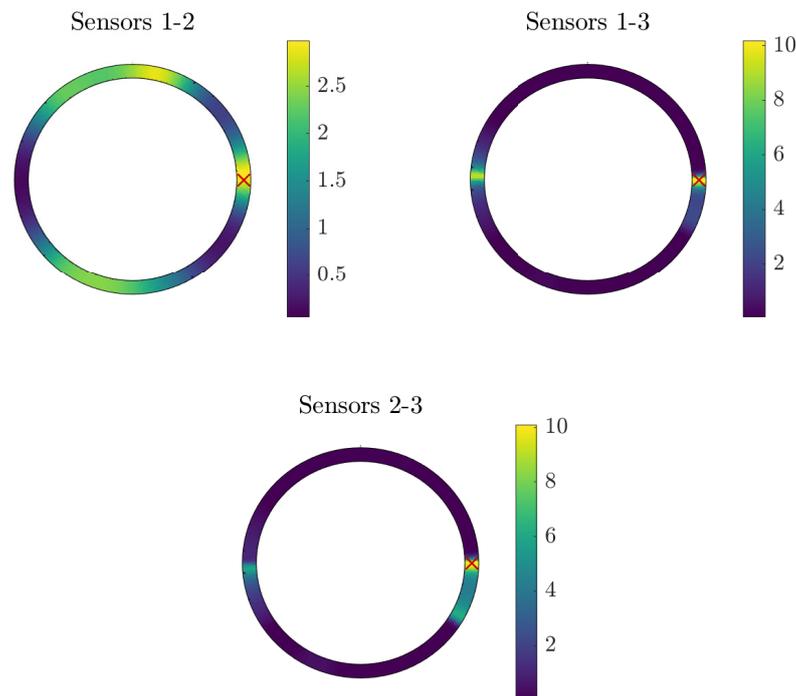


Figure 6.8: Conditional likelihood of emission location for a test event. The true location of the event is marked by the red cross.

Again, how well the  $\Delta T$  behaviour has been learnt for these pairings can be viewed as a primary indicator for the quality of the conditional likelihood map. For sensors 1 & 3 and 2 & 3, there is less variation about the mean onset times, and so the resulting GP inference over these latent functions obtains a lower predictive uncertainty. The corresponding conditional likelihood maps that are returned, therefore, allow for a better identification of the possible event locations arising from the observed onset difference times. It can also be seen that these two pairs return a similar likelihood surface, caused by sensors 1 and 2 being positioned on a similar line of trajectory relative to the bearing itself.

At this point, the argument can be made for maximising the information gained from a set of sensors by optimally distributing the sensors over the structure. Clearly, this is not a trivial task, and requires the consideration of a number of factors. For instance, what cost function should be used, how should the geometry of the structure be incorporated into the model. Despite not having access to what would

be optimal sensor positions for the experimental set up here, it is not much of a stretch for the reader to appreciate that the sensor configuration used here is almost certainly not optimal, and, therefore, improvements in sensor positioning could be made so that the structure may be better covered. Further discussion on this matter will be considered in the conclusions and future work chapter.

Marginalising over the individual models, a combined likelihood can be calculated, with contributions from all the sensor pairs. This likelihood map is shown in Figure 6.9. The figure shows that the highest likelihood corresponds with the true source location, indicating an accurate prediction. There does also exist areas of high likelihood (relative to the majority of the bearing) away from the event origin, both beneath the true location and on the opposite side of the bearing. Given that the  $\Delta T$  functions are continuous, locations that are spatially similar have similar onset times, resulting in this higher likelihood area below the true location. For the latter, an explanation can be arrived at through consideration that the bearing is geometrically symmetrical, with sensor 2 and 3 placed at an equal but opposite distance away from the bearing. As such, there will be elements of symmetry arising in the likelihood surfaces. One approach to reduce the relative likelihood mass at positions opposite to the true emission location would be to add a fourth sensor, increasing the number of individual models that can provide information on the source location. However, it is demonstrated here that accurate location is still very much possible with the minimum number of pairings required for unique location.

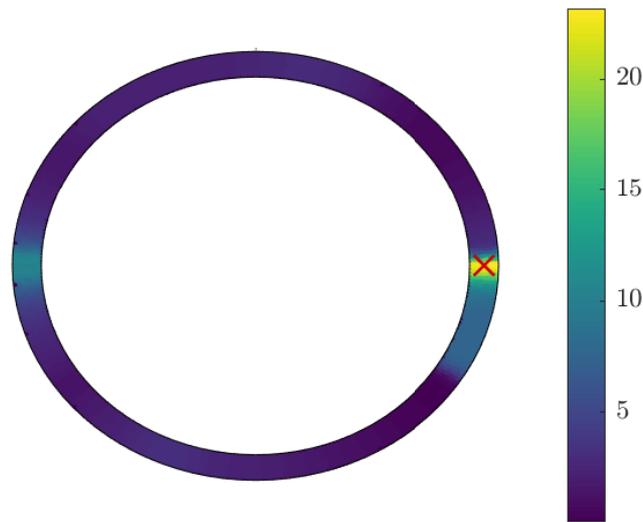


Figure 6.9: Marginal likelihood of source location for a test event. The true location is marked with a red cross.

Let us now consider the prediction of a second event, but this time at  $\theta = 180^\circ$ . As there is more variation in the onset times at this location than the others, predicting at this position likely presents the biggest challenge. Figure 6.10 shows the marginal likelihood for a test event at this location. The figure shows that there is both a high region of likelihood at the true location, but also at a shifted location of  $30^\circ$  to the right. Unlike for the previous event, both of these regions produce a similar likelihood, and, therefore, there is not an obvious unique location prediction. Scenarios such as this, however, demonstrate the power of the probabilistic interpretation offered by the method. With the ability to flag multiple potential emission locations, a user may infer that there is less certainty regarding a location prediction when more than one region of similarly high likelihood is returned. Each of these regions may then be inspected for damage. For a deterministic model, one would be limited to recovering only a single location (which in some cases is equivalent to the maximum likelihood solution), and so does not have the capacity to account for different possible source locations. As seen here, this will be particularly problematic when incoming  $\Delta T$ s are highly variable and/or contain significant levels of measurement noise.

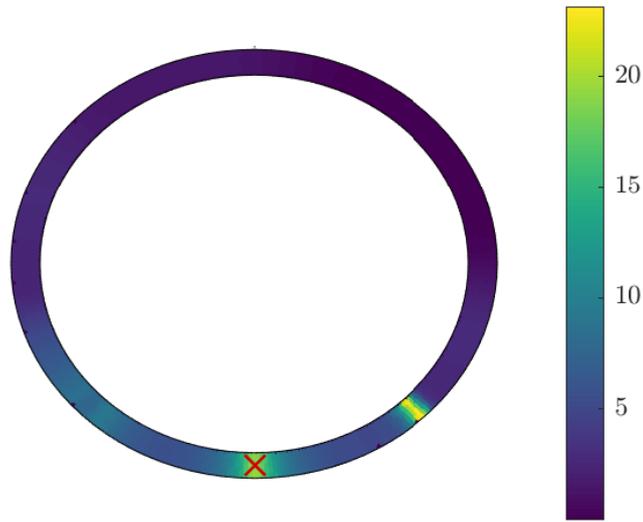


Figure 6.10: Marginal likelihood of source location for a second test event with a less certain prediction. The true location is marked with a red cross.

Given that repeated measurements were collected at each location, it is possible to access a second test event at  $180^\circ$ , for which the marginal likelihood is plotted in Figure 6.11. Here, a more distinct single region of high likelihood is now returned in comparison to Figure 6.10, indicating that the model has been able to better identify a unique location that the event was likely generated at. However, the peak likelihood is actually lower than for the previous event, and results from the  $\Delta T_{1,2}$  for the event in Figure 6.11 being further away from the mean learnt in training than that corresponding to the event in Figure 6.10.

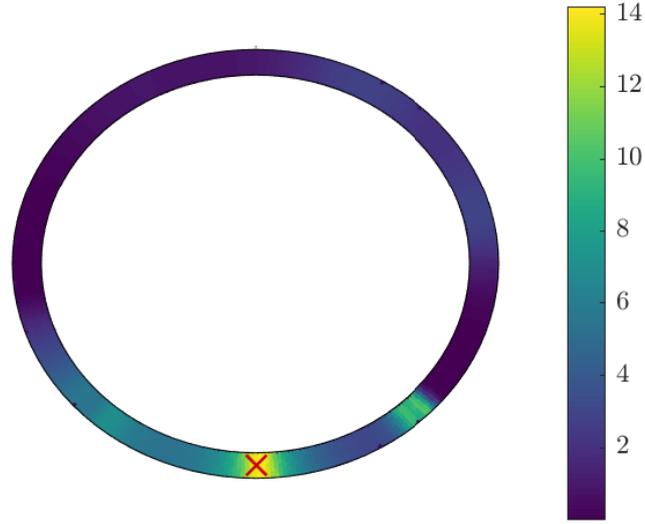


Figure 6.11: Marginal emission location likelihood for a test event at the same location as the preceding figure, except with a more unique event location prediction. The true location is marked with a red cross.

With regard to assessing the localisation performance across a range of positions on the bearing, the most likely emission location was obtained for each of the test events, equalling a total of 100 acoustic emission events that are composed of 10 repeats at each of the 10 positions. To quantify the mean performance at each of the positions, the circular mean and standard deviation were calculated, which are expressed as,

$$\bar{\theta} = \tan^{-1} \left( \frac{1}{N} \sum_i \sin(\theta_i), \frac{1}{N} \sum_i \cos(\theta_i) \right), \quad (6.5)$$

$$\sigma_{\theta} = \sqrt{-2 \ln(R)}, \quad (6.6)$$

where,

$$R = \sqrt{\frac{1}{N} \sum_i \cos^2(\theta_i) + \frac{1}{N} \sum_i \sin^2(\theta_i)}. \quad (6.7)$$

These results are presented in Figure 6.12. Investigating the figure, it can be seen that the majority of predictions yield a mean value that is in good agreement with the true location, with a number of positions consistently returning an accurate prediction across the repeats. For positions in the range of  $90\text{-}270^\circ$ , although the mean prediction generally matches that of the true value, there exists a fairly significant level of fluctuation across the predictions at each location. This increased variation in the predictive mean is a consequence of the increased range of  $\Delta T$  values that occurs at these positions, and highlights the importance of accounting for a distribution of onset times at a particular location, as proposed here. Geometrically, it makes sense these positions result in a wider range of onset times, with angles  $90\text{-}270^\circ$  located on the lower half of the bearing, which is the area the furthest away from the three sensors. The propagation path of these positions is also much more complex, requiring the elastic waves to firstly propagate partially along the raceway, and then through the fixing bracket that holds the bearing, before being received at the sensors. For the upper half of the bearing, there is a direct propagation path between the positions and all of the sensors, which results in a simpler onset time identification. Following on from the discussion in section 6.3, it is likely that localisation performance at angles with larger prediction variance would be significantly improved if placement of sensors 1 and 2 were better distributed over the bearing, or if an additional sensor was used.

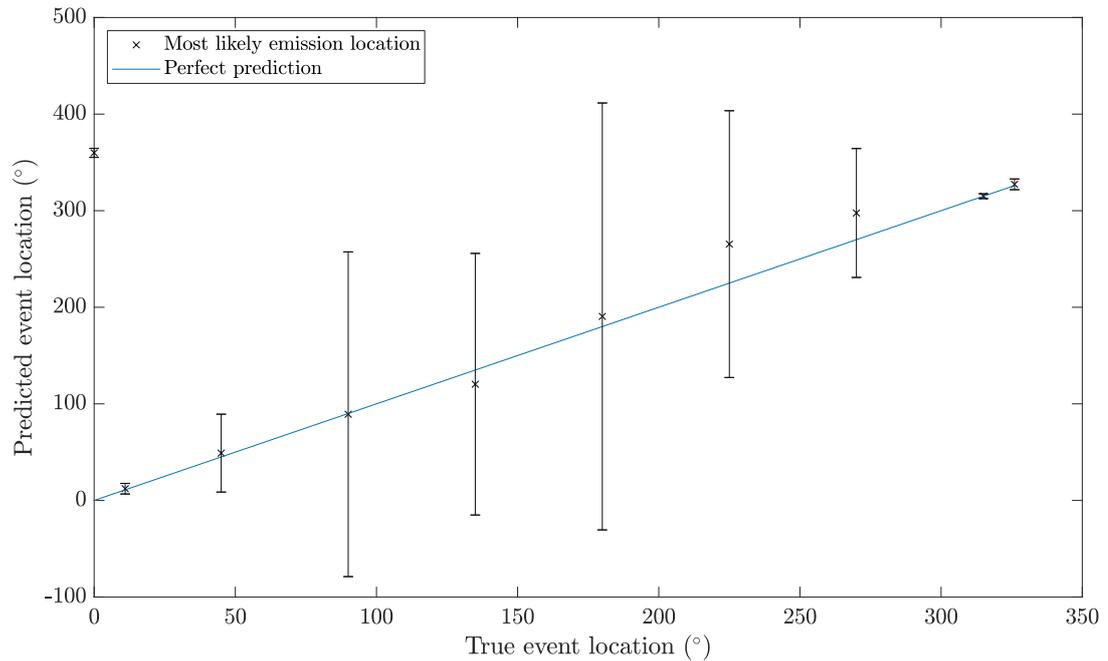


Figure 6.12: Circular mean of the most likely emission location obtained for each position, marked with  $\pm 3$  standard deviation. Note that the location prediction at  $0^\circ$  is almost exactly  $360^\circ$ .

## 6.5 Conclusions

In this chapter, the localisation framework established earlier in the thesis has been extended to non-Euclidean geometry, allowing for the probabilistic localisation of AE in spherical engineering components such as bearings. To generalise the methodology to such domains, the Gaussian process is constructed as an eigenfunction expansion, with the resulting kernel expression becoming consistent with the topology of domains that are not most aptly represented in Euclidean space.

To investigate the proposed approach, an experimental case study consisting of a roller element bearing that replicates a scaled-down version of those found in wind turbine gearboxes was considered. From the collected data, not only were the usual  $\Delta T$  maps able to be learnt, but the acquisition of multiple readings at each position enabled an investigation in the variability present on the  $\Delta T$  features over successive events generated at the same location. This analysis enabled discussions on the effect of sensor placement on the ability to learn these maps held in previous chapters to be more closely investigated, as well as consider predictive performance when the

incoming onset times at repeated positions have a high degree of variation. In an operational context, it is likely that the  $\Delta T$  times will contain differing amplitudes of measurement noise, leading to  $\Delta T$  times that instead of being unique for a given location, will deviate about some mean value. Through use of the likelihood of emission location maps, the method demonstrated the ability to suitably handle this type of behaviour; a significant benefit of the proposed approach. Where there was increased variance on location predictions, a more appropriate sensor setup would have helped considerably.

In this chapter, the bearing used as a case study requires only one polar dimension (two-dimensional in Cartesian coordinates) to be modelled; an angular value. To progress towards larger bearings, consideration of a radial direction will likely be required, extending the problem to two dimensions in a polar coordinate system. Future work should, therefore, investigate the extension towards an additional spatial dimension, where bearings of a greater size can be investigated. As the size of the bearings considered grows (or indeed any structure), the issue of sensor placement will become more important, where providing that one aims to minimise the total number of sensors used, placement must be optimal.

# CONCLUSIONS AND FUTURE WORK

As health monitoring methods continue to mature and develop in response to the desire for improved safety and maintenance scheduling in engineering structures, the ability to localise defects will remain a critical component in such strategies. When equipped with the location that some detected damage has occurred, an operator gains a greater insight into the health state of the system than if its presence is simply detected, enabling more informed maintenance decisions to be made. This may then contribute to reducing operation and maintenance costs, in addition to the (desired) eventual quantification of remaining useful life and prognosis. The development of more robust and informative damage localisation tools is therefore paramount in the progression towards more intelligent structural health monitoring strategies.

In the work contained in this thesis, a novel probabilistic method for localising structural damage in complex and demanding structures through the use of acoustic emission measurements has been developed. By adopting a Bayesian framework, the proposed approach allows a damage (emission) location likelihood to be quantified across the surface of a structure, providing both a most probable estimate of damage location, as well as an uncertainty associated with the prediction. In the framework, the task is to first learn a mapping from spatial coordinates to  $\Delta T$  times, which here, is achieved through the placement of a Gaussian process prior over the onset time functions. Distributions are then able to be recovered across possible  $\Delta T$  times. Upon the observation of an acoustic emission event with an unknown origin, the likelihood of the emission location may be assessed at trial locations, resulting

in the calculation of an emission location likelihood surface over the structure. A number of important extensions are then made to the proposed framework. The main contributions of the thesis are as follows:

1. The development of a model that allows acoustic emission signals to be probabilistically localised in structures with complex geometries and material properties.
2. A statistically principled treatment of uncertainty on the onset time features is presented.
3. The development of a solution for reducing the burden of data acquisition for forming  $\Delta T$  maps through the use of physics-informed machine learning.
4. Extension of the framework to geometries that do not naturally arise in Euclidean geometry, allowing localisation in the context of condition monitoring.

## 7.1 Localisation in complex structures

One of the most significant barriers to the progression of AE localisation strategies is extending beyond structures that are isotropic and have simple geometries that do not have any complex geometrical features that obscure the propagation of the ultrasonic waves. In these cases, analytical laws that describe how acoustic emission waves behave become harder to derive. Finite difference solvers may be called upon, but also come coupled with their own challenges such as ensuring all necessary physical interactions are included in the numerical analysis, as well as in validation. In this thesis, a data-driven approach is presented to overcome this challenge, where the driving mechanism of the propagation behaviour can be indirectly learnt from measured data. By learning a relevant mapping from features derived from these data, the underlying phenomena may be suitably captured, regardless of how complex that behaviour may be (provided there exists sufficient data). Not only does this allow localisation in media where propagation paths are highly nonlinear, but in the possession of a suitable training set, the method can be readily applied to differing structures without requiring there to be any similarity in material properties or geometry.

In Chapter 3, the proposed localisation strategy was applied to a complex plate structure that presents a very challenging localisation environment. A high predictive

accuracy was obtained, outperforming a number of current state-of-the-art approaches that are currently available in the field. Not only were most probable location estimates available, but the method returns a likelihood of emission location map across the full surface of the plate, allowing a degree of (un)certainly to be assigned to the predictions. For instance, where one small region of high likelihood emerges, one can infer that the prediction has been made with high certainty. Conversely, a less certain location estimate will return diffused, or even multiple regions of increased likelihood. Having possession of location likelihoods, will, therefore, provide more insight into the damage localisation than deterministic, single-point estimates. For example, where multiple spots of likely source locations arise, an operator will be able to inspect each of these areas, with the likelihood surface helping inform how large an area needs to be considered for inspection to identify to potential defect. Away from inspection, it may also be possible to utilise the full location likelihood mappings to feed into probability of failure calculations.

Whilst the results in Chapter 3 are very promising, the assumption that the noise variance can be assumed globally constant limits the expressiveness of the model. Additionally, throughout the thesis, the sensor pair models are treated as independent to one another. In reality, there will exist some correlation between different pairings that is a function of the relative position between sensors, particularly between pairs that share a sensor. It is hoped that by including this correlation, the total number of sensor pairings used can be reduced.

## 7.2 Towards a more-informative noise model

When building models, it is often important for one to attempt to capture associated uncertainty. In engineering, there are many sources of uncertainty; measurement noise, irregularities in material properties, model error, amongst many others. By treating the  $\Delta T$  functions as a random variable corrupted with some noise, progression has been made towards representing some of these source of uncertainty within a principled statistical framework.

In Chapter 3, the noise associated with the  $\Delta T$  times is assumed to be drawn from a fixed variance noise process, which enforces a global noise level. Within Chapter 4, this assumption was removed, and replaced with a noise level that is able to vary locally across the input domain. Simply put, transition is made from an input

independent to an input dependent noise model. In the context of  $\Delta T$  mapping, a model with the capacity to account for local noise variation is more reflective of the underlying physical behaviour, and allows for uncertainty that results from the relative path between a source location and sensor to be captured. When propagating this measure into the localisation, predictions are essentially weighted according to how well a particular sensor pair covers a given input location, allowing sensors positioned nearby an input location to contribute more to the combined emission location likelihood. It can be seen that through the use of the heteroscedastic model, the accuracy of the maximum likelihood estimates are improved, as well returning more concentrated likelihood surfaces, which, generally, are more convergent about the true emission origin.

### 7.3 Physics-informed machine learning for AE localisation

In Chapter 3 and Chapter 4, the methods developed for learning the  $\Delta T$  maps are black-box; the two Gaussian process models are pure machine learners, where the ability to learn and make accurate predictions is entirely reliant on the training data fed into the learning process. Where the training data are sufficiently representative of the process one is interested in learning, which in the case of  $\Delta T$  mapping, requires a set of dense onset time measurements that have been collected over the full spatial bounds of the structure, then a black-box learner will have the capacity to do well. However, in situations where the availability of data is reduced, then it is highly likely that a machine learner will struggle. Given that acquiring physical measurements is time-consuming, with it often being impractical to gain access to all areas of a structure in operation, e.g. in systems comprised of many components, it is desirable in  $\Delta T$  mapping to be able to call upon methods that work in scenarios with limited training data.

To this end, Chapter 5 proposes the use of a grey-box model for improving the practicality of constructing  $\Delta T$  maps, where insight derived from physical knowledge is included in the learning process. The nature of the knowledge is that of boundary conditions, where the Gaussian process priors placed over the onset time functions are constrained to satisfy the known behaviour of the  $\Delta T$  times at the spatial boundary of the structure, thus guaranteeing that predictions made at the boundary are consistent

with what occurs physically. It is shown that the constraints allow for an improved predictive performance when learning the  $\Delta T$  maps with training grids with limited spatial coverage, as well as reduced grid densities, therefore, increasing the viability of learning  $\Delta T$  maps. This improved performance is observed through both a reduction in the mean error of the predictions, as well as reducing the associated uncertainty at boundary locations.

As our engineering structures continue to grow in both size and complexity, it is important that the methods developed for monitoring them are scalable. As such, the ability to develop localisation tools in sparse training conditions is important. An additional advantage of the proposed approach is that the level of understanding of the wave propagation one requires is relatively simple, with boundary conditions knowledge generally much easier to acquire than an entire analytical or numerical quantification of the full wavefield. However, boundary constraints are just one type of knowledge an engineer may possess, and additional approaches for including further physical knowledge into modelling  $\Delta T$  maps should be explored. Additionally, the noise model used in the grey-box GP should be extended to the heteroscedastic case.

## 7.4 Localisation beyond Euclidean geometries

An implicit assumption made when using many standard Gaussian process kernels such as the Matérn used in much of this work is that the similarity between input points can be expressed as a function of their Euclidean distance. However, in the real world, we are often interested in components that are spherical. For example, the vast majority of rotating systems and machinery will contain many bearings, all of which possess a spherical topology. It is, therefore, important to extend beyond thinking in Euclidean space. With this in mind, the methodology developed in this thesis is extended in Chapter 6 to include geometries that are not naturally characterised in Euclidean space.

In the chapter, a bearing case study is introduced that was designed such that it replicates a scaled-down wind turbine gearbox bearing. It is demonstrated how by constructing the kernel function as an eigenfunction expansion, the need to specify the covariance between two points as a function of their Euclidean distance is bypassed, and is instead approached through spectral methods. A good localisation

performance in the bearing is obtained, which given that many of the  $\Delta T$  times contain significant variation across repeated measurements, is positive. Additionally, the ability to achieve an accurate two-dimensional location with only three sensors is promising. However, the bearing case study in this chapter was fairly small scale, with the method still to be demonstrated on a larger bearing.

## 7.5 Future work

The work in this thesis has identified many useful and interesting avenues of further possible work. To discuss these extensions, three potential packages of work that may proceed this thesis that have been identified by the author are defined as:

- Efficient learning for localisation models
- Developing physics-informed models for AE localisation
- Scalability of localisation methods

Beginning with the task of learning more efficiently, the premise here is that in the model learning stage, we maximise the insight a machine learner gains regarding the learning task at hand from some fixed amount of training data. In other words, how can we extract the most information from a set amount of collected data to learn a model. For instance, the localisation methodology proposed in this thesis models one Gaussian process for each sensor pair, and then assumes independence between the models. In reality, we know that there will be some correlation in the  $\Delta T$  times between pairs based on how similarly the individual sensors are positioned. Two sensors placed close together will observe very similar propagation paths, and, thus, time of flights. Additionally, many pairings are composed of the same sensors a number of times, introducing another source of correlation between different sensor pairs. To improve the model, this correlation between individual sensor pair models should be captured. For geometries where the propagation paths are generally well understood, it will be possible to explicitly derive correlation expressions between different sensor pairs. For more complex structures such as those of the experimental case study considered in this work, one can attempt to learn these terms in conjunction with the overall GP model [152].

In the current form of the work, uncertainty on the onset times is handled in combination with the uncertainty arising from measurement noise. Whilst the Bayesian methodology pursued here is able to jointly model both sources, it could be of benefit for model interpretability (and possibly performance) to separate the two uncertainty contributions, which could be constructed by switching to a probabilistic onset time identification strategy over AIC.

Another possible avenue of work to explore within this area is the development of a principled strategy for selecting the locations that the training onset time measurements are collected at. We may pose this challenge as: how can onset time measurements be optimally sampled across the structure of interest? Throughout the thesis, the training grids that are considered contain measurements that are equally spaced with respect to adjacent measurements. For non-complex homogenous structures where there is no need to more densely capture measurements in a particular area of the structure, a training grid of this type will likely be close to optimal. However, as components become more complex, particularly in terms of geometrical features that obstruct propagation paths, then it is unlikely that an equispaced training grid will be the best choice, and will either require very dense grids where some areas will be overly sampled, or the behaviour of the onset times will not be sufficiently captured in areas where the wavefield becomes highly nonlinear such as at boundaries between different types of media.

A final route of potential work within the topic of learning more efficiently is the issue of optimal sensor placement. As one may realise, there is much overlap between the problem of optimally sampling over the input space and optimally placing receivers that sense the data. The problem is very similar; given some allocated resources, how do I optimally place my sensors such that the predictive accuracy of the  $\Delta T$  maps is maximised, with the number of sensors used minimised. The advantage here is that a structure may be covered with as few sensors as possible, reducing set up times and costs, whilst still sufficiently reconstructing the  $\Delta T$  fields.

Another area of work that the contributions made in this thesis have opened up is the use of physics-informed machine learning for localisation of damage. An initial and fairly obvious extension is to move the constrained GP model to a heteroscedastic setting, combining both the power of a more expressive noise model with the benefits of the inclusion of boundary conditions knowledge. It is expected that a local noise model would be particularly beneficial to the work of Chapter 6, where a minimal number of sensors are used.

Next, in this thesis, constraining the covariance function of the Gaussian process was considered for including boundary condition information. However, boundary conditions are just one example of understanding that an engineer may have about the propagation behaviour of the onset time functions. When forming a Gaussian process regression, in addition to the covariance, one also has the option to specify a mean function, which may be formed from any number of inputs. This opens up the possibility of guiding the GP predictions independent of training data, either through some analytical expression for the expected behaviour of the ultrasonic waves or through the output of a numerical model. There are then a number of ways this type of model can be interpreted; some examples include increasing the extrapolative power of the Gaussian process predictions, or learning bias corrections in some analytical or numerical model. Exploitation of a mean function to be incorporated into the models developed in this work, therefore, serves as a good candidate to embed further physics into the machine learners.

Finally, more informative expressions for the covariance function can also be considered, which instead of employing kernel functions that have the capacity to model any continuous function, use kernel forms that better represent the autocovariance of the process of interest. For instance, it has already been demonstrated how for single-degree of freedom oscillators, an exact covariance can be derived [136]. Extending such a way of thinking into localisation models would, therefore, be desirable.

A final consideration, for which it is fair to say is an ongoing problem within structural health monitoring more broadly, is how these approaches may be scaled to large structures. With engineering structures growing in size, particularly in the offshore wind sector, there is an increased necessity for health monitoring strategies such as damage localisation methods to be scalable. One perspective of such a task is ensuring that the models developed are computationally cheap, as was discussed in Chapter 5. Additionally, reducing the amount of training data that is necessary will aid training times, with discussions on this matter considered in Chapter 3,5 and this section.

Away from computational scalability, overcoming physical limitations imposed by larger structures is an equal, if not a more important challenge (if one assumes that Moore's law will sufficiently handle the required increase in computational resources, both in processing power and memory) in scaling health monitoring models. From the perspective of high frequency data such as acoustic emission, signal attenuation presents a big barrier to the use of monitoring systems in larger

---

structures. Sensor positioning will be critical in reducing the impact of attenuation, bringing optimal sensor placement into the question. Another approach would be to combine data extracted from phenomena at different frequencies such as low frequency vibration data with that of acoustic emission or other ultrasonic techniques such as ultrasonic-guided waves. The use of different frequency ranges will allow safety critical components/zones more prone to fatigue damage to be more closely monitored through the use of a local network of ultrasonic sensors, whilst the rest of the structure can be monitored with a lower frequency vibration-based strategy, which although is a less sensitive damage indicator, is more practical for large structures. The topic of scalability is, however, a very open one, and is key to seeing a universal adoption of health monitoring methods.



# BASIC BAYESIAN PROBABILITY THEORY

At the core of (Bayesian) probability theory are a number of basic relationships. The first of these is the *sum rule*, stated as,

$$p(\mathbf{x}) = \int p(\mathbf{x}, \mathbf{y}) d\mathbf{y}, \quad (\text{A.1})$$

where  $\mathbf{x}, \mathbf{y}$  are both random variables. Equation (A.1) is also referred to as the marginalisation property in the sense that all possible states of  $\mathbf{y}$  are marginalised (integrated out). Note that in the case of a *discrete* probability distribution, the sum rule may be expressed equivalent as a sum operation over  $\mathbf{y}$ .

The second rule is that of the *product rule*, stated as,

$$p(\mathbf{x}, \mathbf{y}) = p(\mathbf{y}|\mathbf{x})p(\mathbf{x}). \quad (\text{A.2})$$

To derive Bayes' theorem, one begins with the product rule. As

$$p(\mathbf{x}, \mathbf{y}) = p(\mathbf{y}, \mathbf{x}), \quad (\text{A.3})$$

equation (A.2) can be written as,

$$p(\mathbf{y}|\mathbf{x})p(\mathbf{x}) = p(\mathbf{x}|\mathbf{y})p(\mathbf{y}), \quad (\text{A.4})$$

$$p(\mathbf{y}|\mathbf{x}) = \frac{p(\mathbf{x}|\mathbf{y})p(\mathbf{y})}{p(\mathbf{x})}, \quad (\text{A.5})$$

which is Bayes' theorem. For convenience,  $p(\mathbf{x})$  is often rewritten as a combination of the product and sum rule,

$$p(\mathbf{x}) = \int p(\mathbf{x}, \mathbf{y})d\mathbf{y} = \int p(\mathbf{x}|\mathbf{y})p(\mathbf{y})d\mathbf{y}, \quad (\text{A.6})$$

returning,

$$p(\mathbf{y}|\mathbf{x}) = \frac{p(\mathbf{x}|\mathbf{y})p(\mathbf{y})}{\int p(\mathbf{x}|\mathbf{y})p(\mathbf{y})d\mathbf{y}}. \quad (\text{A.7})$$

---

## BIBLIOGRAPHY

- [1] P. J. Shull. *Nondestructive Evaluation: Theory, Techniques, and Applications*. CRC press, 2002.
- [2] C. J. Hellier. *Handbook of Nondestructive Evaluation*. McGraw-Hill Education, 2013.
- [3] C. U. Grosse and M. Ohtsu. *Acoustic emission testing*. Springer Science & Business Media, 2008.
- [4] C. R. Farrar and K. Worden. *Structural Health Monitoring: A Machine Learning Perspective*. John Wiley & Sons, 2012.
- [5] A. Deraemaeker and K. Worden. *New trends in vibration based structural health monitoring*, volume 520. Springer Science & Business Media, 2012.
- [6] 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.
- [7] A. J. Hughes, L. A. Bull, P. Gardner, R. J. Barthorpe, N. Dervilis, and K. Worden. On risk-based active learning for structural health monitoring. *Mechanical Systems and Signal Processing*, 167:108569, 2022.
- [8] J. Hensman, K. Worden, M. Eaton, R. Pullin, K. Holford, and S. Evans. Spatial scanning for anomaly detection in acoustic emission testing of an aerospace structure. *Mechanical Systems and Signal Processing*, 25(7):2462–2474, 2011.
- [9] R. Fuentes. *On Bayesian Networks for Structural Health and Condition Monitoring*. PhD thesis, University of Sheffield, 2017.

- 
- [10] M. R. Jones, T. J. Rogers, N. Dervilis, K. Worden, and E. J. Cross. Assessing the likelihood of damage at the start of a structural health monitoring campaign. *Structural Health Monitoring 2019*, 2019.
- [11] T. Söderström and P. Stoica. *System Identification*. Prentice-Hall, 1988.
- [12] T. B. Schön, A. Wills, and B. Ninness. System identification of nonlinear state-space models. *Automatica*, 47(1):39–49, 2011.
- [13] P. Green and K. Worden. Bayesian and Markov chain Monte Carlo methods for identifying nonlinear systems in the presence of uncertainty. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373(2051):20140405, 2015.
- [14] J.-P. Noël and G. Kerschen. Nonlinear system identification in structural dynamics: 10 more years of progress. *Mechanical Systems and Signal Processing*, 83:2–35, 2017.
- [15] T. J. Rogers. *Towards Bayesian System Identification: With Application to SHM of Offshore Structures*. PhD thesis, University of Sheffield, 2019.
- [16] A. Rytter. *Vibrational Based Inspection of Civil Engineering Structures*. 1993.
- [17] K. Worden and J. M. Dulieu-Barton. An overview of intelligent fault detection in systems and structures. *Structural Health Monitoring*, 3(1):85–98, 2004.
- [18] K. Worden and G. Manson. The application of machine learning to structural health monitoring. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1851):515–537, 2007.
- [19] G. Park, A. C. Rutherford, H. Sohn, and C. R. Farrar. An outlier analysis framework for impedance-based structural health monitoring. *Journal of Sound and Vibration*, 286(1-2):229–250, 2005.
- [20] P. Rizzo, M. Cammarata, D. Dutta, H. Sohn, and K. Harries. An unsupervised learning algorithm for fatigue crack detection in waveguides. *Smart Materials and Structures*, 18(2):025016, 2009.
- [21] J. Toivola, M. A. Prada, and J. Hollmén. Novelty detection in projected spaces for structural health monitoring. In *International Symposium on Intelligent Data Analysis*, pages 208–219. Springer, 2010.

- [22] N. Dervilis, E. Cross, R. Barthorpe, and K. Worden. Robust methods of inclusive outlier analysis for structural health monitoring. *Journal of Sound and Vibration*, 333(20):5181–5195, 2014.
- [23] L. Bull, K. Worden, R. Fuentes, G. Manson, E. Cross, and N. Dervilis. Outlier ensembles: A robust method for damage detection and unsupervised feature extraction from high-dimensional data. *Journal of Sound and Vibration*, 453:126–150, 2019.
- [24] L. A. Bull. *Towards Probabilistic and Partially-Supervised Structural Health Monitoring*. PhD thesis, University of Sheffield, 2019.
- [25] E. Cross. *On Structural Health Monitoring in Changing Environmental and Operational Conditions*. PhD thesis, University of Sheffield, 2012.
- [26] E. J. Cross, K. Worden, and Q. Chen. Cointegration: a novel approach for the removal of environmental trends in structural health monitoring data. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2133):2712–2732, 2011.
- [27] A. Santos, E. Figueiredo, M. Silva, C. Sales, and J. Costa. Machine learning algorithms for damage detection: Kernel-based approaches. *Journal of Sound and Vibration*, 363:584–599, 2016.
- [28] T. Rogers, K. Worden, R. Fuentes, N. Dervilis, U. Tygesen, and E. Cross. A Bayesian non-parametric clustering approach for semi-supervised structural health monitoring. *Mechanical Systems and Signal Processing*, 119:100–119, 2019.
- [29] L. Bull, T. Rogers, C. Wickramarachchi, E. Cross, K. Worden, and N. Dervilis. Probabilistic active learning: an online framework for structural health monitoring. *Mechanical Systems and Signal Processing*, 134:106294, 2019.
- [30] M. G. Baxter, R. Pullin, K. M. Holford, and S. L. Evans. Delta T source location for acoustic emission. *Mechanical Systems and Signal Processing*, 21(3):1512–1520, 2007.
- [31] J. Hensman, R. Mills, S. Pierce, K. Worden, and M. Eaton. Locating acoustic emission sources in complex structures using Gaussian processes. *Mechanical Systems and Signal Processing*, 24(1):211–223, 2010.

- 
- [32] A. Ebrahimkhanlou, B. Dubuc, and S. Salamone. A generalizable deep learning framework for localizing and characterizing acoustic emission sources in riveted metallic panels. *Mechanical Systems and Signal Processing*, 130:248–272, 2019.
- [33] J. Kaiser. *Untersuchungen über das Auftreten von Geräuschen beim Zugversuch*. PhD thesis, Technische Hochschule München, 1950.
- [34] K. Ono. Acoustic emission. *Encyclopedia of Acoustics*, 68:797–809, 1997.
- [35] J. L. Rose. Ultrasonic guided waves in structural health monitoring. In *Key Engineering Materials*, volume 270, pages 14–21. Trans Tech Publ, 2004.
- [36] K. Worden. Rayleigh and lamb waves-basic principles. *Strain*, 37(4):167–172, 2001.
- [37] 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.
- [38] M. Haywood-Alexander. *Development of Novel Tools for Application of Ultrasonic Guided Waves in Fibre-Composites*. PhD thesis, University of Sheffield, 2022.
- [39] H. Vallen. AE Testing Fundamentals, Equipment, Applications. *The e-Journal of Nondestructive Testing*, 7(9), 2002.
- [40] P. D. Wilcox, C. Lee, J. J. Scholey, M. I. Friswell, M. Wisnom, and B. Drinkwater. Progress towards a forward model of the complete acoustic emission process. In *Advanced Materials Research*, volume 13, pages 69–76. Trans Tech Publ, 2006.
- [41] C. B. Scruby. An introduction to acoustic emission. *Journal of Physics E: Scientific Instruments*, 20(8):946, 1987.
- [42] Z. Nazarchuk, V. Skalskyi, and O. Serhiyenko. Acoustic Emission. *Foundations of Engineering Mechanics*, 2017.
- [43] S. Grigg. *Development of a Wireless Structural Health Monitoring System for Aerospace Application*. PhD thesis, Cardiff University, 2018.
- [44] B. Eftekharnajad, M. Carrasco, B. Charnley, and D. Mba. The application of spectral kurtosis on acoustic emission and vibrations from a defective bearing. *Mechanical Systems and Signal Processing*, 25(1):266–284, 2011.

- [45] M. Elforjani and D. Mba. Accelerated natural fault diagnosis in slow speed bearings with acoustic emission. *Engineering Fracture Mechanics*, 77(1):112–127, 2010.
- [46] D. Ozevin and J. Harding. Novel leak localization in pressurized pipeline networks using acoustic emission and geometric connectivity. *International Journal of Pressure Vessels and Piping*, 92:63–69, 2012.
- [47] N. N. Nelson. Acoustic emissions simulator, US Patent 4,018,084, April 19 1977.
- [48] A. Nielsen. *Acoustic emission source based on pencil lead breaking*. 1980.
- [49] K. Tamada and H. Tanaka. Occurrence of brittle flaking on bearings used for automotive electrical instruments and auxiliary devices. *Wear*, 199(2):245–252, 1996.
- [50] R. Fuentes, R. Dwyer-Joyce, M. Marshall, J. Wheals, and E. Cross. Detection of sub-surface damage in wind turbine bearings using acoustic emissions and probabilistic modelling. *Renewable Energy*, 147:776–797, 2020.
- [51] L. Rayleigh. On waves propagated along the plane surface of an elastic solid. *Proceedings of the London Mathematical Society*, 1(1):4–11, 1885.
- [52] H. Lamb. On waves in an elastic plate. *Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character*, 93(648):114–128, 1917.
- [53] B. Pavlakovic, M. Lowe, D. Alleyne, and P. Cawley. Disperse: A general purpose program for creating dispersion curves. In *Review of Progress in Quantitative Nondestructive Evaluation*, pages 185–192. Springer, 1997.
- [54] J. L. Rose. *Ultrasonic waves in solid media*, 2000.
- [55] J. L. Rose. *Ultrasonic guided waves in solid media*. Cambridge university press, 2014.
- [56] M. J. Lowe. Matrix techniques for modeling ultrasonic waves in multilayered media. *IEEE transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(4):525–542, 1995.

- 
- [57] J. J. Hensman. *Novel Techniques for Acoustic Emission Monitoring of Fatigue Fractures in Landing Gear*. PhD thesis, University of Sheffield, 2010.
- [58] J. Naumann. *Acoustic Emission Monitoring of Wind Turbine Bearings*. PhD thesis, University of Sheffield, 2016.
- [59] M. Baxter. *Damage Assessment by Acoustic Emission (AE) During Landing Gear Fatigue Testing*. Cardiff University (United Kingdom), 2007.
- [60] G. Kitagawa and H. Akaike. A procedure for the modeling of non-stationary time series. *Annals of the Institute of Statistical Mathematics*, 30(2):351–363, 1978.
- [61] Y. Morita and H. Hamaguchi. Automatic detection of onset time of seismic waves and its confidence interval using autoregressive model fitting. *Journal of the Seismological Society of Japan*, 1984.
- [62] N. Maeda. A method for reading and checking phase times in autoprocessing system of seismic wave data. *Zisin*, 38:365–379, 1985.
- [63] J. H. Kurz, C. U. Grosse, and H.-W. Reinhardt. Strategies for reliable automatic onset time picking of acoustic emissions and of ultrasound signals in concrete. *Ultrasonics*, 43(7):538–546, 2005.
- [64] B. Schechinger and T. Vogel. Acoustic emission for monitoring a reinforced concrete beam subject to four-point-bending. *Construction and Building Materials*, 21(3):483–490, 2007.
- [65] H. Zhang, C. Thurber, and C. Rowe. Automatic p-wave arrival detection and picking with multiscale wavelet analysis for single-component recordings. *Bulletin of the Seismological Society of America*, 93(5):1904–1912, 2003.
- [66] Z. Zhou, R. Cheng, Y. Rui, J. Zhou, and H. Wang. An improved automatic picking method for arrival time of acoustic emission signals. *IEEE Access*, 7: 75568–75576, 2019.
- [67] T. Kundu, H. Nakatani, and N. Takeda. Acoustic source localization in anisotropic plates. *Ultrasonics*, 52(6):740–746, 2012.
- [68] A. Tobias. Acoustic-emission source location in two dimensions by an array of three sensors. *Non-Destructive Testing*, 9(1):9–12, 1976.

- [69] F. Ciampa and M. Meo. Acoustic emission localization in complex dissipative anisotropic structures using a one-channel reciprocal time reversal method. *The Journal of the Acoustical Society of America*, 130(1):168–175, 2011.
- [70] R. Ernst, F. Zwimpfer, and J. Dual. One sensor acoustic emission localization in plates. *Ultrasonics*, 64:139–150, 2016.
- [71] A. Ebrahimkhanlou and S. Salamone. A probabilistic framework for single-sensor acoustic emission source localization in thin metallic plates. *Smart materials and structures*, 26(9):095026, 2017.
- [72] A. Ebrahimkhanlou and S. Salamone. Single-sensor acoustic emission source localization in plate-like structures using deep learning. *Aerospace*, 5(2):50, 2018.
- [73] L. M. Rios and N. V. Sahinidis. Derivative-free optimization: a review of algorithms and comparison of software implementations. *Journal of Global Optimization*, 56(3):1247–1293, 2013.
- [74] R. Miller, R. Carlos, R. Findlay, V. Godinez-Azcuaga, M. Rhodes, F. Shu, and W. Wang. Acoustic emission source location. *ASNT*, pages 121–146–127, 2005.
- [75] H. Yamada, Y. Mizutani, H. Nishino, M. Takemoto, and K. Ono. Lamb wave source location of impact on anisotropic plates. *Journal of Acoustic Emission*, 18:51, 2000.
- [76] T. Kundu, S. Das, S. A. Martin, and K. V. Jata. Locating point of impact in anisotropic fiber reinforced composite plates. *Ultrasonics*, 48(3):193–201, 2008.
- [77] H. Nakatani, T. Hajzargarbashi, K. Ito, T. Kundu, and N. Takeda. Impact localization on a cylindrical plate by near-field beamforming analysis. In *Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems*, volume 8345, page 83450Y. International Society for Optics and Photonics, 2012.
- [78] F. Ciampa and M. Meo. A new algorithm for acoustic emission localization and flexural group velocity determination in anisotropic structures. *Composites Part A: Applied Science and Manufacturing*, 41(12):1777–1786, 2010.
- [79] R. K. Ing, N. Quieffin, S. Catheline, and M. Fink. In solid localization of finger impacts using acoustic time-reversal process. *Applied Physics Letters*, 87(20):204104, 2005.

- 
- [80] B. Park, H. Sohn, S. E. Olson, M. P. DeSimio, K. S. Brown, and M. M. Derriso. Impact localization in complex structures using laser-based time reversal. *Structural Health Monitoring*, 11(5):577–588, 2012.
- [81] F. Ciampa and M. Meo. Impact detection in anisotropic materials using a time reversal approach. *Structural Health Monitoring*, 11(1):43–49, 2012.
- [82] K. M. Holford and D. Carter. Acoustic emission source location. In *Key Engineering Materials*, volume 167, pages 162–171. Trans Tech Publ, 1999.
- [83] H. Achdjian, E. Moulin, F. Benmeddour, J. Assaad, and L. Chehami. Source localisation in a reverberant plate using average coda properties and early signal strength. *Acta Acustica united with Acustica*, 100(5):834–841, 2014.
- [84] K. M. Holford, M. J. Eaton, J. J. Hensman, R. Pullin, S. L. Evans, N. Dervilis, and K. Worden. A new methodology for automating acoustic emission detection of metallic fatigue fractures in highly demanding aerospace environments: An overview. *Progress in Aerospace Sciences*, 90:1–11, 2017.
- [85] M. J. Eaton, R. Pullin, and K. M. Holford. Acoustic emission source location in composite materials using Delta T mapping. *Composites Part A: Applied Science and Manufacturing*, 43(6):856–863, 2012.
- [86] S. K. Al-Jumaili, M. R. Pearson, K. M. Holford, M. J. Eaton, and R. Pullin. Acoustic emission source location in complex structures using full automatic Delta T mapping technique. *Mechanical Systems and Signal Processing*, 72: 513–524, 2016.
- [87] T. Schumacher, D. Straub, and C. Higgins. Toward a probabilistic acoustic emission source location algorithm: a Bayesian approach. *Journal of Sound and Vibration*, 331(19):4233–4245, 2012.
- [88] E. D. Niri and S. Salamone. A probabilistic framework for acoustic emission source localization in plate-like structures. *Smart Materials and Structures*, 21(3):035009, 2012.
- [89] D. Sen, K. Erazo, and S. Nagarajaiah. Bayesian estimation of acoustic emissions source in plate structures using particle-based stochastic filtering. *Structural Control and Health Monitoring*, 24(11):e2005, 2017.

- [90] R. Madarshahian, P. Ziehl, and J. M. Caicedo. Acoustic emission Bayesian source location: onset time challenge. *Mechanical Systems and Signal Processing*, 123:483–495, 2019.
- [91] E. D. Niri, A. Farhidzadeh, and S. Salamone. Nonlinear Kalman filtering for acoustic emission source localization in anisotropic panels. *Ultrasonics*, 54(2):486–501, 2014.
- [92] C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*. Number 3. MIT press Cambridge, MA, 2006.
- [93] T. Rogers, P. Gardner, N. Dervilis, K. Worden, A. Maguire, E. Papatheou, and E. Cross. Probabilistic modelling of wind turbine power curves with application of heteroscedastic Gaussian process regression. *Renewable Energy*, 148:1124–1136, 2020.
- [94] R. Fuentes, E. Cross, A. Halfpenny, K. Worden, and R. J. Barthorpe. Aircraft parametric structural load monitoring using Gaussian process regression. In *7th European Workshop on Structural Health Monitoring*, 2014.
- [95] M. Civera, C. Surace, and K. Worden. Detection of cracks in beams using Treed Gaussian processes. In *Structural Health Monitoring & Damage Detection, Volume 7*, pages 85–97. Springer, 2017.
- [96] R. B. Gramacy. *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. Chapman and Hall/CRC, 2020.
- [97] P. Jiang, Q. Zhou, and X. Shao. *Surrogate model-based engineering design and optimization*. Springer, 2020.
- [98] R. K. Pandit, D. Infield, and J. Carroll. Incorporating air density into a Gaussian process wind turbine power curve model for improving fitting accuracy. *Wind Energy*, 22(2):302–315, 2019.
- [99] L. D. Avendano-Valencia, E. N. Chatzi, and D. Tcherniak. Gaussian process models for mitigation of operational variability in the structural health monitoring of wind turbines. *Mechanical Systems and Signal Processing*, 142:106686, 2020.
- [100] K. Chandrasekhar, N. Stevanovic, E. J. Cross, N. Dervilis, and K. Worden. Damage detection in operational wind turbine blades using a new approach based on machine learning. *Renewable Energy*, 168:1249–1264, 2021.

- 
- [101] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC, 1995.
- [102] P. M. Lee. *Bayesian Statistics*. Oxford University Press London:, 1989.
- [103] K. P. Murphy. *Machine Learning: A Probabilistic Perspective*. MIT press, 2012.
- [104] A. Terenin. Gaussian Processes and Statistical Decision-making in Non-Euclidean spaces. *arXiv preprint arXiv:2202.10613*, 2022.
- [105] M. L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Science & Business Media, 2012.
- [106] D. Duvenaud. *Automatic Model Construction with Gaussian Processes*. PhD thesis, University of Cambridge, 2014.
- [107] C. Rasmussen and Z. Ghahramani. Occam’s razor. *Advances in Neural Information Processing Systems*, 13, 2000.
- [108] R. Fletcher. *Practical Methods of Optimization*. John Wiley & Sons, 2013.
- [109] A. P. Engelbrecht. *Computational Intelligence: An Introduction*. John Wiley & Sons, 2007.
- [110] D. H. Wolpert and W. G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, 1997.
- [111] J. Sun, W. Xu, and B. Feng. A global search strategy of quantum-behaved particle swarm optimization. In *IEEE Conference on Cybernetics and Intelligent Systems, 2004.*, volume 1, pages 111–116. IEEE, 2004.
- [112] J. Sun, B. Feng, and W. Xu. Particle swarm optimization with particles having quantum behavior. In *Proceedings of the 2004 Congress on Evolutionary Computation (IEEE Cat. No. 04TH8753)*, volume 1, pages 325–331. IEEE, 2004.
- [113] K. Worden, R. Barthorpe, E. Cross, N. Dervilis, G. Holmes, G. Manson, and T. Rogers. On evolutionary system identification with applications to nonlinear benchmarks. *Mechanical Systems and Signal Processing*, 112:194–232, 2018.

- 
- [114] E. B. Flynn and M. D. Todd. A Bayesian approach to optimal sensor placement for structural health monitoring with application to active sensing. *Mechanical Systems and Signal Processing*, 24(4):891–903, 2010.
- [115] M. A. Vega and M. D. Todd. A variational Bayesian neural network for structural health monitoring and cost-informed decision-making in miter gates. *Structural Health Monitoring*, page 1475921720904543, 2020.
- [116] C. Brooks, S. P. Burke, and G. Persand. Benchmarks and the accuracy of garch model estimation, 2001.
- [117] C. T. Brownlees, R. F. Engle, and B. T. Kelly. A practical guide to volatility forecasting through calm and storm. *Available at SSRN 1502915*, 2011.
- [118] H. Sun, N. Zhao, X. Zeng, and D. Yan. Study of solar radiation prediction and modeling of relationships between solar radiation and meteorological variables. *Energy Conversion and Management*, 105:880–890, 2015.
- [119] W. Jiang, Z. Yan, D.-H. Feng, and Z. Hu. Wind speed forecasting using autoregressive moving average/generalized autoregressive conditional heteroscedasticity model. *European Transactions on Electrical Power*, 22(5):662–673, 2012.
- [120] P. Goldberg, C. Williams, and C. Bishop. Regression with input-dependent noise: A Gaussian process treatment. *Advances in Neural Information Processing Systems*, 10, 1997.
- [121] K. Kersting, C. Plagemann, P. Pfaff, and W. Burgard. Most likely heteroscedastic gaussian process regression. In *Proceedings of the 24th international conference on Machine learning*, pages 393–400, 2007.
- [122] M. Lázaro-Gredilla and M. K. Titsias. Variational heteroscedastic Gaussian process regression. 2011.
- [123] T. Salimans, D. Kingma, and M. Welling. Markov chain Monte Carlo and variational inference: bridging the gap. In *International Conference on Machine Learning*, pages 1218–1226. PMLR, 2015.
- [124] D. Chan, R. Kohn, D. Nott, and C. Kirby. Locally adaptive semiparametric estimation of the mean and variance functions in regression models. *Journal of Computational and Graphical Statistics*, 15(4):915–936, 2006.

- 
- [125] F. M. Dekking, C. Kraaikamp, H. P. Lopuhaä, and L. E. Meester. *A Modern Introduction to Probability and Statistics: Understanding why and how*, volume 488. Springer, 2005.
- [126] J. Willard, X. Jia, S. Xu, M. Steinbach, and V. Kumar. Integrating physics-based modeling with machine learning: A survey. *arXiv preprint arXiv:2003.04919*, 2020.
- [127] E. J. Cross, S. J. Gibson, M. R. Jones, D. J. Pitchforth, S. Zhang, and T. J. Rogers. *Structural Health Monitoring based on Data Science Techniques*, chapter Physics-informed machine learning for Structural Health Monitoring. Springer, 2022.
- [128] L. P. Swiler, M. Gulian, A. L. Frankel, C. Safta, and J. D. Jakeman. A survey of constrained Gaussian process regression: Approaches and implementation challenges. *Journal of Machine Learning for Modeling and Computing*, 1(2), 2020.
- [129] C. Agrell. Gaussian processes with linear operator inequality constraints. *arXiv preprint arXiv:1901.03134*, 2019.
- [130] J. Riihimäki and A. Vehtari. Gaussian processes with monotonicity information. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 645–652. JMLR Workshop and Conference Proceedings, 2010.
- [131] S. Da Veiga and A. Marrel. Gaussian process modeling with inequality constraints. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 21, pages 529–555, 2012.
- [132] H. Maatouk. Finite-dimensional approximation of Gaussian processes with inequality constraints. *arXiv preprint arXiv:1706.02178*, 2017.
- [133] A. F. López-Lopera, F. Bachoc, N. Durrande, and O. Roustant. Finite-dimensional Gaussian approximation with linear inequality constraints. *SIAM/ASA Journal on Uncertainty Quantification*, 6(3):1224–1255, 2018.
- [134] A. Pensoneault, X. Yang, and X. Zhu. Nonnegativity-enforced Gaussian process regression. *Theoretical and Applied Mechanics Letters*, 10(3):182–187, 2020.

- [135] C. Jidling, N. Wahlström, A. Wills, and T. B. Schön. Linearly constrained Gaussian processes. *arXiv preprint arXiv:1703.00787*, 2017.
- [136] E. J. Cross and T. J. Rogers. Physics-derived covariance functions for machine learning in structural dynamics. *IFAC-PapersOnLine*, 54(7):168–173, 2021.
- [137] M. Alvarez, D. Luengo, and N. D. Lawrence. Latent force models. In *Artificial Intelligence and Statistics*, pages 9–16. PMLR, 2009.
- [138] M. Raissi and G. E. Karniadakis. Hidden physics models: Machine learning of nonlinear partial differential equations. *Journal of Computational Physics*, 357:125–141, 2018.
- [139] E. J. Cross, T. J. Rogers, and T. J. Gibbons. Grey-box modelling for structural health monitoring: physical constraints on machine learning algorithms. *Structural Health Monitoring 2019*, 2019.
- [140] A. Solin and S. Särkkä. Hilbert space methods for reduced-rank Gaussian process regression. *Statistics and Computing*, 30(2):419–446, 2020.
- [141] C. Chatfield and H. Xing. *The Analysis of Time Series: An Introduction with R*. CRC press, 2019.
- [142] D. Pitchforth, T. Rogers, U. Tygesen, and E. Cross. Grey-box models for wave loading prediction. *Mechanical Systems and Signal Processing*, 159:107741, 2021.
- [143] S. Zhang, T. J. Rogers, and E. J. Cross. Gaussian process based grey-box modelling for shm of structures under fluctuating environmental conditions. In *European Workshop on Structural Health Monitoring*, pages 55–66. Springer, 2020.
- [144] R. J. LeVeque. *Finite difference methods for ordinary and partial differential equations: steady-state and time-dependent problems*. SIAM, 2007.
- [145] A. Solin and M. Kok. Know your boundaries: Constraining gaussian processes by variational harmonic features. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2193–2202. PMLR, 2019.
- [146] J. Ribrant and L. Bertling. Survey of failures in wind power systems with focus on swedish wind power plants during 1997-2005. In *2007 IEEE Power Engineering Society general meeting*, pages 1–8. IEEE, 2007.

- 
- [147] L. Rogers. The application of vibration signature analysis and acoustic emission source location to on-line condition monitoring of anti-friction bearings. *Tribology International*, 12(2):51–58, 1979.
- [148] I. Martinez. *A Machine Learning Approach for Damage Detection and Localisation in Wind Turbine Gearbox Bearings*. PhD thesis, University of Sheffield, 2020.
- [149] A. Feragen, F. Lauze, and S. Hauberg. Geodesic exponential kernels: When curvature and linearity conflict. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3032–3042, 2015.
- [150] H. König. *Eigenvalue Distribution of Compact Operators*, volume 16. Birkhäuser, 2013.
- [151] A. Reid. *Strobo-Strain: Stroboscopic Neutron and X-ray Strain Measurements in Dynamically Loaded Engineering Components*. PhD thesis, University of Sheffield, 2019.
- [152] M. A. Alvarez, L. Rosasco, N. D. Lawrence, *et al.* Kernels for vector-valued functions: A review. *Foundations and Trends in Machine Learning*, 4(3): 195–266, 2012.