



THE UNIVERSITY OF SHEFFIELD

PhD Mechanical Engineering

On Regression in Population based Structural Health Monitoring

by

Daniel R. Clarkson

July 2026

Dr Aidan J. Hughes

Dr Timothy J. Rogers

Prof Nikolaos Dervilis

Thesis submitted to the School Of Mechanical, Aerospace and Civil Engineering,
University of Sheffield for the degree of Doctor of Philosophy

Acknowledgements

Thank you to my wonderful supervisors Tim, Nikos, and Aidan. You've provided endless support, guidance and encouragement, I couldn't have asked for a better team. Thank you to the rest of the Dynamics Research Group, to have such a concentration of wisdom and kindness has truly been a blessing, I feel very lucky to have studied here. To my friends and family, thank you for keeping my spirits high in the dark winter months, your endless outpouring of love and support made the write-up a breeze. To Riann, thank you for bearing with me through it all, I'm sure it hasn't been easy.

This thesis is dedicated to my parents Guy Clarkson and Moira Macfarlane.

Abstract

Complex engineering infrastructure is commonplace among modern society. Fatigue, erosion, and adverse environmental factors degrade these infrastructure over time, requiring maintenance to restore proper function. Structural Health Monitoring (SHM) provides proactive solutions to ensure structural safety and reliability, SHM systems consist of data acquisition and processing systems to enable the detection of damage in monitored structures. Data-based approaches to SHM have demonstrated considerable promise, but their widespread adoption remains limited by a fundamental challenge: the scarcity of labelled data. Acquiring observations that capture the full spectrum of structural health conditions — from the healthy state through to advanced damage — is costly, disruptive, and often practically unfeasible for in-service assets.

Population-based structural health monitoring has emerged as a principled response to this challenge. By treating a collection of structures jointly rather than as isolated assets, PBSHM enables information to be shared across population members — allowing data-rich assets to inform inference for data-scarce ones. While considerable research effort has been directed toward developing population-based technologies for classification tasks, comparatively little attention has been paid to regression, where the quantity of interest is a continuous health variable rather than a discrete damage category. To continue the development of PBSHM systems, regression based technologies are required. Regression further enables damage severity to be quantified, degradation trajectories to be tracked, and predictions to be made about the remaining useful life of a structure. These capabilities are essential for the informed, cost-effective operation and maintenance of structural assets, and their development for population-based settings is the central concern of this thesis.

Contents

Acknowledgements	i
Abstract	ii
Publications	vi
1 Introduction	1
1.1 Structural Health Monitoring	2
1.1.0.1 Open Challenges in SHM	3
1.1.0.2 Data Scarcity in SHM	3
1.2 Population-based Structural Health Monitoring	4
1.2.1 Transfer Learning	5
1.2.2 Active Learning	5
1.3 Uncertainty in SHM	6
1.4 Thesis Contributions	7
1.5 Thesis Outline	8
2 Background Theory	10
2.1 Structural Health Monitoring	10
2.1.1 SHM approaches	11
2.1.2 Unsupervised learning	12
2.1.3 Supervised learning	12
2.1.4 Partially supervised learning	13
2.2 Probability Theory	14
2.2.1 Bayes Theorem	15
2.2.2 Prior distributions	16
2.2.3 The likelihood	17
2.2.4 Posterior distribution	17
2.3 Problems in SHM	17
2.4 Population-based SHM	18
2.4.1 Homogeneous and heterogeneous populations	18
2.4.2 Similarity quantification in heterogeneous populations	20
2.5 Transfer Learning and Multi-task Learning	23
2.5.1 Multi-task Learning	24
2.5.2 Applications of Transfer Learning and Multi-task Learning	25
2.5.2.1 General Machine Learning Applications	25

2.5.2.2	Engineering Applications	26
2.5.2.3	Transfer Learning in Structural Health Monitoring: PB-SHM	27
2.6	Active Learning	28
2.6.1	Batch, Pool, and Online Active Learning	29
2.6.1.1	General Machine Learning Applications	29
2.6.1.2	Engineering Applications	30
2.6.2	Active Learning in Engineering and Structural Health Monitoring	30
2.6.3	Active Learning for Regression	31
2.7	Conclusion	31
3	Literature Survey: Transfer Learning and Active Learning in SHM	32
3.1	Transfer Learning	32
3.2	Active Learning	36
3.2.1	Optimal Design of Experiments	36
3.2.2	Surrogate modelling	37
3.3	Summary	37
4	Active Learning for Regression in Engineering Populations	39
4.1	Hierarchical Bayesian modelling	40
4.1.1	Decision Theory	42
4.2	Active Learning	44
4.2.1	Risk-based Active Learning	45
4.3	Case Study - A Population of Machining Tools	45
4.3.1	The dataset	45
4.3.2	The hierarchical model	47
4.3.3	Decision Process and Active Learning	52
4.4	Results	56
4.5	Discussion	59
4.6	Conclusion	60
5	Continuous Damage in Engineering Populations	62
5.1	Experiment Design	63
5.1.1	Initial Design	64
5.1.2	The beams	64
5.1.3	The clamping force	66
5.1.4	Final set-up	66
5.2	Data Processing	67
5.3	First analysis	68
5.3.1	The input	69
5.3.2	Frequency Response Functions	70
5.3.3	Modal Analysis	72
5.3.3.1	Natural Frequencies	72
5.3.3.2	Damping Ratios	74
5.3.3.3	Mode Shapes	75
5.4	Damage features	78
5.4.1	Principal component analysis	78

5.4.2	Limitations of the experiment	80
5.5	Conclusion	80
6	Transfer with continuous health states	83
6.1	Principal Component Analysis	83
6.1.1	PCA in PBSHM	85
6.2	Naive PCA	86
6.2.1	Fitting a model to the source domain	86
6.2.2	Finding a mapping	90
6.2.2.1	Pre-processing	91
6.2.3	Learning the rotation	92
6.2.4	PCA stability	93
6.3	Shared Feature Space	96
6.4	Non-dimensionalised frequency	97
6.5	Population-based PCA	100
6.5.1	Personalised PCA in PBSHM	103
6.6	Discussions and Conclusions	103
6.6.1	The Instability of Data-Driven Alignment in Sparse Regimes . . .	105
6.6.2	Non-dimensionalisation as a stabilisation mechanism	105
6.6.3	Decoupling shared and personal behaviours	106
6.7	Conclusion	106
7	Conclusions and Future Work	108
7.1	Summary and Conclusions	108
7.2	Limitations, Challenges and Future Works	110
7.3	Closing remarks	113
A	Chapter 5 - additional material	115
	Bibliography	118

Publications

Journal Publications

Clarkson, D.R., Bull, L.A., Dardeno, T.A., Wickramarachchi, C.T., Cross, E.J., Rogers, T.J., Worden, K., Dervilis, N. and Hughes, A.J., 2025. Active learning for regression in engineering populations: a risk-informed approach. *Data-Centric Engineering*, 6, p.e16.

Conference Publications

Clarkson, D.R., Bull, L.A., Dardeno, T.A., Wickramarachchi, C.T., Cross, E.J., Rogers, T.J., Worden, K., Dervilis, N. and Hughes, A.J., “Sharing Information Between Machine Tools to Improve Surface Finish Forecasting” . In *Proceedings of the Fourteenth International Workshop on Structural Health Monitoring (IWSHM)*, 12-14 September, 2023

Clarkson, D.R., Bull, L.A., Dardeno, T.A., Wickramarachchi, C.T., Cross, E.J., Rogers, T.J., Worden, K., Dervilis, N. and Hughes, A.J., “A spin on active learning analysis for health monitoring”. In *Proceedings of the 11th European Workshop on Structural Health Monitoring (EWSHM 2024)*, 10-13 June, 2024

Clarkson, D.R., Rogers, T.J., Worden, K., Dervilis, N. and Hughes, A.J., “Domain Adaptation for Regression in SHM”. IN *10th International Conference on Computational Methods in Structural Dynamics and Earthquake Engineering (COMPDYN 2025)*, 15-18 June, 2025

Chapter 1

Introduction

Modern society is fundamentally reliant on the safe, continuous operation of highly complex engineering infrastructure. Due to fatigue, degradation, or adverse environmental factors, these infrastructures can become damaged over time, requiring maintenance to restore proper function. Maintenance is a foundational aspect of engineering practice that underpins this reliability across diverse sectors such as manufacturing, transportation, civil infrastructure, and energy. Over previous decades, this discipline has evolved significantly. Traditionally, maintenance relied on reactive measures—where action is taken only after a failure occurs, often resulting in significant downtime and unanticipated costs.

To reduce the likelihood of failure, the industry shifted toward scheduled maintenance, which introduces physical inspections and interventions at fixed time intervals, regardless of the actual condition of the asset. Although effective, this approach can sometimes lead to unnecessary servicing. In the early stages of operational life, regular servicing may not be required at all, and towards the end of life, the fixed intervals may miss critical degradation. Today, the drive for improved efficiency and reduced operational costs is leading engineering fields toward condition-based and predictive maintenance. These approaches rely on measurable indicators of system health and the ability to anticipate degradation, allowing interventions to be performed at the most opportune time.

This transition to predictive maintenance is no longer just an operational preference; it is a necessity. The global engineering landscape has reached a scale where traditional scheduled maintenance is no longer physically or economically feasible. This is starkly visible in the renewable energy sector; for instance, global wind power capacity surpassed 1.2 Terawatts by the end of 2024 [1], with over 117 GW of new capacity installed globally in that year alone [2]. Sending human crews to physically inspect every asset in

these rapidly expanding, distributed networks—often located in hostile offshore environments—is prohibitively expensive and inherently dangerous. Consequently, the reliance on human-centric, scheduled inspections must be replaced by automated systems capable of assessing asset integrity remotely and continuously.

1.1 Structural Health Monitoring

Structural health monitoring (SHM) offers solutions to dated inspection methods to ensure the safety and reliability of various infrastructure. The aim of SHM is to help inform decision-making, particularly for the operation and maintenance (O&M) of high-value and safety-critical infrastructure. SHM systems consist of data acquisition and processing systems to enable the detection of damage in monitored structures. The level of information provided by an SHM system depends on the specific system and structure being monitored. Rytter [3] provides a hierarchy of information that an SHM system might be able to provide, which was later extended by Worden [4].

- **Detection:** Is there damage?
- **Localization:** Where is the damage?
- **Classification:** What kind of damage is it?
- **Quantification:** How severe is the damage?
- **Prognosis:** What does this damage mean for the remaining life of the structure?

Each level in the hierarchy generally requires more information, and this information is generally more difficult to acquire. Damage is typically difficult or impossible to measure directly [5], sensors often capture indirect quantities—such as vibrations, strain measurements, or images—rather than direct structural flaws, meaning that SHM systems must rely on models to interpret these data. These models exist on a spectrum: at one extreme, physics-based models [6–8] that use the underlying laws of physics to describe the observed behaviours of the system. At the other extreme are data-based models [9–12] that learn the relationship between sensor measurements and structural damage directly from data. In-between these two approaches are *physics-informed* [13–18] approaches that combine physical laws and data-based approaches.

1.1.0.1 Open Challenges in SHM

Several significant research challenges remain for SHM systems, which can be broadly categorised into three areas. First, the acquisition of data and extraction of damage-sensitive features remains a largely bespoke procedure, requiring specialist knowledge to design sensor systems and interpret the resulting measurements. Feature extraction frequently relies on manual processes such as modal analysis, and identifying features that are sensitive to a range of damage locations, types, and severities remains an open problem. Second, confounding influences — most notably environmental and operational variability (EoVs) — can mask the subtle changes in structural response that indicate damage. If left unaddressed, these effects risk both obscuring true damage signatures and producing false indications, and so methods that normalise or remove their influence are a necessary component of any robust SHM pipeline. Third, and most relevant to the work presented here, SHM models often rely on large, varied datasets to learn reliable representations of structural behaviour. This requirement is particularly prohibitive for models that aim to provide detailed diagnostics — such as damage localisation, classification, or severity estimation — since these tasks demand supervised learning approaches with labelled observations spanning all damage states of interest.

1.1.0.2 Data Scarcity in SHM

Acquiring extensive, labelled datasets that capture a structure’s behaviour across the full spectrum of health conditions and degradation stages is prohibitively costly and often practically unattainable. Because critical engineering infrastructure is designed to be highly reliable, assets rarely experience severe damage under normal operation. Consequently, the data collected by operational SHM systems are overwhelmingly skewed toward the healthy state. This inherent data imbalance means models lack the necessary examples of deterioration, leading to unreliable predictions when faced with novel or progressing damage scenarios. Since acquiring damage labels from operational structures is often unfeasible, engineers take any opportunity to test decommissioned structures, such as the Z24 bridge [19], or laboratory-scale models [20]. Due to physical or experimental limitations, many SHM datasets only showcase the behaviours of physical systems through discrete realisations of their health states — outputs are measured, some form of damage is applied (typically adding masses or removing physical sections of the structure), and outputs are measured again. The practical difficulty of capturing continuously progressing damage, combined with the operational priority of distinguishing safe from unsafe states, has historically directed research attention toward classification technologies in SHM. While this is a natural starting point — and remains an important capability — the requirements of industry are evolving as the field matures. A

categorical safe/unsafe determination is increasingly insufficient for asset managers who must plan maintenance schedules, allocate inspection resources, and make economically informed decisions about the remaining useful life of a structure. These demands require not just the detection or localisation of damage, but a quantitative estimate of its severity and trajectory — tasks that classification alone cannot provide. While unsupervised machine learning algorithms, such as Autoencoders [21, 22] or One-Class Support Vector Machines [23, 24], can be used to detect the mere presence of damage (Level 1 in Rytter’s hierarchy), and multi-class classification algorithms [25] can be used to localise (Level 2) or identify (Level 3) the type of damage, regression is a natural approach to quantify (Level 4) the extent of damage and provide prognosis (Level 5) for the life of the structure.

The fundamental lack of adequate damage data on an individual asset level has given rise to a new approach: Population-based structural health monitoring (PBSHM) [26–29].

1.2 Population-based Structural Health Monitoring

The data scarcity challenges outlined above motivate a paradigm shift in how SHM systems are conceived: rather than treating each structure as an isolated asset, PBSHM considers a collection of structures jointly [26–29]. This approach is predicated on the observation that structures sharing common design principles, environmental conditions, loading regimes, or ageing mechanisms will exhibit statistical commonalities in their response data — commonalities that can be exploited to transfer information between population members. Where a single structure may never accumulate sufficient examples of damaged behaviour to train a reliable model, a population of nominally similar structures collectively represents a far richer source of information. PBSHM therefore reframes data scarcity as a problem of information routing: the data exists within the population, and the challenge lies in transferring it appropriately between assets. The primary methodological tool for achieving this is transfer learning. Transfer learning seeks to leverage knowledge acquired from one structure (the *source*) to improve inference on a related but distinct structure (the *target*). Domain adaptation is one of the most common branches of transfer learning and specifically addresses the distributional mismatch that arises between source and target — for instance, due to differences in geometry, material properties, or sensor configuration. Transfer learning methods form the technical backbone of PBSHM and their development is the central concern of this thesis. For a thorough introduction to the theoretical foundations of PBSHM, the reader is directed to [26–29].

1.2.1 Transfer Learning

Transfer learning is a machine learning paradigm in which knowledge acquired in one context is reused to improve inference in a related but distinct context [30]. Rather than training a model from scratch on limited local data, transfer learning allows a model to draw on richer information from elsewhere — either from a different dataset, a different domain, or a different but related task. In the context of PBSHM, this “elsewhere” is the population itself: other structures whose data can inform predictions about a target asset. One of the most widely-explored forms of transfer learning in PBSHM is *domain adaptation*, in which feature data are mapped from a label-rich source domain to a label-scarce target domain. The objective is to find a shared latent representation in which the source and target distributions are sufficiently aligned that label information — such as damage state or severity — can be transferred meaningfully from one to the other. The source and target domains may differ due to variations in geometry, material properties, sensor placement, or operational conditions, and so this alignment is rarely trivial. It is worth distinguishing between two broad regimes of transfer. In *homogeneous* transfer, the source and target share the same feature space and the challenge is purely distributional — the data look similar but are not identically distributed. In *heterogeneous* transfer, the source and target occupy different feature spaces entirely, requiring a more fundamental mapping between representations before any knowledge can be shared. Both regimes are relevant to structural populations, where assets may range from near-identical units on a wind farm to structurally distinct bridges sharing only a broad family resemblance.

1.2.2 Active Learning

Active learning (AL) presents another avenue to address the limitations imposed by data scarcity in SHM [31]. Conventional supervised machine learning requires labelled examples spanning the full range of conditions a model may encounter in deployment — a requirement that is rarely satisfiable in structural monitoring, where labels correspond to damage states that are costly or dangerous to induce deliberately. Active learning reframes this problem: rather than passively accepting whatever labelled data happens to be available, an active learning system selectively identifies which unlabelled observations would be most informative to label next, given what the model currently knows. By concentrating labelling effort on the most valuable examples, the goal is to achieve strong model performance with far fewer labels than conventional supervised approaches would require. The selection of which observations to query is governed by an *acquisition function*, which encodes some notion of informativeness. Common strategies

include querying points of maximum predictive uncertainty — where the model is least confident — or points expected to most reduce the overall uncertainty across the input space. More sophisticated criteria account for the cost of acquiring a label, balancing informativeness against practical constraints. Active learning can be applied in two broad settings. In the *offline* setting, a fixed pool of unlabelled data has already been collected, and the active learner selects a subset to label from this pool. In the *online* setting, data arrive sequentially and the system must decide in real time whether each new observation warrants the acquisition of its label [32]. The online setting is particularly relevant to SHM, where monitoring systems continuously stream data throughout the operational life of a structure and inspections are expensive, disruptive, or logistically constrained. Additionally, damage labels cannot be acquired retrospectively. An asset manager can only acquire these labels as and when they occur. In this context, an active learning system can be understood as a principled decision-making tool for inspection planning — one that directs human attention toward the moments most likely to yield meaningful diagnostic information.

1.3 Uncertainty in SHM

Uncertainty is a fundamental feature of almost every quantitative field — physics, economics, medicine, and engineering among them. It arises whenever a conclusion must be drawn from incomplete or imperfect information, and managing it well is a prerequisite for reliable decision-making. In the context of SHM, uncertainty is not merely an inconvenience but an inherent property of the problem: structural damage must typically be inferred from indirect measurements, sensor data is subject to finite precision and systematic error, and predictions must be made about future states of a system that is continuously evolving under poorly characterised loading and environmental conditions.

It is useful to distinguish between two fundamental sources of uncertainty. *Epistemic* uncertainty arises from a lack of knowledge — in principle, it could be reduced given more data, better models, or more thorough inspection. *Aleatoric* uncertainty, by contrast, reflects genuine stochasticity in the system and cannot be eliminated regardless of how much information is gathered. In practice, both are present simultaneously in SHM problems: a model may be uncertain about the current damage state because it has seen little data from the target structure (epistemic), while the structural response itself may vary unpredictably due to turbulent loading or environmental fluctuations (aleatoric). Distinguishing between these sources matters for decision-making — epistemic uncertainty motivates further data collection or model refinement, whereas aleatoric uncertainty must simply be accommodated.

Representing uncertainty formally requires a mathematical language. Several frameworks exist, including interval analysis [33, 34] and fuzzy logic [35, 36], but the current thesis adopts probability theory [37] as its foundation. Probability offers a coherent and well-developed calculus for reasoning about uncertain quantities, and crucially, it connects naturally to the broader framework of statistical decision theory — a principled approach to making optimal choices under uncertainty. This connection is particularly valuable in SHM, where the ultimate purpose of a model is not prediction in isolation but to inform consequential decisions about inspection, maintenance, and intervention. A model that quantifies its own uncertainty enables a decision-maker to weigh the costs of action against the risks of inaction in a formally grounded way, rather than relying on deterministic outputs that offer no indication of their own reliability.

1.4 Thesis Contributions

While classification-based approaches can determine the presence or general category of damage, regression enables a finer resolution of damage conditions. As the SHM field matures and industry demands move beyond binary safe/unsafe assessments, regression technologies become increasingly central to the discipline. Within the PBSHM paradigm, considerable research effort has been directed toward developing population-based technologies for classification tasks — transferring discrete damage labels across populations. By contrast, comparatively little attention has been paid to regression tasks, where the quantity of interest is a continuous health variable rather than a discrete category. This gap is consequential: the same data scarcity challenges that motivate PBSHM in the classification setting apply equally to regression, and the population-based tools needed to address them must be developed and validated accordingly. This thesis aims to address this gap by developing and demonstrating population-based SHM technologies specifically for regression tasks. Three core contributions of this thesis are outlined:

- **Contribution 1: Extend transfer learning technologies for PBSHM to regression tasks.** Transfer learning is the primary mechanism by which PBSHM exploits population-level information. This thesis develops and demonstrates transfer learning methodologies capable of mapping continuous damage states across structurally distinct source and target domains, addressing the challenges that arise when the quantity to be transferred is a continuous variable rather than a discrete label.
- **Contribution 2: Expand active learning to regression problems in population monitoring** Active learning offers a principled approach to data-efficient

monitoring by directing inspections toward the most informative observations. This thesis extends active learning to the population setting, demonstrating how population-based regression models can be used to improve inspection resource allocation across a fleet of structures — with explicit consideration of the economic consequences of monitoring decisions.

- **Contribution 3: Develop a benchmark dataset capturing continuous degradation across a population of structures.** The development and validation of regression-based PBSHM technologies requires experimental data that reflects continuous, progressive damage rather than discrete, artificially induced states. This thesis presents a novel dataset collected from a population of beams of varying material and geometry, in which degradation is introduced continuously — providing a benchmark for the community to further develop and evaluate population-based regression methods.

Together, these contributions represent a coherent programme of research aimed at advancing the capability of PBSHM beyond classification and toward the quantitative, prognostic technologies that the field — and industry — increasingly requires.

1.5 Thesis Outline

The remainder of this thesis is structured as follows:

Chapter 2 provides the theoretical background underpinning the machine learning methodologies developed throughout the thesis. This includes the probabilistic foundations of Bayesian inference, transfer learning approaches that enable information sharing across a population of structures, and active learning approaches that support budget-efficient inspection planning.

Chapter 3 presents a review of the relevant literature, surveying the current state of the art in transfer learning and active learning for structural health monitoring. The limitations and open challenges identified in this review directly motivate the research objectives of the thesis.

Chapter 4 fulfils *Contribution 2* by developing a risk-based active learning framework for populations of structures. A Bayesian hierarchical model is used to model the degradation of a population of machining tools, enabling data-rich members to inform predictions for data-scarce members. A decision-theoretic inspection planning process

is then defined, in which inspections are allocated according to expected utility — providing a principled and natural solution to the problem of budget allocation across a structural population.

Chapter 5 fulfils *Contribution 3* by introducing a novel experiment designed to capture continuous structural degradation in a population-based SHM setting. Progressive damage is applied to a population of beams of varying material and geometry, and the evolution of their dynamic response is monitored throughout. The experimental design is described in detail, and damage-sensitive features extracted from the resulting data are characterised and analysed.

Chapter 6 fulfils *Contribution 1* by developing a transfer learning framework for continuous damage states, applied to the dataset introduced in Chapter 5. Principal component analysis is used to reduce the dimensionality of the frequency response functions, and a novel domain alignment technique is proposed to map between structurally distinct source and target beams. The particular challenges that arise when transferring continuous, rather than discrete, damage information are discussed throughout.

Chapter 7 concludes the thesis by highlighting the key findings and contributions of each chapter. The advantages and limitations of the proposed methodologies are critically assessed, and directions for future research are identified.

Chapter 2

Background Theory

Structural Health Monitoring seeks to assess the condition of engineering structures through the analysis of measured data, with the aim of supporting reliable inspection and decision-making. As monitoring campaigns increasingly involve multiple structures operating under varying environmental and operational conditions, the focus of SHM has expanded from single-structure analysis to population-based approaches. Population-Based Structural Health Monitoring leverages information shared across a population of similar structures to improve diagnostic performance, particularly in scenarios where data from an individual asset are limited. The development of PBSHM has been closely linked with advances in machine learning, which provide flexible tools for modelling complex, high-dimensional data. Many of these methods rely on probability theory to quantify uncertainty and variability across structures within a population. However, differences between source and target structures, as well as the high cost of labelled data, present significant challenges, motivating the use of transfer learning and active learning to enable robust knowledge transfer and data-efficient learning. This chapter reviews the background theory to establish the foundation of the methods proposed in this thesis.

2.1 Structural Health Monitoring

Structural Health Monitoring aims to apply damage identification strategies to assess the condition of a structure. SHM systems reduce operation and maintenance (O&M) costs over the lifetime of an asset by alerting operators to emerging damage and indicating when maintenance might be necessary. These systems enable operators to move away from cost-inefficient O&M practices such as periodic inspections, towards more optimal

data-informed condition-based strategies. The level of information provided by an SHM system depends on the specific system and structure being monitored.

2.1.1 SHM approaches

Generally SHM models can be classified into physics-based models and data-driven models. Physics-based (or white-box) models rely on the governing physical laws of the system to interpret sensor measurements and infer health states. Data-driven (or black-box) models learn directly the relationship between measured data and system health states. Physics-informed (or grey-box) models are somewhere in between, where physics is encoded directly into the machine learning model such that the model follows the physical laws that govern the system it is modelling.

Physics-based models [6–8] are built on fundamental physical laws. Generally, these models offer some interpretability, where the parameters of the model have some physical meaning. This is helpful because learning the parameters of the system can give insight into underlying physical mechanisms. Physics-based modelling requires detailed understanding of the system (boundary conditions, material properties, geometry), including physics in the model that isn't representative of the underlying physical system will reduce the predictive capability of the model. Equally, physics-based models can predict behaviours under new conditions even if the behaviours are not captured in the data, as long as the behaviour is captured in the physics.

Data-driven models [9–12] learn patterns directly from measured data. These models don't, necessarily, require explicit knowledge of governing physics. However, data-driven models will only ever be able to make predictions about behaviours represented in the dataset they learn from. This means that data-driven models may require larger volumes of higher-quality data to generalise well. These models are capable of modelling complex system behaviour that can be difficult to capture explicitly in equations, but can lack the interpretability provided by physics-based models.

Physics-informed models [13–18] reduce the data requirements of black-box models by including physics in the model. For example, to model a fluid flowing through a pipe, a black-box model might need millions of datapoints to learn the behaviour and the model may not generalise well to different fluids or different pipes. A physics-informed model might be programmed to contain the equations of fluid dynamics directly in the model, allowing the model to predict better with less data and generalise to new domains.

2.1.2 Unsupervised learning

In unsupervised machine learning it is assumed that only input data are available, giving a training dataset $\{\mathbf{x}_i\}_{i=1}^{n_u}$ with n_u unlabelled samples. The goal is to learn “interesting patterns” in the data. Since there are no labels, y , there is no obvious error-metric to use. Typical uses for unsupervised machine learning include dimensionality reduction, density estimation and finding groups in data (clustering) [38]. In SHM, unsupervised learning can be used for *novelty detection*, where measurements of a damage sensitive feature can be recorded while the structure is ‘healthy’ to estimate a baseline distribution of the measurements under ‘normal’ conditions. Then during monitoring, measurements that deviate from the baseline will indicate damage. These methods are useful because one of the biggest problems in SHM is the scarcity of damage labels, the nature of novelty detection means that one only ever needs to take measurements from the undamaged structure making novelty detection less resource intensive than many supervised machine learning methods. Despite this the diagnostic capabilities of unsupervised machine learning are limited to simple detection [5].

2.1.3 Supervised learning

Supervised machine learning aims to learn the relationship between a set of input features x and a target variable y such that a model is trained on a dataset of input features and corresponding target variables $\{x_i, y_i\}_{i=1}^{n_l}$, with n_l labelled data. When the target variable is continuous the learning task is regression and when the target variable is discrete categories, the learning task is classification. In contrast to unsupervised learning, supervised approaches typically enable deeper diagnostics because the presence of labels provides more information to guide model training. Historically, many of the ML algorithms developed in SHM have been classification based [10, 39, 40]. A key reason is that SHM datasets obtained from laboratory experiments, field campaigns, or finite element models often contain discrete health states or a clear healthy–damaged distinction [19, 20, 41–43]. Implementing controlled levels of damage in experiments is challenging, making continuous damage labels rare and discrete labels far more common.

The availability of categorical datasets, coupled with the wider machine learning community’s influence, meant that many popular algorithms were directly adopted for SHM classification tasks. These include support vector machines [44–46] and convolutional neural networks [9–11], among others. Performance metrics such as the F1-score and ROC curves also provided straightforward ways of comparing model performance. This historical emphasis on classification provides useful diagnostics but also highlights the

need for methods capable of modelling continuous, evolving damage processes—an area where regression-based approaches are increasingly important.

2.1.4 Partially supervised learning

Obtaining fully labelled datasets is often infeasible in SHM; however, labels can be obtained for a subset of measurements through inspections during a monitoring campaign. As a result, SHM datasets often consist of a large quantity of unlabelled data and a small subset of labelled samples. This scenario is well-suited to partially-supervised learning methods [47], which aim to develop regression or classification models by integrating information from both labelled and unlabelled data. By incorporating unlabelled data into the learning process, partially-supervised methods can reduce the reliance on extensive labelled datasets.

Partially-supervised learning can generally be categorised into two approaches: *semi-supervised learning* and *active learning* [48]. Semi-supervised learning leverages both labelled data,

$$\mathcal{D}_l = \{(x_i, y_i)\}_{i=1}^{n_l},$$

and unlabelled data,

$$\mathcal{D}_u = \{x_i\}_{i=1}^{n_u},$$

by training models on the combined dataset

$$\mathcal{D} = \mathcal{D}_l \cup \mathcal{D}_u,$$

thus incorporating information from both types of samples. In contrast, active learning does not directly use unlabelled data during training. Instead, a predictive model is first trained on the available labelled data, then used to estimate the information content or uncertainty of unlabelled samples, guiding the labelling procedure to efficiently select the most informative unlabelled samples for annotation. The method a learner uses to choose which sample to label is called the query regime (or acquisition function), there are many different measures one can use such as risk [49], query-by-committee [50] and uncertainty sampling [51, 52]. Active learning will be covered more thoroughly in Section 2.6.

Whether to help guide labelling procedures or to quantify the uncertainty associated with measurements of structures, a probabilistic framework that can represent uncertainty explicitly is often preferred for ML in SHM [38].

2.2 Probability Theory

Measured data in SHM is inherently uncertain to some degree. This uncertainty can come from experimental sources, such as sensor accuracy, human error or model sources such as model discrepancy and parameter variability. This data is used to inform decision-making, without dealing with this uncertainty appropriately asset managers risk safety and economic consequences. As such, uncertainty should be mitigated during data acquisition and quantified within models [53]. Probabilistic frameworks offer methods to quantify this uncertainty. In SHM, signals measured from structures of interest can be treated as realisations of a random variable X , that can take values in a domain $\mathcal{X}$, i.e. $x \in \mathcal{X}$. For continuous variables, $\mathcal{X}$ is typically the set, or a subset, of real numbers, $X \subseteq \mathbb{R}$; for discrete variables, X is a finite set of discrete outcomes. The probability of a specific outcome describes the likelihood it will occur relative to all other possible outcomes, where $P(X) = 0$ indicates the outcome will definitely not occur and $P(X) = 1$ that it is certain to occur. Therefore, it is bounded within the interval $[0, 1]$, and the sum over all possible outcomes must equal one [54].

A continuous random vector is associated with a probability density function (p.d.f) $p(x)$, which describes the probability that X will take a value x in $a < X < b$ as follows,

$$P(a < X < b) = \int_a^b p(x) dx, \quad \text{s.t.} \quad \int_{\mathcal{X}} p(x) dx = 1, \quad 0 \leq p(x) \leq 1. \quad (2.1)$$

Often, the objective of SHM is to infer the health state y of the structure from a set of observations x . In much of the SHM literature y is a discrete category that describes the state of the structure. y can take one of K classes in a finite set $y \in \mathcal{Y} = \{1, 2, \dots, K\}$. This is called a *label space* in machine learning. A more realistic formulation considers y as a *continuous* health variable; Y taking values in a domain $\mathcal{Y} \subseteq \mathbb{R}$ (or $\mathbb{R}^d$ for multivariate states). This continuous description of health enables modelling of gradual degradation and damage progression. The random variable Y and its corresponding probability density function must satisfy the conditions stated in Equation 2.1. $p(y | x)$ denotes the *conditional* probability distribution, the probability of the state y given that x has taken a certain value. Given two random variables X and Y , the probability that they each take a certain value simultaneously, the *joint* probability, is given by $p(x, y)$. The joint probability can be decomposed as follows:

$$p(x, y) = p(y | x)p(x). \quad (2.2)$$

Independence between y and x would imply that knowledge of one value would not restrict the range of values the other would be expected to take. Meaning that Equation 2.2 could be further simplified as:

$$p(x, y) = p(y)p(x). \quad (2.3)$$

2.2.1 Bayes Theorem

Bayesian statistics is a probability framework that has had widespread adoption in data analysis because of its rigorous and coherent treatment of uncertainty. The typical Bayesian workflow consists of three main steps: Determine the prior distribution by capturing information about the parameters of the model before any data is collected; determine the likelihood function, using information about the parameters available in the observed data; finally, combining the prior distribution and the likelihood function to calculate the posterior distribution. This work flow is neatly contained in Bayes' theorem, shown in Equation 2.4.

$$\begin{aligned} P(\theta | X) &= \frac{P(X | \theta) \cdot P(\theta)}{P(X)} \\ \text{Posterior} &= \frac{\text{likelihood} \times \text{prior}}{\text{average likelihood}} \end{aligned} \quad (2.4)$$

Where θ is the parameters of the model and X is the available data. This equation expresses how a belief, expressed as a probability, should rationally change to account for the availability of evidence. That is, the prior distribution is updated to the posterior distribution after being exposed to data. This can be seen in Figure 2.1, where the posterior is somewhere in between the prior and the likelihood. The evidence has updated the prior and effectively moved it closer to the evidence in the form of the posterior.

Often the denominator on the right-hand side of Equation 2.4 is overlooked, it is a normalisation factor that ensures the posterior probabilities add up to one. The denominator is an integral over all possible parameter values and if θ is high-dimensional, the denominator becomes very difficult to calculate.

$$P(X) = \int P(X | \theta)P(\theta)d\theta \quad (2.5)$$

Several ways of addressing this problem have been developed. Historically, researchers were limited to conjugate likelihood-prior pairs [56], where for certain prior-likelihood




FIGURE 2.1: A graphical representation of a prior distribution being updated to a posterior distribution after exposure to evidence [55].

pairs, the posterior distribution belongs to the same family as the prior, forming a conjugate prior. More recent developments allow us to avoid computing the denominator entirely using Markov-chain Monte Carlo (MCMC) [57–59], approximating the posterior using Variational Inference [60, 61] or Laplace approximation [62] as well as Approximate Bayesian Computation (ABC) [63] and nested sampling [64] methods. A thorough explanation of Bayesian data analysis can be found in [65].

2.2.2 Prior distributions

A prior probability distribution, often called simply a prior, is a key part of Bayesian inference. A prior is the assumed distribution of an uncertain parameter before some evidence is taken into account. The prior should be considered as any another part of the model and so should be chosen, evaluated and revised [66]. The key issues in setting up a prior distribution are [67]:

1. What information is going into the prior distribution?
2. What are the properties of the resulting posterior distribution?

Prior information often comes from previous experiences, it is common to use the posterior distribution from one model as the prior distribution for the next model. The prior can come from many sources and is used to encode the prior belief about the parameter. In the past priors were contained mostly to conjugate pairs, specific sets of prior and likelihood combinations that result in analytical solutions for the posterior distribution. Now, with the development of Markov-Chain Monte Carlo sampling, the posterior distribution can be approximated for a wider range of priors. When there is limited understanding of the true nature of the parameters of the system, it is suggested to use

‘weakly informative priors’, these are priors that express partial or vague information about a variable. The purpose of weakly informative priors is for regularization and to very loosely constrain variables to sensible values [66]. An important consideration for priors are which values are included in the support of the distribution. If certain values are excluded in the prior with 0 probability, then no matter how strong the evidence is for the values, they will retain 0 probability in the posterior.

2.2.3 The likelihood

The likelihood, $P(X | \theta)$, is the probability of observing the set of data given a specific set of parameters. Likelihood functions can be chosen by the analyst, and the function the analyst chooses can have a large effect on the posterior distribution and parameter estimation of the analyst’s model. For example, choosing a student-t distribution for the likelihood can imbue robustness to a model [66].

2.2.4 Posterior distribution

The posterior distribution, $P(\theta | X)$, reflects one’s updated knowledge, balancing prior knowledge with observed data [65]. It describes the uncertainty about statistical parameters conditional on a collection of observed data. Posterior distributions should be a better reflection of the underlying generative process of the data than a prior since the posterior includes more information. From the posterior distribution estimates, such as maximum a posteriori (MAP), can be derived. As previously mentioned, if new data becomes available, the posterior can become the new prior to incorporate new information into the analysis.

2.3 Problems in SHM

Perhaps the largest problem in SHM is data-scarcity. Creating damage labels requires damaging the structure of interest. This can be expensive, high-risk and is often not possible for real structures. This limits the usefulness of supervised machining learning methods particularly for regression and deep learning approaches. Additionally, even when labels are available the datasets are imbalanced due to healthy states being much more common than damaged ones. This leads to implicit biases for models trained on this data. A similar problem occurs due to environmental and operational variability. Temperature, humidity and loading (e.g. traffic, wind) can change measured signals, masking or mimicking damage, making robust detection difficult without comprehensive

datasets. The difficulties imposed by inadequate data motivates sharing information between similar assets [68].

2.4 Population-based SHM

PBSHM aims to address the problem of data-scarcity in SHM. By adopting a population-based approach to SHM, such that knowledge and information can be transferred between similar structures, there is potential for improved diagnostic and prognostic capabilities [69]. The foundations of PBSHM have been presented in a series of journal papers by Worden *et al.*, each detailing the fundamental concepts of the approach; homogeneous populations [26], heterogeneous populations [27], mapping and transfer [28], the geometric spaces in which structures exist [29], and embedding these technologies in software [70].

2.4.1 Homogeneous and heterogeneous populations

In a population of structures — such as the wind turbines comprising an offshore wind farm — members may share common geometries, topologies, materials, and boundary conditions. Where structures are built to the same design specification, using the same components and manufacturing processes, they can be considered *nominally identical*, and a population consisting exclusively of such structures is referred to as a *homogeneous* population. In principle, homogeneous populations offer the most straightforward setting for information sharing: the statistical commonalities between members are strong, and knowledge acquired from one asset can be transferred to another with minimal correction. In practice, however, even nominally identical structures will exhibit variation. Manufacturing tolerances, differences in installation conditions, and localised environmental effects all introduce heterogeneity that is absent from the design specification. Offshore wind turbines, for instance, may be structurally identical yet experience markedly different geotechnical conditions due to variation in the geological composition of the seabed beneath each foundation. These differences modify the effective boundary conditions of the turbine towers and produce measurable changes in their dynamic behaviour — meaning that even a nominally homogeneous population will exhibit some degree of structural variability in practice [71]. Accounting for this variability is therefore important even in the homogeneous case, and becomes increasingly critical as the diversity of the population grows. An example of a homogeneous population can be seen in Figure 2.2.




FIGURE 2.2: A representation of a homogeneous population of wind turbines [71].

In contrast, a *heterogeneous* population encompasses structures that are not exclusively nominally identical — admitting greater diversity in design, geometry, material properties, and boundary conditions, and potentially including multiple distinct structural types. Despite the more significant differences between members, heterogeneous populations can still exhibit commonalities that are valuable for knowledge transfer. The challenge lies in identifying and exploiting these shared characteristics while accounting for the structural dissimilarity that distinguishes the population from the homogeneous case.

Returning to the offshore wind farm example, consider an operating company that manages two wind farms at different locations — one equipped with three-bladed turbines and one with four-bladed turbines. Treating each farm as a separate homogeneous population is a natural starting point, but there are additional benefits to considering them jointly as a single heterogeneous population. The tower designs of the two turbine types may be closely related, for instance, and by pooling observations across both farms it becomes possible to develop improved predictive models for this shared substructure — models that would be harder to learn from either population in isolation. The same principle applies across a wide range of engineering contexts: an aircraft fleet comprising multiple models, a portfolio of suspension bridges with varying span configurations, or a collection of offshore platforms of different designs may all contain transferable structure that a population-based approach can exploit, provided that the mapping between domains is handled appropriately.

This distinction between homogeneous and heterogeneous populations is important for the transfer learning methodology developed in this thesis. The beam population introduced in Chapter 5 is deliberately heterogeneous — spanning multiple materials and

geometries — and the transfer learning framework of Chapter 6 is designed specifically to accommodate this variability, addressing the more challenging and practically relevant case in which source and target structures are related but not identical. An example of a heterogeneous population can be seen in Figure 2.3.




FIGURE 2.3: A representation of a heterogeneous population of wind turbines [71].

For further discussion on homogeneous and heterogeneous populations, the reader is directed to [26] and [27]. Although the similarities between structures have been described qualitatively so far, it is helpful to quantify this similarity to better understand where information transfer can be effective.

2.4.2 Similarity quantification in heterogeneous populations

Knowing when to do transfer is just as important as knowing how to do transfer. If an engineer does no assessment of similarity across a population, there is a chance of *negative transfer* where the use of transfer learning decreases the predictive capability of the model. Quantifying similarity between structures is useful because it can be assumed that, typically, the more similar structures are the more likely it is to successfully transfer information between them. Negative transfer has particularly critical consequences in SHM where decisions made on predictions by SHM models can have large economic and safety implications. Graph theory provides a powerful and flexible framework for representing and comparing discrete structured objects. Within the PBSHM literature, Gosliga *et al.* [27] introduced the *irreducible element* (IE) model as a structural representation, offering considerably lower dimensionality than alternatives such as finite element or CAD models while retaining the information most relevant to health monitoring applications. The IE representation abstracts a structure into a set of discrete elements — such as beams, plates, and shells — each carrying geometric and material

attributes, connected by joints that encode the topological relationships between them. By capturing the essential character of a structure in this reduced form, the IE model provides a common language for comparing members of a population, quantifying their similarity, and determining the degree to which knowledge acquired from one structure can be meaningfully transferred to another. An example of an IE model can be seen in Figure 2.4




FIGURE 2.4: An expanded IE representation of a wind turbine with the elements labelled A to I and the connections labelled 1 to 9. The shaded node is a special element representing the ground, where the boundary condition is defined in the attributes for Joint 9 [27].

Once an IE representation of a structure is obtained, the information can be encoded into an attributed graph (AG). While the IE model presents essential characteristics of a structure in a readable format for humans, the purpose of the AG is to embed structures into a space that enables efficient pairwise comparisons. By embedding structures into a metric space via AGs, graph-matching algorithms can be utilized to identify common subgraphs among sets of structures. These common subgraphs represent shared substructures within the sets of structures and provide insights into potential areas for knowledge transfer. Measures of proximity within the AG space, such as the Jaccard index as was used in [27], can be used as measures of similarity. An example of an AG can be seen in Figure 2.5.

In summary, graphical representations of structures allow comparison, identification of common substructures, and quantification of similarity.




FIGURE 2.5: The resulting Attributed Graph from the IE representation shown in Fig. 2.4. It is possible to see how the topology is defined by the element set in the joints [13].

Another measure of similarity is the modal assurance criteria (MAC) [14]. The MAC is a physics-based tool that assesses the similarity of mode shapes between structures, it is a normalised scalar product between each pair of modal vectors $\psi_s^{(i)}$ and $\psi_t^{(i)}$, where the $_s$ denotes the source domain and $_t$ denotes the target domain (subjects that will be covered in the following sections). The elements of the MAC matrix are given By

$$\text{MAC}(i, j) = \frac{\left| \psi_s^{(i)T} \psi_t^{(j)} \right|^2}{\psi_s^{(i)T} \psi_s^{(i)} \psi_t^{(j)T} \psi_t^{(j)}}. \quad (2.6)$$

where $\text{MAC}(i, j) \in [0, 1]$, with 0 indicating no correspondence and 1 showing complete correspondence [72]. It has been shown that the structural similarity measured by the MAC is an indicator for successful transfer.

2.5 Transfer Learning and Multi-task Learning

Standard machine learning (ML) models are only able to make suitable predictions about scenarios they have seen before (e.g. predictions on test data will be poor if that data is not represented in the training data), this is a major limitation of standard ML. Human learning rarely occurs like this; rather, it relies on previous experiences to improve performance for new skills. For instance, an athlete proficient in basketball possesses transferable skills that allow them to learn netball more efficiently. Despite lacking direct experience in the new sport, the similarity between tasks enables the athlete to leverage previous training to enhance performance. This principle is analogous to Transfer Learning (TL), where knowledge derived from a data-rich source task is utilized to improve predictive accuracy on a related target task with limited data. ML algorithms can also benefit from mechanisms that leverage relationships between similar domains. This logic drives the field of Transfer Learning. This section outlines the principles of Transfer Learning and motivates its application within PBSHM.

Following the foundational survey [30], two core objects must be defined: a *domain* and a *task*.

Definition 2.1 (Domain). A domain, $\Omega = \{\mathcal{X}, P(X)\}$, is defined by a feature space $\mathcal{X}$ and a marginal probability distribution $P(X)$ on that space, where $X = \{x_1, \dots, x_n\} \in \mathcal{X}$.

Definition 2.2 (Task). A task for a given domain is defined by $\mathcal{T} = \{\mathcal{Y}, f(\cdot)\}$, where $\mathcal{Y}$ is the label space and $f(\cdot)$ is a predictive function learnt from a finite sample $\{x_i, y_i\}_{i=1}^n$, where $x_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$. In a probabilistic setting, the predictive function can also be viewed as modelling the conditional distribution $P(Y|X)$.

The most common paradigm in Transfer Learning considers a single source and a single target domain. While scenarios exist involving multiple source domains (Multi-Source Transfer Learning), for simplicity, this section assumes a pairwise transfer.

Definition 2.3 (Transfer Learning). Given a source domain Ω_s and corresponding task $\mathcal{T}_s$, and a target domain Ω_t and corresponding task $\mathcal{T}_t$, the objective of Transfer Learning is to use information contained within Ω_s and $\mathcal{T}_s$ to improve the target predictive function $f_t(\cdot)$ in $\mathcal{T}_t$, where $\Omega_s \neq \Omega_t$ and/or $\mathcal{T}_s \neq \mathcal{T}_t$.

Typically, it is assumed that the source domain contains sufficient labelled data to learn a predictive function, while the target domain has significantly fewer data. The specific approach to transfer learning varies depending on the availability of labels in the source and target domains. Pan et al. [30] classify the main forms as follows:

- **Unsupervised Transfer Learning (Inductive/Transductive):**

The source domain consists of labelled data, $D_s = \{x_{s,i}, y_{s,i}\}_{i=1}^{n_s}$, which are assumed to be generated following the underlying source joint distribution $P_s(Y, X)$. The target domain consists only of unlabelled data, $D_{t,u} = \{x_{t,j}\}_{j=1}^{n_{t,u}}$, generated from the underlying target joint distribution $P_t(Y, X)$ [73]. Unsupervised transfer learning focuses on solving unsupervised learning tasks in the target domain, such as clustering, dimensionality reduction, and density estimation [74, 75]. In this case, there are no labeled data available in both source and target domains in training. This topic is typically only used in the context of *feature-representation-transfer* and does not form a part of this thesis.

- **Supervised Transfer Learning:**

Both the source dataset, D_s , and the target dataset, $D_t = \{x_{t,j}, y_{t,j}\}_{j=1}^{n_t}$, contain labelled data, though the number of target samples n_t is typically small. Similar to conventional machine learning, this setting can be extended to semi-supervised transfer learning by incorporating both labelled and unlabelled data into the learning process.

2.5.1 Multi-task Learning

Closely related to transfer learning is the paradigm of *multi-task learning* (MTL), in which multiple related tasks are learned jointly rather than sequentially. In contrast to transfer learning, where a distinction is made between a data-rich source task and a data-scarce target task, multi-task learning treats all tasks symmetrically, with the objective of improving generalisation performance across the entire set of tasks through shared representations or parameters.

Formally, given a collection of tasks $\{\mathcal{T}_k\}_{k=1}^K$ defined over corresponding domains $\{\Omega_k\}_{k=1}^K$, multi-task learning seeks to learn predictive functions $\{f_k(\cdot)\}_{k=1}^K$ simultaneously by exploiting commonalities between tasks, typically through shared latent structures, regularisation, or hierarchical priors [76]. By coupling the learning processes, information from data-rich tasks can implicitly regularise and inform data-poor tasks, reducing overfitting and improving predictive performance.

From a probabilistic perspective, multi-task learning is naturally expressed within a Bayesian framework, where task-specific models are linked through shared hyperparameters or hierarchical priors. This formulation enables principled information sharing across tasks while preserving task-specific uncertainty, and closely aligns with population-based modelling approaches in structural health monitoring [77]. In this sense, multi-task learning may be viewed as a special case of transfer learning in which no single task is

designated as the target, and knowledge transfer occurs bidirectionally across the task set.

The distinction between transfer learning and multi-task learning is therefore largely one of emphasis and problem formulation. While transfer learning is often adopted when interest is focused on improving performance for a specific target task, multi-task learning is particularly well suited to settings involving collections of related systems or experiments, where learning across a population is desired. In this work, multi-task learning is treated as a closely related special case of transfer learning, in which multiple related tasks are learned jointly through shared representations.

2.5.2 Applications of Transfer Learning and Multi-task Learning

Transfer learning and multi-task learning have enabled the deployment of high-performance models in data-constrained environments. While its foundations lie in general computer science applications, TL has increasingly been adopted in engineering, including structural health monitoring, where labelled data are often scarce.

2.5.2.1 General Machine Learning Applications

In the broader AI community, TL is the standard solution for initializing models in domains where training from scratch is computationally prohibitive or data is scarce.

- **Computer Vision:** The most prominent application of TL involves models pre-trained on massive datasets such as ImageNet [12]. Architectures like ResNet [78] are trained to recognize generic features (e.g., edges, textures) on millions of images. These pre-trained weights are then “frozen” or fine-tuned to detect specific objects in a target domain with limited data, such as identifying defects in solar panels or classifying medical imaging. MTL is often applied to jointly predict object categories, bounding boxes, and segmentation masks, improving generalisation across tasks [76].
- **Natural Language Processing (NLP):** Large Language Models (LLMs) like BERT [79] and GPT [80] are pre-trained on vast text corpora and then fine-tuned for domain-specific tasks, such as sentiment analysis or legal document classification. MTL can simultaneously learn related NLP tasks, e.g., part-of-speech tagging, named entity recognition, and syntactic parsing, improving overall model performance [81].

- **Healthcare / Medical Applications:** TL allows models trained on large public datasets to be adapted to smaller, specialised datasets, e.g., predicting disease outcomes from medical images. MTL can leverage shared structures in the data to predict multiple disease markers or anatomical measurements from the same patient scans, enhancing predictive accuracy across all tasks [81].

2.5.2.2 Engineering Applications

In engineering contexts, data acquisition is often constrained by the cost of destructive testing, rarity of failure modes, or limited operational data. TL bridges the gap between different operating conditions or physical assets, MTL enables simultaneous learning across multiple related tasks, improving generalisation and reducing a model's susceptibility to outliers.

- **Sim-to-Real Transfer in Robotics:** Training reinforcement learning agents directly on hardware is often unsafe or impractical. TL enables policies learned in physics simulations (source) to be transferred to physical robots (target) using techniques like Domain Randomization [82]. MTL can also be applied to jointly learn multiple control tasks—e.g., navigation, object manipulation, and obstacle avoidance—within the same model, improving overall adaptability [81].
- **Fault Diagnosis in Rotating Machinery:** TL compensates for varying operating conditions, e.g., adapting a bearing fault classifier from one speed or motor to another [83, 84]. MTL can extend this by jointly predicting multiple fault types or conditions from shared vibration signals, improving diagnostic accuracy across all tasks.
- **Biomedical Engineering: Brain-Computer Interfaces:** TL facilitates Subject-to-Subject Transfer, reducing calibration time by adapting models trained on previous users to a new individual [85].
- **Electrical Engineering: Battery Health Estimation:** TL enables Cross-Condition Transfer of degradation models from lab-aged batteries to operational EVs [86]. MTL can be applied to jointly predict multiple battery metrics—State of Health, Remaining Useful Life, and internal resistance—allowing the model to exploit correlations between degradation indicators and improve prediction accuracy [87–89].
- **Smart Grids and Energy Forecasting:** TL allows models trained on existing wind or solar farms to forecast output at new locations [90]. MTL can improve

performance by jointly forecasting multiple outputs (e.g., wind, solar, and load demand) across a network of sites, capturing interdependencies between energy sources. [91]

- **Materials Science and Discovery:** The discovery of new materials, such as high-entropy alloys, is often bottlenecked by the high cost of synthesis and characterization. Transfer Learning aids this process by enabling property prediction for novel chemical compounds where data is sparse. The source domain generally comprises large datasets of simple, well-characterised materials (e.g., standard crystals), while the target domain includes small datasets of novel, complex materials. By leveraging the fact that relationships between atomic descriptors and physical properties learned from simple structures often hold true for more complex systems, TL accelerates the screening process for new materials [92].

2.5.2.3 Transfer Learning in Structural Health Monitoring: PBSHM

In SHM, the “Data Problem” is the primary driver for Transfer Learning. It is rare to possess historical damage data for every specific structure in a fleet; thus, knowledge must be transferred from the few structures that have experienced failure.

- **Transfer in Engineering Populations:** This approach treats a fleet of structures (e.g., wind turbines) as a population. TL facilitates the transfer of damage detectors from a source structure with labelled failure data to a target structure that is data-poor [26]. This is critical for scaling SHM, as it accounts for the inevitable differences in resonance frequencies caused by manufacturing tolerances and foundation stiffness [27]. MTL can also be used to transfer information across a fleet of engineering systems, where data-scarce members can borrow statistical strength from those that are data-rich [68]. Recent developments show that even highly-disparate (heterogeneous) populations can be transferred between [93].
- **Environmental and Operational Compensation (EOC):** Variations in environmental conditions, such as temperature, can cause feature distributions to drift, potentially triggering false alarms. Domain Adaptation, a subfield of TL, is used to align the statistical distributions of features collected under different environmental states (e.g., Summer vs. Winter), ensuring the damage detector remains valid without requiring continuous re-labelling of data [94, 95].
- **Experimental-to-Operational Transfer:** A common strategy involves damaging a scaled laboratory model to learn failure patterns. TL is required to project

these learned patterns onto the full-scale operational structure, bridging the “domain gap” between the controlled laboratory environment and the noisy operational environment [96].

- **Sensor Signal Reconstruction:** MTL can be used to reconstruct missing sensor measurements by learning correlations between different sensor time series [97]. MTL captures correlations among different signal tasks (e.g., sensors or frequency components), improving reconstruction accuracy and computational efficiency over independent models [98].

2.6 Active Learning

Another method to address data scarcity is active learning (AL). AL is a field within machine learning that aims to improve machine learning capabilities in data-scarce scenarios by strategically selecting the most informative samples for annotation. This is especially valuable in domains where obtaining labels is costly or requires expert knowledge, such as engineering diagnostics and structural health monitoring. Rather than passively learning from a fixed labelled dataset, active learning algorithms iteratively query an ‘oracle’ to obtain labels expected to maximally improve model performance [52].

Active learning methods are commonly categorised based on their query strategies.

- **Uncertainty-Based Sampling:** Uncertainty sampling selects instances for which the model is least confident in its predictions. Early work in this area introduce methods such as least confidence, margin sampling, and entropy-based measures [99, 100]. These methods are simple and widely applicable but can focus excessively on outliers or noisy data.
- **Query-by-committee:(QBC)** maintains a committee of models representing different hypotheses consistent with labelled data. Samples that induce high disagreement among committee members are selected for labeling [50]. QBC is theoretically sound but can be computationally demanding for complex models.
- **Expected Model Change and Error Reduction:** These approaches select samples expected to induce the largest change in model parameters or yield maximal reduction in generalisation error [101, 102]; While appealing in theory, exact computation is often intractable, requiring approximations.

- **Information-Theoretic / Bayesian Methods:** Information-based strategies aim to maximize the mutual information between queried labels and model parameters or predictions [103, 104]. Probabilistic models such as Gaussian processes naturally lend themselves to these approaches.

2.6.1 Batch, Pool, and Online Active Learning

In some settings, labelling one sample at a time is often impractical. Batch active learning selects multiple samples per iteration while reducing redundancy [105, 106]. Pool-based AL [107] assumes a fixed unlabelled data set, while online (stream-based) AL [32, 48] considers sequentially arriving data. Online AL is most applicable for SHM as data collected through SHM systems will arrive sequentially throughout the life of the structure.

2.6.1.1 General Machine Learning Applications

Active Learning (AL) addresses the challenge of limited labelled data by intelligently selecting the most informative samples to label. This reduces annotation costs while maintaining model performance, and is particularly valuable in domains where labelling is expensive or time-consuming.

- **Text and Image Classification:** AL is widely used to train image classifiers and object detectors efficiently. Models query the most uncertain or representative images from unlabelled datasets to improve medical imaging (tumor detection) [108], defect detection [109], or wildlife monitoring [52]. Similar methods are used to reduce queries in sentiment analysis [110], spam detection [111], and named entity recognition [112].
- **Material Science:** AL is used to efficiently explore chemical or materials spaces [113]. By querying experiments that are expected to be most informative, researchers reduce the number of lab experiments needed to predict properties like solubility or conductivity [114].
- **Robotics and Control Systems:** AL reduces the number of trials a robot must perform to learn a task by querying the most informative states or actions [115]. This is especially important when real-world trials are costly or dangerous [116].

2.6.1.2 Engineering Applications

In engineering systems, data acquisition and labelling are often expensive, time-consuming, or destructive. Active learning is therefore used to optimise where, when, and how data are collected, ensuring that each new observation maximally improves model performance or decision quality.

- **Experimental Design:** In experimental mechanics and materials engineering, active learning guides the selection of experimental conditions (e.g., load levels, material compositions) to efficiently explore parameter spaces [114]. This reduces the number of costly experiments required to identify governing relationships [117, 118]
- **System Identification and Model Calibration:** Active learning is used to efficiently explore a feature space [119]. In many engineering systems, parameters are estimated sequentially as new data become available. Active learning guides when and where new measurements should be taken to refine the model efficiently. [116]

2.6.2 Active Learning in Engineering and Structural Health Monitoring

In SHM, active learning is employed to reduce the cost and intrusiveness of data acquisition by selectively guiding experiments, measurements, and inspections. AL has been applied to a range of tasks including system identification, sensor placement, damage detection, surrogate modelling, and inspection planning.

- **Optimal Sensor Placement and Measurement Selection:** Active learning is closely related to optimal experimental design in structural identification, where sensor locations or measurement configurations are selected to maximise information gain about uncertain model parameters [117, 120]. By prioritising informative sensor placements, AL reduces instrumentation costs while preserving identifiability of structural dynamics [118, 120].
- **Damage Detection and Condition Assessment:** Active learning has been applied to damage detection by prioritising data acquisition in regions of high predictive uncertainty or ambiguity between healthy and damaged states [95, 121]. This is particularly relevant for rare damage scenarios, where labelled examples are scarce and costly to obtain [5].

- **Inspection Planning and Decision-Theoretic SHM:** AL is increasingly used to support inspection planning by determining when and where inspections should be performed to maximise information about structural condition while minimising cost. Bull et al. [121] propose a probabilistic online AL framework for SHM, while Hughes et al. [49] formulate inspection planning as a decision-theoretic active learning problem incorporating inspection costs and risk.
- **Surrogate Modelling and Computational Efficiency:** In computationally expensive simulations, AL is used to construct surrogate models by adaptively selecting simulation points that improve predictive accuracy. This approach is common in uncertainty quantification and reliability analysis, where full model evaluations are prohibitive [114, 122].

2.6.3 Active Learning for Regression

While early research focused on classification, regression-based active learning has gained prominence, especially in physical systems. Predictive variance often guides sample selection, making Gaussian process regression a natural choice [123, 124]. Criteria such as variance maximization [125], integrated variance reduction [126], and Bayesian experimental design [127] have been widely studied, particularly in engineering applications where outputs are continuous and uncertainty quantification is essential. Recent directions include AL with deep learning [106], cost-aware querying [128], and meta-learning for query strategy optimization [107].

2.7 Conclusion

This chapter has introduced the theoretical and methodological foundations underpinning the contributions of this thesis. These include foundations of SHM and Probability Theory, PBSHM and transfer learning, and active learning in engineering applications. These foundations provide the theoretical language and tools required to understand and evaluate the contributions of this thesis. The following chapter surveys the existing literature on transfer learning and active learning for SHM in greater depth, identifying the specific gaps that motivate the research presented in Chapters 4, 5, and 6.

Chapter 3

Literature Survey: Transfer Learning and Active Learning in SHM

The background presented in Chapter 2 establishes the theoretical foundations that underpin the data-efficient methodologies explored in this thesis — from the probabilistic treatment of uncertainty to the principles of transfer learning and active learning. Historically, traditional SHM approaches have been unable to resolve the problem of label scarcity at the individual asset level; the shift towards Population-based SHM has reframed this challenge by enabling information to be shared across a population of structures, such that data-rich assets can support inference for data-scarce ones. This core tenet of PBSHM means the field is naturally oriented towards data-efficient technologies — methods that extract maximum value from limited observations and direct data collection intelligently. The current chapter surveys the literature on these methodologies, situating the contributions of this thesis within the broader research landscape. The review is organised into two principal categories: *transfer learning* — encompassing domain adaptation, hierarchical modelling, and neural network fine-tuning — and *active learning* — with a focus on query strategy design, budget-constrained inspection planning, and population-level resource allocation.

3.1 Transfer Learning

Multi-task learning and in particular hierarchical Bayesian modelling has been shown to be an effective tool for knowledge transfer in PBSHM [68, 91, 97, 129–133]. A recent example of probabilistic information sharing is the work of Dhada et al. [134],

who employed a Weibull statistical hierarchical model to facilitate collaborative prognosis. By treating asset-specific failure parameters as samples from a shared population distribution, they effectively mitigated the high variance associated with sparse failure data—a common challenge in SHM. However, the reliance on a fixed parametric form (Weibull) limits its applicability to degradation processes that strictly adhere to standard reliability laws. Additionally, often these types of models only work for homogeneous populations and can struggle when the population becomes disparate. Dhada et al. [135] addressed the ‘cold start’ problem in unsupervised anomaly detection by proposing a hierarchical statistical model for industrial fleets. Recognising that independent models suffer from high variance when data is sparse, they employed a Normal-Inverse-Wishart prior to facilitate collaborative learning within latent subpopulations. This partial pooling approach effectively stabilized the baseline estimation for new assets by borrowing covariance information from similar peers. However, the reliance on a Multivariate Gaussian assumption limits its efficacy to systems with unimodal, linear operating regimes. In the domain of dynamic characterisation, Dardeno et al. [136] applied hierarchical Bayesian modelling to Frequency Response Functions across a population of helicopter blades. By employing a partial pooling strategy, they demonstrated that fleet-wide information could constrain the uncertainty of modal parameters for data-sparse assets. The reliance on computationally intensive MCMC sampling currently limits its deployment in high-frequency, real-time monitoring scenarios. Brealy et al. [137] demonstrated the utility of hierarchical modelling for capturing spatial dependencies in wind farms. By framing the power curve estimation of each turbine as a related task within a Multi-Task Learning framework, they successfully modelled the complex ‘wake effects’ that cause performance variations across the fleet. This approach allows for the prediction of power generation behaviour for unobserved turbines by leveraging their spatial coordinates and the learned correlations from neighbouring assets. Tsialiamanis et al. [138] tackled the problem of damage prognosis in the context of PBSHM. Two methods were proposed, firstly a functional subspace was defined that includes the potential behaviour of members of the population subjected to the phenomenon of damage evolution. Secondly, a *meta-learning* method, where *deep kernel transfer* was used to enhance the predictive performance of a Gaussian process in a crack growth case study.

Sedehi et al. [139] propose a framework for integrating physics-based models within a Hierarchical Bayesian Gaussian Process. By embedding the governing equations directly into the GP’s mean function and deriving kernels from the physical gradients, they ensured that the learned discrepancy terms remained physically interpretable. Furthermore, their hierarchical treatment of time-series segments allowed for the accommodation of non-stationary errors, effectively bridging the gap between simplified physical

models and complex real-world structural behaviour. Di Francesco et al. [140] integrated decision-theoretic planning into hierarchical modelling. By coupling hierarchical Bayesian modelling with Value of Information (VoI) analysis, they demonstrated that the ‘value’ of a datapoint is not intrinsic, but dependent on its ability to reduce uncertainty across correlated assets. However, their work also highlights the computational bottleneck of full Bayesian decision analysis, motivating the need for more efficient, proxy-based active learning heuristics in large-scale populations. While most PBSHM literature focuses on the macro-scale variability between assets, Papadimas and Dodwell [141] applied hierarchical Bayesian modelling to the micro-scale challenge of stochastic material calibration. Dealing with additively manufactured steel, they utilised Bayesian Quadrature to efficiently marginalise over individual test samples, thereby isolating the inherent aleatory uncertainty of the material from epistemic measurement noise. This work highlights that the ‘population’ concept is scale-invariant; the mathematical machinery used to model a fleet of wind turbines is equally applicable to a batch of tensile test specimens. One of the first attempts to tackle heterogeneous populations in PBSHM was Gardner et al. [142] who introduced kernelised Bayesian transfer learning to the PBSHM framework. By projecting data from heterogeneous structures into a common latent subspace via task-specific kernels, they enabled knowledge transfer between assets with differing sensor counts and geometries. However, its reliance on kernel methods imposes an $O(N^3)$ computational cost, currently restricting its application to condensed feature sets.

While deep learning approaches dominate recent literature, typical DA methods rely on nonparametric distance metrics, which require sufficient data to perform density estimation. Poole et al. [143] demonstrated that simpler Statistic Alignment (SA) methods can often outperform complex kernel-based techniques in PBSHM. By explicitly aligning only the first and second-order statistics of the normal operating condition—via Normal Condition Alignment (NCA)—they mitigated the risk of negative transfer caused by class imbalance. This paper highlights that NCA is a key preprocessing tool for successful transfer. However, because SA relies on linear transformations of the feature space, its efficacy is likely limited to populations where the structural differences can be approximated by affine mappings, rather than complex non-linear distortions (for example, a bridge that behaves differently in winter and summer might be multimodal). Dardeno et al. [93] used NCA as a preprocessing tool for transfer. This work tackles the problem of heterogeneous populations in PBSHM through transfer via intermediary structures. By generating a continuous sequence of intermediary, synthetic structures that morph geometrically from a bridge to an aircraft, they established a ‘chain’ of high-fidelity transfer tasks. This work highlights that positive transfer between disparate systems is possible if one can construct a smooth manifold of physics-based precursors connecting

the source and target domains. As with most of the literature for DA in SHM, this work considers classification problems. Wickramarachchi et al. [144] extended the utility of statistic alignment to the temporal domain, specifically addressing the ‘Repair Problem’ in long-term monitoring. Structural repairs often induce a distributional shift that invalidates existing diagnostic models; however, by employing NCA to map post-repair features back to the pre-repair baseline, they demonstrated that legacy damage classifiers could be reused. This approach effectively salvages the ‘sunken cost’ of labelled pre-repair data, preventing the need to restart the learning process after every maintenance intervention. Poole et al. [14] showed that physics-based metrics can be used to inform feature selection and measure similarity, showing that the MAC can be used to predict positive (or negative) transfer.

Rezazadeh et al. [145] monitors composite plates across varying temperatures. They use graph attention networks to capture spatial relationships between sensors and combines Maximum Mean Discrepancy (MMD) and correlation alignment (CORAL) losses to match the source and target domains. Graph networks such as these have been successful in several transfer scenarios [27, 70, 146, 147], but can require large datasets to learn from. Sun et al. [148] employ vision-based SHM to track crack growth in concrete. Anisotropic diffusion and RANSAC is used to physically align images taken from different angles and distances over time, solving the domain shift caused by camera perspective. Giglioni et al. [149] transfer damage classifiers between two real-world bridge datasets using a Domain adversarial Neural Network (DANN). While adversarial strategies like these demonstrate potential for unsupervised label transfer, they often lack uncertainty quantification that is so important for decision-making in SHM (which is fundamentally a tool to aid decision-making) and often require large amounts of data that are not available in many real monitoring scenarios. Crucially, standard adversarial alignment techniques are designed to match the marginal distributions of the source and target domains. As highlighted by Chen et al. [150], this process often distorts the feature topology—the relative geometric ordering of data points. For classification, this topological distortion is acceptable as long as class separability is maintained. However, for regression, the relative distance between points encodes physical severity (e.g., the magnitude of stiffness loss). Adversarial forcing can effectively ‘scramble’ this metric information to achieve distributional overlap, rendering the aligned features unsuitable for continuous prognostics.

Zhang et al. [151] use CORAL for global alignment of source and target features and a contrastive loss function to pull similar faults into tight clusters for classification. This works well to aid class separability but minimising intra-class variance means that severity estimation is impossible. Enforcing extreme compactness within fault classes actively destroys the continuous topological structure required for damage quantification.

Additionally, black box models such as these lose physical interpretability, here they are aligning the covariance of abstract neural features which, in tandem with the lack of rigorous uncertainty quantification, limits decision-making capabilities. The work on DA for regression is significantly sparser than the classification literature and most existing work focuses on prognostic health monitoring (PHM) rather than structural parameter or health state estimation [152–154, 154].

3.2 Active Learning

In SHM, Active Learning (AL) is primarily employed to alleviate the scarcity of labelled data. By leveraging a predictive model to identify the most informative data points, AL minimises annotation costs while maximising model performance. In the context of inspection planning, this equates to triggering physical inspections to acquire health state labels from specific assets of interest. The foundational work by Bull et al. [48] utilises Dirichlet processes to detect novel damage classes and trigger inspections to label these ‘unknowns’. A key motivation for this framework is the requirement for AL to be ‘online’, allowing the model to adapt as data arrives throughout the asset’s life. Initially, these approaches relied on uncertainty sampling [31]. However, later work integrated Decision Theory, moving beyond pure information gain to encompass the economic and safety implications of inspection. Hughes et al. [155, 156] demonstrated that decision-theoretic approaches, where decisions should be made to maximise the expected utility, prove to be a useful tool for balancing the cost of inspection against the risk of failure. A key benefit to a decision-theoretic approach is that it naturally lends itself to budgets and explicitly defines the costs or benefits of a given action.

More recently, Poole et al. [95] introduced one of the first AL frameworks for population settings. However, a critical limitation persists across this body of work: it predominantly treats damage as a discrete state. Within SHM there is very limited research applying AL in regression settings which is essential for prognosis and severity estimation.

3.2.1 Optimal Design of Experiments

A closely related field is Optimal Design of Experiments (ODoE), which has traditionally served two distinct purposes in SHM: static sensor placement and dynamic active sensing. Regarding static design, Papadimitriou et al. [120, 157] established that sensor locations can be optimised to maximise the information content for system identification. While this handles single structures, there is no existing work on ODoE over populations

in Civil Engineering—i.e., optimising which structures to instrument in a fleet. However, such frameworks exist in biostatistics under the guise of Nonlinear Mixed Effects Models (NLME). Researchers like Tod et al. [158] and Mentré et al. [159] optimise experimental designs by expanding the Fisher Information Matrix to capture both the fixed effects (population mean) and random effects (inter-subject variability).

3.2.2 Surrogate modelling

The final intersection of AL and engineering is Adaptive Experimental Design, often termed surrogate modelling. In reliability analysis, calculating the probability of failure requires evaluating expensive Finite Element (FE) models thousands of times [160, 161]. To mitigate this, Echard et al. [162] and Sudret [163] pioneered methods like Active Kriging-Monte Carlo Simulation (AK-MCS). Instead of random sampling, these methods minimise a *learning function* to identify input configurations where the surrogate is uncertain regarding the limit state boundary. While methodologically similar to inspection planning, the objective differs fundamentally. In surrogate modelling, the ‘Oracle’ is a deterministic computer code, and the cost is computational time. In SHM, the oracle is a maintenance engineer, and the costs are those associated with inspections.

3.3 Summary

While transfer learning has been extensively applied to SHM, the literature is overwhelmingly dominated by discrete classification tasks. The few existing regression approaches typically focus on black-box remaining useful life prediction for rotating machinery, lacking the physical parameter estimation required for structural identification. Secondly, while the foundational work on Active Learning in SHM has established the utility of probabilistic querying, it has remained heavily constrained to discrete classification and novelty detection. The prevailing frameworks utilise generative models to identify new data clusters or refine the decision boundaries between ‘healthy’ and ‘damaged’ states. Even where new queries measures have been proposed, the underlying objective is typically to classify the discrete health state of an asset. There is a distinct lack of frameworks that apply Active Learning to the regression of continuous physical parameters, effectively ignoring the continuous nature of structural degradation in favour of simplified categorical states. While current approaches have often treated Transfer Learning and Active Learning in isolation—typically focusing on single systems—comparatively little attention has been paid to resource allocation and information sharing in population settings. The combination of these complimentary technologies is a natural fit, as active learning can select the most beneficial data and transfer learning can share that

information across the population. The following chapter extends active learning to the population setting, demonstrating how population-based regression models can be used to improve resource allocation across a fleet of structures.

Chapter 4

Active Learning for Regression in Engineering Populations

Data scarcity remains one of the most significant barriers to the widespread adoption of data-based structural health monitoring. As outlined in Chapter 1, acquiring labelled observations from in-service structures is costly, disruptive, and in many cases practically unfeasible — yet without sufficient data, the predictive models that underpin SHM systems cannot be trained or validated reliably. This chapter addresses the data scarcity problem through two complementary mechanisms: information sharing across a population of structures, and active learning for improved resource allocation.

The first mechanism exploits the statistical commonalities that exist between similar engineering assets. By modelling a population of structures jointly using a Bayesian hierarchical model, information from data-rich members can be shared with data-scarce ones — reducing the number of observations required to achieve reliable predictions for any individual asset. The second mechanism addresses the question of which observations are worth acquiring in the first place. Rather than inspecting assets according to a fixed periodic schedule, active learning directs inspection effort toward the observations that are most valuable given the current state of the model — and, crucially, given the economic consequences of the decisions that model supports.

Together, these two mechanisms form a unified, population-aware framework for risk-based inspection planning. The framework is demonstrated on an experimental case study involving a population of machining tools, where the quantity of interest is the surface roughness of machined workpieces — a continuous, regression target that degrades progressively over the tool’s operational life. An inspection and maintenance decision process is defined in terms of the expected utility of available actions, and this

decision process is used to construct the active learning algorithm: inspections are triggered not when uncertainty is highest, but when the expected benefit of the resulting observation — accounting for the cost of inspection and the risk of undetected degradation — justifies the expenditure.

The proposed methodology is benchmarked against two alternatives: an uninformed approach to label acquisition, and independent modelling of each regression task without population-level information sharing. It is shown that the proposed approach achieves superior performance in terms of expected cost — maintaining predictive performance across the population while significantly reducing the number of inspections required. This result demonstrates that active learning and transfer learning are a powerful combination: by jointly exploiting population structure and directing data collection intelligently, the framework makes the most of the limited data available while grounding its decisions in the practical realities of industrial asset management. The current chapter combines two avenues to address the issue of data scarcity in SHM – information sharing and active learning. Firstly, hierarchical Bayesian modelling facilitates the sharing of information between similar engineering systems, reducing the number of data required to make suitable predictions about a given system. Secondly, active learning is used to improve decision-making for inspections, only triggering inspections that will be most beneficial to the monitoring system. The proposed methodology is demonstrated using an experimental case study where multiple regressions are performed over a population of machining tools, the quantity of interest is the surface roughness of the workpieces. An inspection and maintenance decision process is defined using these regression tasks which is in turn used to construct the active-learning algorithm. The novel methodology proposed is benchmarked against an uninformed approach to label acquisition and independent modelling of the regression tasks. It is shown that the proposed approach has superior performance in terms of expected cost – maintaining predictive performance while reducing the number of inspections required.

4.1 Hierarchical Bayesian modelling

Consider data recorded from a population of K engineering structures. The population can be denoted,

$$\{\mathbf{x}_k, \mathbf{y}_k\}_{k=1}^K = \left\{ \{x_{ik}, y_{ik}\}_{i=1}^{N_k} \right\}_{k=1}^K \quad (4.1)$$

where y_k is a target response vector for inputs x_k and $\{x_{ik}, y_{ik}\}$ are the i^{th} pair of observations in group k . There are N_k observations in each group and thus $\sum_{k=1}^K N_k$

observations in total. The aim is to learn a set of K predictors related to a regression or classification task. This thesis focuses on regression, where the tasks satisfy,

$$\{y_{ik} = f_k(x_{ik}) + \epsilon_k\}_{k=1}^K \quad (4.2)$$

and the output y_{ik} is determined by evaluating one of K latent functions with separate additive noise ϵ_k . The mapping f_k is assumed to be correlated between members in the population. The models should be improved by learning the parameters in a joint inference over the whole population. In machine learning this is referred to as multi-task learning (MTL); in statistics, such data are usually modelled with hierarchical models [164, 165].

In practice, some members in a population may possess extensive historical data, while members that may have been more recently deployed will have very limited data for training. Learning separate independent models for each group might lead to unreliable results for data-poor members, while a single regression model of all the data (complete pooling) would result in poor generalisation. Hierarchical Bayesian models can learn separate models for each member, while encouraging the parameters of these models to be correlated [38]. The established theory is summarised here.

Consider K linear regression models,

$$\{\mathbf{y}_k = \mathbf{\Phi}_k \boldsymbol{\alpha}_k + \epsilon_k\}_{k=1}^K \quad (4.3)$$

where $\mathbf{\Phi}_k = [\mathbf{1}, \mathbf{x}_k]$ is the $N_k \times 2$ design matrix; $\boldsymbol{\alpha}_k$ is the 2×1 vector of weights; and the noise vector is $N_k \times 1$ and normally distributed $\epsilon_k \sim N(0, \sigma_k^2 \mathbf{1})$. $\mathbf{1}$ is a vector of ones, $\mathbf{I}$ is the identity matrix, and $N(m, s)$ is the normal distribution with mean m and variance s . The likelihood of the target response vector is then

$$\mathbf{y}_k \mid \mathbf{x}_k \sim \mathcal{N}(\mathbf{\Phi}_k \boldsymbol{\alpha}_k, \sigma_k^2 \mathbf{I}) \quad (4.4)$$

$$\therefore y_k \mid x_k \sim \mathcal{N}(\alpha_1^{(k)} + \alpha_2^{(k)} x_{ik}, \sigma_k^2) \quad (4.5)$$

following the Bayesian methodology, one can set a common hierarchy of prior distributions over the weights (slope and intercept) for each member in the population, typically normal distributions are used for the weights of each group and inverse-Gamma for the variance parameter. The normal distribution is symmetric around its peak, which works

well to represent a best guess for a centre point with symmetric uncertainty. Additionally, adjusting the variance of a normal distribution can model anything from strong certainty (a very narrow peak) to large uncertainty (a very flat, wide peak). Inverse-Gamma distributions encode the belief that the parameter you are modelling can never be negative, which is true of variances.

$$\{\alpha_K\}_{k=1}^K \stackrel{i.i.d.}{\sim} \mathcal{N}(\mu_\alpha, \text{diag}\{\sigma_\alpha^2\}) \quad (4.6)$$

$$\mu_\alpha \sim \mathcal{N}(\mathbf{m}_\alpha, \text{diag}\{\mathbf{s}_\alpha\}) \quad (4.7)$$

$$\sigma_\alpha \stackrel{i.i.d.}{\sim} \mathcal{IG}(a, b) \quad (4.8)$$

In words, Equation 4.6 assumes that the weights $\{\alpha_K\}_{k=1}^K$ are normally distributed $\mathcal{N}(\cdot)$ with mean μ_α and covariance $\text{diag}\{\sigma_\alpha^2\}$. Equation 4.7 states that prior expectation of the weights α_k is normally distributed with mean μ_α and covariance $\text{diag}\{\mathbf{s}_\alpha\}$. Equation 4.8 states that the prior deviation of the slope and intercept is inverse-Gamma distributed with shape a and scale b . A general representation of hierarchical regression can be seen in the direct graphical model in Figure 4.1.




FIGURE 4.1: A graphical model representing the linear mixed model with partial pooling.

The parent nodes $\{\mu_\alpha, \sigma_\alpha^2\}$ are inferred from the data so Equations (4.6)–(4.8) encode prior belief of the dependence between latent variables. If these parent nodes were a fixed value, rather than inferred, each model would be independent, preventing the flow of information between domains, this structure allows data-sparse domains to borrow statistical strength from those that are data rich [38].

4.1.1 Decision Theory

Hierarchical models allow one to quantify beliefs about the states of interest and do reasoning under uncertainty. In engineering these predictions are used to take actions in

the real world. Decision theory gives a framework to use the uncertainty quantification provided by a Bayesian approach to make optimal choices in many engineering problem settings.

The idea of a ‘rational’ decision-maker is defined as a decision-maker that acts to maximise the expected-utility of their actions; this is expressed via the von Neumann-Morgenstern theorem [166], which states that, for two decidable actions a and b :

$$a \succeq b \iff EU(a) \geq EU(b) \quad (4.9)$$

where $\succeq$ denotes a weak preference, indicating that a decision-maker favours a at least as much b ; and $EU(\cdot)$ denotes the expected utility associated with doing an action. In [166], the expected utility is derived. Consider a stochastic event X . Which has mutually-exclusive outcomes of which $x \in \mathcal{X}$ are conditionally dependent on a decision D between actions a and b . The expected utility of action a is computed as follows,

$$EU(a) = \int_{-\infty}^{\infty} P(X = x \mid D = a) \cdot U(X = x, D = a) dx \quad (4.10)$$

$P(X \mid D = a)$ is the probability of the outcome of X given that action a is executed. U denotes a utility function that maps $U : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$. Conveniently, if the utility of D is independent of the variable X , $U(X, D)$ can be expressed as the sum of two utility functions, $U(X)$ and $U(D)$ That separately describe the utilities associated with outcomes and actions. Equation 4.10 can then be written as,

$$EU(a) = \left[\sum_{x \in \mathcal{X}} P(X = x \mid D = a) \cdot U(X = x) \right] + U(D = a) \quad (4.11)$$

For a single decision D over a finite set of actions $\mathcal{A}$, an optimal action a^* can be defined such that the maximum expected utility (MEU) is achieved, where,

$$MEU(D) = \max_{a \in \mathcal{A}} EU(a) \quad (4.12)$$

and,

$$a^* = \arg \max_{a \in \mathcal{A}} EU(a) \quad (4.13)$$

From Equations 4.11 and 4.12, one can see an equivalence between expected utility and risk, which is defined as the product of a probability and a cost. One limitation of utility/risk-based decision theory is that it can be difficult to obtain the probability distributions in Equations 4.10 and 4.11. However, via the Bayesian framework outlined in Section 4.1 one can acquire posterior distributions, based on one's beliefs about an action, that can be used as the probabilities required for Equations 4.12 and 4.11. Additionally, the costs or utilities required for these equations can be elicited from asset-owners, allowing the expected utility of an action to be estimated.

4.2 Active Learning

To implement a Bayesian hierarchical model, such as the one proposed in Section 4.1, labelled data is required. For many engineering applications, and particularly for SHM, collecting labels for data is very expensive. It requires sending a domain expert to inspect the physical asset and often requires the operation of the asset to be temporarily halted. These costs motivate reducing the number of inspections and only inspecting measurements that would most improve the predictive model. Active learning is a set of tools that do this. Additionally, an offline active learner would not be suitable for SHM. SHM systems are required to analyse data as it arrives throughout the life of the monitored system and decisions about interventions can be required immediately upon arrival of this data. One of the major challenges of an *online* active learner in SHM is that once the decision is made *not* to query a label, access is lost to this information. It is not possible to retrieve the label once the opportunity has passed.

Consider data, $X = \{\mathbf{x}_i\}_{i=1}^N$, which have hidden labels, $Y = \{\mathbf{y}_i\}_{i=1}^N$, which can be acquired by paying the costs associated with inspection. The process of choosing and labelling these data points is referred to as *querying*. An active learner aims to learn a mapping of the observations, X , to the labels, Y , while keeping queries to a minimum. A general heuristic for active learning is presented by [121] and adapted to regression in Figure 4.2.




FIGURE 4.2: An active learning heuristic adapted from [121].

4.2.1 Risk-based Active Learning

The measure that an active learner uses to decide which unlabelled data to query is important. Information theoretic approaches use information measures such as entropy and uncertainty to guide querying, these types of approaches are common. [49] suggests a risk-based active learner for SHM. Risk-based approaches uses the expected utility of an action, or the ‘Value of Information’, to guide querying. The costs associated with interventions and monitoring are built-in directly to the decision-making process. Decisions on inspections and maintenance are made to maximise the benefit of the expected outcome. While risk-based active learning has been explored in the classification setting [155], the current chapter proposes an approach to risk-based active learning for the regression problem in which queries are guided according to the expected utility.

4.3 Case Study - A Population of Machining Tools

In this section, the framework outlined in Sections 4.2 and 4.1 will be applied to a case study. A population of machining tools will be modelled by a hierarchical Bayesian model and a risk-based decision process will be used to actively inspect the tools. The decision-theoretic approach to inspection planning will be compared to a periodic inspection plan.

4.3.1 The dataset

A dataset [41] measures deterioration over the life of machining tools during a turning process. The experimental set-up is shown in Figure 4.3. The workpiece is rotated around the dashed A-line and the tool makes four passes along the workpiece. Each pass starts at point S and ends at point E. After four passes, the tool is inspected, and measurements are taken of the workpiece and tool. This process continued until tool failure. The investigation was repeated for seven nominally identical tools, which form a population.

The deterioration is measured indirectly from the roughness of the workpiece. The surface roughness of a machined part can be used as an indicator of tool degradation because the tool’s cutting edge directly dictates the micro-topography of the workpiece. As the tool cuts the workpiece, the tool begins to wear through a number of mechanisms (flank wear, built-up edge, nose radius wear etc.). This leads to a loss of edge sharpness and a change in tool geometry. This can cause chatter and vibrations that lead to an increase in surface roughness and a deterioration in workpiece quality. The measurements from this experiment can be seen in Figure 4.4. It can be seen that, in general,




FIGURE 4.3: Schematic showing the experimental set-up used for data acquisition [41].

surface roughness increases with the distance a given tool has cut along the workpiece, this distance is termed *sliding distance*. Several tools show a sharp decrease in surface roughness at the second measurement point — this reading is believed to occur due to an initial sharpening of the tool early in the cutting process.




FIGURE 4.4: Experimental surface roughness measurements.

Because of the nature of the experiment, the measurements of surface roughness are very noisy; which can lead to robustness issues when modelling. Combined with the high noise, the shallow gradient of the deterioration makes it difficult to learn the parameters of the regression. This motivates the use of a Bayesian hierarchical model adapted from Section 4.1.

4.3.2 The hierarchical model

Because of the natural degradation of tools during the machining process, and its effect on surface finish, tools must be replaced regularly. While each tool may be produced to the same specification and made from the same materials, there will be variations in the manufacturing process that manifest as variations between the physical properties of the tools and differences in cutting behaviour. This can be an issue for standard modelling techniques where a model trained on one tool, may not generalise well to other tools. However, this variation lends itself well to a hierarchical Bayesian model because these types of model account for variations within a population while taking advantage of the statistical similarities between them. An additional benefit of hierarchical models is their suitability to the online setting and sparse datasets. This is particularly useful for tool condition monitoring where researchers may need to make predictions as soon as the machining process has begun and with only a few data points to learn a model. Additionally, the usual benefits of Bayesian modelling apply (uncertainty quantification, prior information etc.). The hierarchical framework set out in Section 4.1 is adapted below to the case study. Here, there is a population of $K = 7$ similar tools, with a target response vector y_k , the roughness measurements for each tool. The input vectors x_k are the sliding distance measurements for each tool (how far the tool has cut across the work piece). $\{x_{ik}, y_{ik}\}$ are the i^{th} pair of observations in tool k . There are N_k observations in each member and thus $\sum_{k=1}^K N_k$ observations in total. The aim is to learn a set of K predictors related to the regression task. The type of hierarchical model used for this analysis is a linear mixed model [165], so for each member in the population, a gradient m_k and intercept c_k are learnt. A graphical representation of this model can be seen in Figure 4.5.

The likelihood of this model is

$$\{y_{ik}\}_{k=1}^K \sim \text{Cauchy}(y_{\text{mean}}, \gamma_k), \quad (4.14)$$

where the location parameter is the equation of a straight line,

$$y_{\text{mean}} = m_k \cdot x_{ik} + c_k. \quad (4.15)$$

The Cauchy distribution was chosen as the likelihood because of the noisy nature of the data. Cauchy distributions are particularly suited to these types of measurements because of the larger probability density at the extremes, compared to normal distributions another more typical choice. This makes the model less susceptible to outliers




FIGURE 4.5: A graphical model representing the linear mixed model with partial pooling.

and extreme values. Following the Bayesian methodology, one can set prior distributions over the slope for the groups, which encode our prior knowledge of the parameter values.

$$\{m_k\}_{k=1}^K \sim \text{Normal}(\mu_m, \sigma_m) \quad (4.16)$$

$$\mu_m \sim \text{Gamma}(k, \theta) \quad (4.17)$$

$$\sigma_m \sim \text{HalfCauchy}(s_{\sigma_m}). \quad (4.18)$$

Where the slopes are normally distributed, with mean μ_m and standard deviation σ_m . Equation 4.17 shows the prior expectation of the slopes is Gamma distributed with shape $k = 1$ and scale $\theta = 1$. The Gamma distribution was a natural choice because we want to encode that the gradients are always positive over the life of a tool. Equation 4.18 shows that the prior deviation of the slope is HalfCauchy distributed with location parameter equal to zero and scale parameter $s_{\sigma_m} = 25$. As recommended by [67], the variance priors for this hierarchical model are set to be *weakly informative*. The priors for the intercepts are

$$\{c_k\}_{k=1}^K \sim \text{Normal}(\mu_c, \sigma_c) \quad (4.19)$$

$$\mu_c \sim \text{Normal}(\bar{\mu}_c, s_{\mu_c}) \quad (4.20)$$

$$\sigma_c \sim \text{HalfCauchy}(s_{\sigma_c}), \quad (4.21)$$

where the intercepts are normally distributed, with mean μ_c and standard deviation σ_c . Equation 4.20 shows the prior expectation of the intercepts is also normally distributed with mean $\bar{\mu}_c = 0$ and standard deviation $s_{\mu_c} = 1$. Equation 4.21 shows that the prior deviation of the intercept is HalfCauchy distributed with location parameter equal to zero and scale parameter $s_{\sigma_c} = 25$.

The systematic application of graph-theoretic algorithms has led to a number of probabilistic programming languages. Here, models are implemented in `NumPyro` [167]. The parameters are inferred using MCMC, via the no U-turn implementation of Hamiltonian Monte Carlo [59]. Throughout, the burn-in period is 1000 iterations and 2000 iterations are used for inference.

To motivate a population-based approach, the hierarchical model will be compared to a model with complete-pooling and an independent model with no pooling at all. In the complete-pooling approach, a single regression for all the data is learnt. For the independent model, a regression is learnt for each member in the population but it is assumed there is no correlation between them. Hierarchical modelling is somewhere in between, where a separate regression can be learnt for each member, while encouraging the parameters of these models to be correlated [38].

To simulate tool replacement, some measurements from Figure 4.4 will be hidden from the models, this emulates new tools with scarce data. In the following figures, the green line shows the samples from the posterior distribution over the (parameterised) latent linear functions, and the grey area shows the area in which 90% of the posterior samples fall.

Figure 4.6 shows that with complete-pooling, the model struggles with poor generalisation and the physical variations between tools means that a single regression performs poorly on many tools. Even with access to a tools full history, such as Tool 2 and Tool 4, a complete-pooling method may not perform well.

The model with no pooling, can be seen in Figure 4.7. For tools that have enough historic data, the mean and variance predictions are reasonable. However, this model performs poorly with scarce data. For these tools, the model struggles to learn the parameters of




FIGURE 4.6: Surface roughness predictions using a complete-pooling method.

the regression. The variance of these tools are over-estimated and makes poor predictions about the hidden data, when there is not enough data the model relies on vague priors. Overpredicting the variance of the surface roughness could be problematic for asset owners if they use this model to inform decisions. For example, unnecessary inspections may be triggered or tools may be replaced prematurely, increasing the costs of production.

Finally, the hierarchical model can be seen in Figure 4.8. For tools with plentiful data, this model behaves comparably to the no-pooling model. However, for data-scarce tools, there is a large reduction in the estimated variation and improvements in the predicted mean. This model is able to draw on the statistical strength of other tools to help predictions. The model *remembers* the data from other tools, captured by the prior shared between models of machining operations, and has learnt how similar tools behave.

It can be seen in Table 4.1 that the information sharing provided by a partial-pooling model can improve regression accuracy, shown through a reduction in the total mean squared error as compared to the complete-pooling and no-pooling models. It should be noted that while partial-pooling approaches can achieve improved performance across a population, for individual members in the population, the partial-pooling will occasionally score worse in terms of MSE. For example, in 'Tool 2' the surface roughness measurements are very elevated compared to any other member in the population, and returns to the population mean towards the end of tool life. The no-pooling approach over fits to this data while the partial-pooling approach does a better job at modelling




FIGURE 4.7: Surface roughness predictions with no pooling.




FIGURE 4.8: Surface roughness predictions using a partial pooling method.

the underlying process, masked by noisy measurement. Importantly, the partial pooling method makes good predictions near the end of tool life, the portion of tool life that is most critical for making decisions about replacing the tool.

TABLE 4.1: A table containing the mean squared error of partial-pooling, complete-pooling and no-pooling methods.

Pooling Type	Mean Squared Error							
	Tool 1	Tool 2	Tool 3	Tool 4	Tool 5	Tool 6	Tool 7	Total
No-Pooling	0.03539	0.00848	0.02781	0.00541	0.00930	0.08680	0.06081	0.23401
Complete-Pooling	0.03844	0.01492	0.03578	0.00716	0.00506	0.06062	0.02525	0.18724
Partial-Pooling	0.03928	0.00629	0.03167	0.00523	0.00575	0.06099	0.02677	0.17599

These results are in accordance with the results of previous work [68, 136] on the topic and motivates sharing information across a population to improve predictions.

4.3.3 Decision Process and Active Learning

Tool health deterioration can lead to tool failure. When monitoring tools, such as the ones that produced the data in this case study, a manufacturer may have an interest in avoiding tool failure because of the safety implications and costs associated with damage to the workpiece. Typically, inspections are used to observe the health state of the tool. However, performing an inspection has its own costs. Tool inspections halt production and labour and expertise are required to conduct an inspection. Optimising this decision process, whether or not to inspect an asset, can improve economic efficiency. Often there is a limited budget for inspections and ideally inspections should only be conducted when necessary. Employing decision-theoretic approach as shown in Section 4.2 to guide decision-making for inspections provides a risk-based active learning approach to more efficiently allocate inspection budgets.

A common quality-control criterion for the machining of engineering components is maintaining a surface roughness below some threshold level S_{crit} , beyond which the component is no longer fit for purpose. A high-quality surface finish can significantly improve the fatigue strength, corrosion resistance and creep life of machined parts [168]. If the surface finish is damaged via inadequate modelling or control of the machining process, the part may need to be discarded or re-machined; this has an associated utility, $C_{\text{workpiece}}$. Additionally, during the machining process an inspection can be conducted to gain access to a noisy observation of the surface roughness, with a utility $C_{\text{inspection}}$. Throughout the machining process the tool can be replaced, with a cost C_{tool} . Here, there is a decision, inaction, to allow the tool to continue machining with a risk to damage the workpiece, and action, to inspect/replace the tool.

The EU, as seen in Equation 4.10, needs to be determined for each of the possible actions, and the action with maximum expected utility should be chosen in accordance with Equation 4.13. One limitation of the dataset used in this case study is that inspections and tool replacements can only be triggered at specific discrete time steps because the data was collected at periodic intervals, a dataset with less restrictive inspections points is under production and will be part of future work.

The expected utility of three actions need to be considered to evaluate this decision process. Inaction, to do nothing and allow the tool to continue machining, to inspect the tool and/or to replace the tool. If an inspection occurs, there is another opportunity to decide to replace the tool based on the new information acquired from the inspection.

Again, Equation 4.10 can be used to calculate the expected utility of a decision. To compute the utility associated with inaction, *not* inspecting the tool at time step t , one can use Equation 4.22. The two outcomes of inaction are the surface roughness reaching or exceeding S_{crit} before the next opportunity to intervene, time step $t + 1$, or the surface roughness *not* reaching or exceeding S_{crit} . The utility of not exceeding S_{crit} is assigned a value of zero, so the second term in Equation 4.22 equals zero. The probability $P_{t+1}(S > S_{\text{crit}})$ can be estimated from the model; here it is the proportion of samples from the Hamiltonian Monte Carlo simulation that exceed S_{crit} before the next opportunity to inspect.

$$EU(D = \text{do nothing}) = P_{t+1}(S > S_{\text{crit}}) \times U(S > S_{\text{crit}}) + P(S < S_{\text{crit}}) \times U(S < S_{\text{crit}}) \quad (4.22)$$

When computing the utility of inspection at a given time step, again there are two outcomes based on whether the surface roughness will exceed S_{crit} . Again, probabilities of the outcomes can be estimated from the HMC samples. The equation to compute the EU of inspection can be seen in Equation 4.23. It includes the probability that the tool needs to be replaced, which is the probability that S is already greater than S_{crit} at that time step, as well as the probability that S is *not* greater than S_{crit} but *will be* by the next time step. For the current linear model, this is equivalent to $P_{t+1}(S > S_{\text{crit}})$ and could be evaluated a such from the HMC samples. Again, these probabilities are estimated from HMC samples.

$$EU(D = \text{inspection}) = C_{\text{tool}} \times (P_{t+1}(S > S_{\text{crit}})) + C_{\text{inspection}}. \quad (4.23)$$

The criteria required to trigger the replacement of a tool can be seen in Equation 4.24. It is the probability at which the risk associated with damaging the work piece (the probability of damage occurring multiplied by the utility associated with it) becomes greater than the cost of replacing the tool.

$$\begin{aligned}
 P(T_f < T) \times C_{\text{workpiece}} &\geq C_{\text{tool}} \\
 P(T_f < T) &\geq \frac{C_{\text{tool}}}{C_{\text{workpiece}}} \\
 P(T_f < T) &\geq \alpha.
 \end{aligned} \tag{4.24}$$

Equation 4.24 shows that tools should be replaced at the earliest time, T , that $P(T_f < T)$ is equal to or greater than the ratio, α . $P(T_f < T)$ is the assessment of time-to-failure after one has updated beliefs with the results of any inspections and where failure is S exceeding S_{crit} . Intuitively, as C_{tool} increases, the probability requirements are increased and so replacements are more difficult to trigger, the system prioritises extending tool life. As $U(S > S_{\text{crit}})$ increases, the probability requirements are reduced and so replacements are more easily triggered and the system prioritises the surface quality of the workpiece.

A complete decision-theoretic approach would consider the expected value of information associated with improving the model based on the new data from an inspection. For the current work, the full value of information calculation is not included because it is very computationally expensive. It is assumed that inspecting the tool would provide zero improvement to the predictions of the model, meaning that the utility of inspecting the tool is understated in this analysis. There is additional utility that is not included from the improvements of the predictive capabilities of the model with the additional data from the inspection. This leaves a systemic bias in this approach that undervalues inspections. The inclusion of the full value of information analysis is left for future work.

At every potential inspection point, the hierarchical model makes predictions about the surface roughness. These predictions are based on the currently available data and are used to inform the decision analysis detailed above. A diagram showing this process is shown in Figure 4.9.

The online decision-theoretic active learning approach for inspection planning presented in this work is compared to a conventional inspection plan featuring periodic inspections. With both approaches, the model has access to the full measurement history of four of the tools and very limited data for the rest of the tools (in this case, without access to at least four full tool datasets, the hierarchical model does not have enough data to draw on for accurate predictions).




FIGURE 4.9: The decision theoretic active-learning heuristic.

In general, the parameters of the decision process, $S_{crit}, C_{inspection}, C_{tool}, U(S > S_{crit})$ can be elicited from expert familiar with the inspection process. Here, the values are set by hand to reflect a situation common in industry, where the cost of damaging the workpiece is much greater than the cost of replacing the tool¹:

$$\begin{aligned}
 U(S > S_{crit}) &= 1 \\
 C_{tool} &= 0.25 \\
 C_{inspection} &= 0.05 \\
 S_{crit} &= 0.9\mu m
 \end{aligned} \tag{4.25}$$

To assess the performance of these monitoring systems, each will be compared to a monitoring system with access to the full measurement history of every tool. This monitoring system will be seen as the gold standard as it represents the limiting case of using all possible information within the dataset and thus the point at which it decides to replace the tool can be considered to be the optimal point of replacement. If the other monitoring systems choose to replace the tool later than the optimal point, it is assumed the roughness has exceeded S_{crit} (despite what the noisy measurements might suggest) and the workpiece will be considered damaged with a cost $C_{workpiece}$. If the monitoring systems choose to replace the tool earlier than the optimal replacement time, then some

¹It is worth noting that optimal actions are invariant under affine transformations of the utility function

portion of the tool life is wasted, the cost of which can be calculated using Equation 4.26. The optimal monitoring system can choose to replace the tool at any point throughout the life of the tools i.e. it can observe several measurements that are greater than S_{crit} and then realise it should have replaced the tool before these measurements. The other monitoring systems are performing in an online manner, so do not have this luxury. The combined costs of both inspections and sub-optimal replacements, will be compared.

$$C_{\text{early replacement}} = \frac{t_{\text{replacement}}}{t_{\text{optimal replacement}}} \times C_{\text{tool}} \quad (4.26)$$

As mentioned previously, the first four tools are to be seen as historic data that the manufacturer has collected from previous tools. The hierarchical model can leverage this data to inform the population-level distributions. The active-learning procedure tool replacements are not calculated for these tools, only Tools 5-7. During the original experiment the measurements were taken at a time step that relates to 6.02 km sliding distance [41]. Figure 4.10 shows the optimal replacement for Tools 5-7 based on the “gold standard” model.




FIGURE 4.10: Benchmark replacements determined using all available information.

4.4 Results

In this section the decision theoretic active learner described in Section 4.3 and a periodic approach to inspection planning will be compared. The point at which these approaches

suggest to replace Tools 5-7, based on the criteria showcased in Equation 4.24, will be compared to the “gold standard” model, for which, the suggested tool replacements are shown in Figure 4.10.

Figure 4.11 shows the suggested tool replacements with periodic inspections. It can be seen that, compared to the optimal replacements in Figure 4.10, every tool is replaced one time step too early, with a total of ten inspections. This result is to be expected because in forecasting the surface roughness predictions with reduced data (as compared to the “gold standard” model which has access to all the information), the uncertainty is inflated thus increasing the estimated risk associated with damage to the workpiece.




FIGURE 4.11: Tool replacement chosen with periodic inspections compared to gold standard replacements.

Figure 4.12 shows the replacements with risk-based inspection planning. Tools 5 and 6 are again inspected one time step before the fully observed case. Additionally, Tool 7 is inspected two time steps before the gold standard. The risk-based monitoring system used a total of five inspections. The reduced number of inspections is because the decision-theoretic approach will, in general not suggest inspections early in the life of a tool when the risk of S exceeding S_{crit} is small. Most inspections will be triggered near the end of tool life when the risk associated with damaging the work piece is greatest. Again, because of the reduced access to information (compared to both the “gold standard” and the periodic inspections), inflated risk increases the likelihood of pre-optimal tool replacements.




FIGURE 4.12: Tool replacement chosen with risk-based inspections compared to gold standard replacements.

Table 4.2 collates the number of inspections and tool replacements. It should be noted that the unit '0.602kms' refers to the sliding distance (how far the tool has cut) between each surface roughness measurement.

TABLE 4.2: A table containing the inspection values.

Inspection Type	Tool replaced at (0.602 kms)			Discrepancy (0.602 kms)			Number of Inspections
	Tool 5	Tool 6	Tool 7	Tool 5	Tool 6	Tool 7	
Full Information	10	10	11	0	0	0	0
Periodic Inspections	9	9	10	1	1	1	10
Risk-based Inspections	9	9	9	1	1	2	5

Table 4.3 compares the performance of the inspection approaches. The 'Cost of inspections' column is the number of inspections multiplied by the cost of an inspection. The 'Cost of Wasted Tool Life' for each inspection approach can be seen in the first column of Figure 4.3 and can be calculated using Equation 4.27.

$$C_{wasted\ tool\ life} = C_{tool} \times \frac{\text{Discrepancy}}{\text{Optimal Replacement}} \quad (4.27)$$

As can be seen in the final column of Table 4.3, using the parameters described in this section, a risk-based approach to inspection planning reduced costs associated with monitoring by 36.95%.

TABLE 4.3: A table containing the monitoring and maintenance costs of each inspection method.

Inspection Type	Cost of Wasted Tool life			Cost of inspections			Total Cost
	Tool 5	Tool 6	Tool 7	Tool 5	Tool 6	Tool 7	
Periodic Inspections	0.025	0.025	0.028	0.15	0.15	0.15	0.573
Risk-based Inspections	0.025	0.025	0.056	0.15	0.05	0.05	0.361

4.5 Discussion

The case study presented in Section 4.3 highlights the effectiveness of a risk-based approach to inspection planning for a population of continuously degrading systems. The Bayesian hierarchical model was used to predict surface finish — a continuously varying quantity that degrades progressively over the operational life of each tool — and it is this continuous, regression-based characterisation of health state that makes the risk-based decision process possible. By maintaining a probabilistic estimate of the current and future surface roughness of each tool, the framework can evaluate the expected utility of inspection and replacement actions in a principled way, quantifying the risk of continued operation against the cost of intervention. This is a qualitatively different capability from classification-based approaches, which can only associate damage to discrete categories — they cannot quantify how close it is to a given threshold, nor how quickly it is approaching it.

The significant reduction in monitoring costs can be attributed to the elimination of unnecessary inspections — those that would not meaningfully improve predictions or alter the recommended action — while maintaining a rate of workpiece damage comparable to that of a periodic inspection strategy. This result underscores a key advantage of decision-theoretic query selection: by anchoring inspection decisions to the expected consequences of action and inaction, the framework avoids the inefficiency of inspecting simply because uncertainty is high, or the time since the previous inspection is too long.

The formulation of the expected utility equations is not always straightforward. A thorough understanding of the decision problem is required to derive expressions that faithfully reflect the costs and risks of the application, and simplifications are often necessary to make the calculations tractable in real time. This is a meaningful barrier to adoption in novel settings, and represents a trade-off inherent to the decision-theoretic approach: the richer the decision model, the more problem-specific expertise is required to construct it.

The choice of a Bayesian hierarchical model introduces further practical considerations. The hierarchical structure allows the continuously evolving surface finish of each tool

to be modelled jointly with the population, such that the degradation trajectory of a data-scarce tool is regularised by the behaviour observed across the fleet. This partial pooling increases the dimensionality of the posterior that a Monte Carlo sampler must explore, which can lead to complex posterior geometries, increased computational costs, and restrictions on the choice of prior distributions. In online monitoring scenarios where decisions must be made instantaneously alongside data acquisition, these costs may be prohibitive. Additionally, the full Value of Information analysis was omitted from the algorithm as presented — its inclusion would provide a more theoretically complete basis for the inspection decisions but would add further computational burden. These considerations motivate future work on scalable inference strategies that preserve the decision-theoretic structure of the framework while reducing its computational footprint. It is also worth reflecting on the generality of the risk-based query strategy. The framework was demonstrated in the context of machining tool degradation, where the cost structure — inspection cost, replacement cost, and the cost of workpiece damage — is well-defined and quantifiable, and where the health variable of interest is a single continuous scalar. In other SHM applications, the health variable may be higher-dimensional, the cost structure harder to specify, or the degradation trajectory less regular. Nevertheless, the underlying principle — that inspections should be directed by the expected utility of their outcomes rather than by information gain alone — is broadly applicable to many real engineering decision-making tasks.

4.6 Conclusion

This chapter has presented a risk-based active learning framework for population-level inspection planning, demonstrated on a population of machining tools undergoing continuous surface degradation. Central to the framework is the use of a Bayesian hierarchical model to predict surface finish — a continuously varying regression target — across the population. By maintaining a probabilistic estimate of the surface roughness of each tool throughout its operational life, the framework is able to evaluate the expected utility of inspection and replacement decisions in a principled and quantitative way. This continuous, regression-based characterisation of structural health is what distinguishes the approach from classification-based alternatives: rather than detecting when a threshold has been crossed, the framework tracks how the health variable is evolving and acts when the predicted trajectory makes intervention economically justified.

The case study demonstrated that this approach yields a significant reduction in monitoring costs relative to both periodic inspection and uninformed label acquisition strategies, without sacrificing predictive performance across the population. In several instances,

the risk-based approach achieved superior predictive performance as well as reduced cost — a consequence of directing inspections toward the observations most consequential for decision-making rather than those that merely reduce model uncertainty in a general sense. This distinction is important: in engineering applications where inspections carry real economic and operational costs, a query strategy grounded in decision theory is more naturally aligned with the objectives of the monitoring system than one based purely on information gain.

The results establish a clear proof of concept for the combination of population-level information sharing and risk-based active learning for continuous health monitoring. By jointly exploiting the statistical structure of the population and directing data collection intelligently toward the most consequential observations, the framework achieves data-efficient monitoring that is grounded in the practical economics of asset management. In doing so, this chapter fulfils **Contribution 2** as defined in Section 1.4: expanding active learning to the population monitoring context for regression tasks, with explicit consideration of the economic consequences of inspection decisions.

The dataset outlined in Section 4.3 was produced from a population of nominally identical machining tools, where health state data — in the form of surface roughness measurements — were collected intermittently and required the tools to be removed from service to be inspected. While sufficient to demonstrate the proposed framework, this setting represents a relatively favourable case for transfer. Additionally, the discrete sampling of health states limits the richness of the regression target. Chapter 5 addresses both of these constraints directly, introducing a new PBSHM benchmark dataset in which the population exhibits more significant heterogeneity, and in which damage state data is available continuously throughout the experiment — providing a dataset designed specifically for the development and evaluation of regression-based transfer learning.

Chapter 5

Continuous Damage in Engineering Populations

Experimentation is a cornerstone of scientific and engineering progress. Beyond design validation and the investigation of complex physical phenomena, experiments play an essential role in machine learning: they generate the empirical data needed to train, test, and evaluate models in realistic conditions. In SHM, this is particularly consequential — models are only as reliable as the data they are trained on, and data that poorly reflects in-service structural behaviour will produce models that fail when deployed on real assets.

A recurring limitation of existing SHM benchmark datasets is that they capture structural behaviour through discrete realisations of damage states. These include the Z24 bridge [19], the Hawk T1A aircraft [20], and the machine tool population dataset used in Chapter 4: in each case, structural outputs are measured, a damage condition is applied — typically the addition of masses or removal of material — and outputs are measured again. Figure 5.1 illustrates this for the Z24 bridge, and Figure 5.2 shows the pseudo-damage locations on the Hawk airframe.

Most structural damage does not manifest as a sequence of discrete states. Degradation mechanisms such as bolt loosening, foundation deterioration, and the progressive failure of welded joints evolve continuously over time, and models trained on discrete datasets learn to recognise predefined damage conditions rather than to track gradual deterioration between them. This motivates the development of regression-based methodologies that treat structural health as a continuous variable.

Despite this, publicly available datasets capturing continuous structural degradation are scarce, and to the author’s best knowledge, none exist that include a population of

to three different lengths. The beams were excited with white noise through a shaker, clamped normally at one end and clamped with a large linear actuator at the other. The clamping force provided by the linear actuator was controlled and used in this experiment to simulate damage, whereby the clamping force was slowly reduced over the experiment. The clamping force can be seen as a health state of the system, at the start of the experiment the clamping force is large and the system is assumed to be healthy. Then, as the clamping force reduces, the system becomes increasingly damaged.

5.1.1 Initial Design

A preliminary mock-up was created in Autodesk Fusion 360 [169] and can be seen in Figure 5.3. The design should limit dynamic interaction between the external environment and the beam. Additionally, it should be easy to remove and replace the beams when it is time to test the next specimen.

One of the critical choices in the design of the experiment was the length of the beams, this decides several other dimensions in the experimental set-up.

5.1.2 The beams

The beams have a cross-section with a height of 6mm and width of 12mm, a readily available cross-section from local retailers. Three lengths of beams were tested, 480mm, 550mm and 620mm. These lengths were chosen such that the beams had suitable dynamic properties and based on the limitations of the testing space.

To define parameters for the experiment, the dynamic properties of the beams were estimated. Using the derivations from [170], the natural frequencies of a beam can be estimated from Equation 5.1

$$f_i = \frac{\lambda_i^2}{2\pi L^2} \left(\frac{EI}{m} \right)^{1/2} \quad (5.1)$$

where f_i is the i -th natural frequency, E is the Young's modulus of the material, L is the length of the beam, m is the mass per unit length of the beam, I is the second moment of area of the beam and λ_i is a mode dependent constant related to the boundary conditions. The second moment of area of the beams can be calculated using the following equation:

$$I = \frac{bh^3}{12} = \frac{0.012 \times 0.006^3}{12} = 2.16 \times 10^{-10} \quad (5.2)$$




FIGURE 5.3: A CAD mock-up of the experimental set up.

The material properties and the estimated natural frequencies of the beams are collated in Figure 5.1

The natural frequencies in Table 5.1 are based on clamped-clamped boundary conditions. At the start of the experiment when the linear actuator is providing a large clamping force, the beams will behave similarly to this.

TABLE 5.1: Material properties and the estimated natural frequencies of the beam population.

Material	Length (m)	Young's Modulus (GPa)	m (kg/m)	ω_{n1} (Hz)	ω_{n2} (Hz)	ω_{n3} (Hz)
Stainless Steel	0.48	210	0.576	137.1	378.1	741.14
Stainless Steel	0.55	210	0.576	104.5	287.9	564.5
Stainless Steel	0.62	210	0.576	82.2	226.6	444.2
Mild Steel	0.48	200	0.565	135.1	372.5	730.2
Mild Steel	0.55	200	0.565	102.9	283.7	556.1
Mild Steel	0.62	200	0.565	80.9	223.2	437.6
Brass	0.48	100	0.612	91.8	253.1	496.2
Brass	0.55	100	0.612	69.9	192.8	337.9
Brass	0.62	100	0.612	55.0	151.7	297.4
Aluminium	0.48	69	0.1944	135.3	373.0	731.3
Aluminium	0.55	69	0.1944	103.1	284.1	557.0
Aluminium	0.62	69	0.1944	81.1	223.6	438.3

5.1.3 The clamping force

Using a linear actuator as one of the clamps for the dynamic testing meant the boundary conditions could be controlled. The clamping force regime can be seen in 5.4. It includes a period at the start where the clamping force remains at 50kN for five minutes, before linearly reducing to 0kN over 20 minutes. The first five minutes allows the use normal condition alignment [171], which has been shown to be a crucial pre-processing step for transfer learning [143]. Data from in-service structures can usually be assumed to be healthy for the first part of its life and used to initiate a mapping for data from different domains (this does not hold for retrofitted assets).

5.1.4 Final set-up

Five accelerometers were placed evenly along the beam at 1/6, 2/6, 3/6, 4/6 and 5/6 the length of the beams, sufficient coverage to calculate the dynamic properties of the beams. To recover higher order modeshapes, a greater number of accelerometers would be required. To limit the damping effects of the accelerometers, the wires were taped to a supporting structure, such that the weight of the wires was supported outside of the immediate experimental set-up. To excite the beams, a LDS permanent magnetic shaker was used with a white noise input. Figure 5.5 shows the experimental set-up used to collect the data in this chapter. The data for each beam was recorded three times for the purposes of averaging and each test ran for 25 minutes.




FIGURE 5.4: The clamping force regime, showing an initial ‘healthy-state’ at 50kN before reducing monotonically to 0kN.

5.2 Data Processing

The system under investigation was a Single-Input Multiple-Output (SIMO) configuration. Frequency response functions (FRFs) were estimated using the H1 estimator, which is appropriate for systems with noisy outputs and provides an unbiased estimate of the transfer function [172]. The H1 estimator is defined as:

$$\hat{H}_1(\omega) = \frac{G_{xy}(\omega)}{G_{yy}(\omega)} \quad (5.3)$$

where $G_{xy}(\omega)$ is the cross-power spectral density between the response X and the input y , and $G_{yy}(\omega)$ is the auto-power spectral density of the input. The resulting FRFs were then processed using the Poly-reference Least-Squares Complex Frequency-domain (p-LSCF) method [173] to extract the system poles, enabling robust identification of natural frequencies and damping ratios and modes, by fitting the FRFs in a least-squares sense [174].




FIGURE 5.5: A picture of the experimental set-up used to collect the data in this chapter.

5.3 First analysis

The following analysis characterizes the structural dynamics of the beams under the evolving clamping force detailing the primary signal processing and parameter estimation steps applied to the raw experimental data. The analysis proceeds in three stages: first, the input excitation is characterized to ensure signal quality; second, the system's FRFs are derived to visualize spectral trends; and finally, a modal analysis is performed

to extract and track the evolution of natural frequencies and damping ratios as the boundary conditions degrade.

5.3.1 The input

The shaker was driven by a band-limited white noise signal over a frequency range of 0–2000 Hz, to avoid aliasing the upper band must be less than half the sampling frequency. This provides a broadband excitation sufficient to excite the first several bending modes of each beam. A representative input signal measured by the force transducer mounted between the shaker stinger and the beam is shown in Figure 5.6. The force transducer measurement ensures that the true input to the structure is recorded rather than the commanded signal, accounting for any dynamics introduced by the stinger and shaker armature.




FIGURE 5.6: The input spectra from the shaker. Highlights the interaction between the shaker and the beam, with drop out around the natural frequencies of the beam.

During experimental modal testing, the interaction between the exciter and the structure can lead to degradation of the measured input force, particularly in the vicinity of the structure's natural frequencies. As the structure approaches resonance, its dynamic impedance decreases, resulting in large structural motion relative to the exciter. This can cause partial loss of contact, non-linear behaviour in the stinger, or saturation effects in the shaker, leading to a reduction or distortion of the measured input force, commonly

referred to as force drop-out [174]. Such effects violate the assumption of a well-defined, broadband excitation and can reduce the accuracy of frequency response function (FRF) estimates, particularly for estimators that assume noise-free input measurements [172]. The issue is exacerbated in lightly damped structures and becomes more pronounced as boundary conditions degrade, increasing structural compliance and reducing force transmissibility at resonance. Care must therefore be taken when interpreting modal parameters extracted near resonance.

5.3.2 Frequency Response Functions

Frequency response functions (FRFs) provide a fundamental representation of the input–output behaviour of a linear dynamic system and form the basis of most vibration-based structural health monitoring techniques [172]. In this study, FRFs are used as an initial diagnostic tool to assess the quality of the experimental data, to verify that the dynamic behaviour of the beam population is physically meaningful, and to examine how progressive boundary degradation manifests in the measured responses.

For each beam, the system was configured in a SIMO arrangement and FRFs were estimated using the H1 estimator, which is appropriate for systems with noisy outputs and provides an unbiased estimate of the transfer function [172].

Figure 5.7 shows a representative magnitude and phase FRF measured at the mid-span accelerometer for a single beam at the start of the experiment: an initial healthy state with a high clamping force. This is the stainless steel beam of length 550mm.

The resulting FRFs exhibit clear resonance peaks and smooth phase transitions, indicating adequate excitation across the frequency range of interest.

Figure 5.8 shows how the FRF develops as boundary conditions change. A zoomed-in look at one of the resonance peaks can be seen in Figure 5.9. As the clamping force reduces, the peak of the resonant peak shifts, showing a reduction in the natural frequencies.

Importantly, the evolution of the FRFs occurs smoothly over the duration of the experiment, with no abrupt changes in resonance structure. This indicates that the imposed damage mechanism produces a continuous degradation process rather than a sequence of discrete damage states. Such behaviour contrasts with many benchmark SHM datasets, where FRFs are typically obtained only at a small number of predefined damage conditions. The continuous evolution observed here provides a more realistic representation of in-service structural degradation and motivates the need for analysis methodologies capable of handling non-stationary and gradually evolving dynamics.




FIGURE 5.7: A ‘healthy state’ magnitude (top) and phase (bottom) FRF of a representative beam.

It can be seen that the variation in the FRFs as the clamping force decreases increases with frequency (frequencies greater than 800Hz show larger variation). This holds true across the population, where higher modes are more sensitive to damage. At low frequencies, only the highest levels of damage induce large variations, whereas at high frequencies medium levels of damage can cause large variations. This can be seen particularly clearly in Figure 5.10 where the magnitude FRFs from all the accelerometers are shown.

The way in which the population varies with length can be seen in Figure 5.11. It shows the magnitude FRF for stainless steel beams of 480mm, 550mm, and 620mm from one accelerometer.

The way in which the population varies with length can be seen in Figure 5.12. It shows the magnitude FRF for all materials in the dataset at a length of 620mm from one accelerometer. The FRF’s for teh entire beam population can be found in Appendix A.




FIGURE 5.8: Evolution of the frequency response function of a representative beam as a function of clamping force. The colour indicates the clamping force, progressing from blue (healthy) to yellow (damaged).

5.3.3 Modal Analysis

Following the estimation of the FRFs, the dynamic properties of the beam population were extracted using frequency-domain modal analysis techniques. Modal analysis enables the decomposition of the measured structural response into a set of modal parameters, namely natural frequencies, damping ratios, and mode shapes. These parameters provide a compact and physically interpretable description of the structural dynamics and are commonly used as indicators of structural degradation.

5.3.3.1 Natural Frequencies

p-LSCF formulates a parametric model of the FRFs in the frequency domain and estimates the system poles by fitting a common-denominator polynomial across all measured




FIGURE 5.9: A zoomed-in look at the resonance peak of the frequency response function. The colour indicates the clamping force, progressing from blue (healthy) to yellow (damaged).




FIGURE 5.10: Evolution of the frequency response function of all accelerometers of a representative beam as a function of clamping force. The colour indicates the clamping force, progressing from blue (healthy) to yellow (damaged).

FRFs. This approach is well suited to multi-output experimental data and provides robust pole estimates in the presence of noise. The first three bending modes were consistently identifiable across all beams, these modes were tracked throughout the experiment as the clamping force was gradually reduced.




FIGURE 5.11: Evolution of the frequency response function for stainless steel beams of length 480mm (top), 550mm (middle) and 620mm (bottom). The colour indicates the clamping force, progressing from blue (healthy) to yellow (damaged).

Figure 5.13 shows the evolution of the first bending natural frequency of a representative beam as a function of clamping force. A progressive reduction in natural frequency is observed as the clamping force decreases from 50kN to 0kN, reflecting a reduction in the effective stiffness of the system due to the softening of the boundary conditions. This behaviour is consistent with the transition from an approximately clamped-clamped configuration towards a more compliant boundary condition as the actuator force decreases.

5.3.3.2 Damping Ratios

Damping ratios reflect the energy dissipation mechanisms present in the structure. Figure 5.14 shows the variation of the damping ratio for the first three bending mode across the population of beams as the clamping force is reduced. After detecting physical poles in the stabilisation diagrams from p-LSCF, these damping ratios are derived from solving the eigenvalue problem of the denominator matrix. The damping ratios are the ratio




FIGURE 5.12: Evolution of the frequency response function across 4 materials of length 620m. The colour indicates the clamping force, progressing from blue (healthy) to yellow (damaged).

of the real part of the pole to its total magnitude. It can be seen that the damping ratios seem to be unchanged by damage. The damage induced in this experiment is primarily expressed as a reduction in stiffness at the boundary, which would not have an effect on the damping in the system. The system damping seems dominated by material specific damping, which seems logical for metal beams with comparatively low levels of excitation.

5.3.3.3 Mode Shapes

The five accelerometers positioned evenly along the beam provided sufficient spatial resolution to reliably identify the first three bending mode shapes. Figure 5.15 shows the first three bending mode shapes of a representative beam without any damage.




FIGURE 5.13: Evolution of the first natural frequency of a representative beam as a function of clamping force. The colour indicates the clamping force, progressing from blue (healthy) to yellow (damaged).

Figure 5.16 shows the first three bending mode shapes of a representative beam throughout the experiment as damage is applied. The first two mode shapes behave exactly as one would expect a clamped-clamped beam to behave. The first mode shape presents the typical ‘half-sine’ shape, a bowed line with the peak in the middle. The second mode shape showcases ‘two half-waves’, the characteristic ‘S-shape’ where there is one internal nodal point near the midspan. The two halves of the beam move in opposite directions — one half moves up while the other moves down. The third mode shape has a larger spread, in some instances it looks similar to the first mode, while in others it conforms to the behaviour one would expect from a clamped-clamped beam where forms the characteristic ‘M-shape’. In the second and third modes especially, the mode shapes invert at higher levels of damage.

The boundary conditions for a clamped clamepd beam are as follows:

$$\phi(0) = 0, \quad \left. \frac{d\phi}{dx} \right|_{x=0} = 0 \quad (5.4)$$

$$\phi(L) = 0, \quad \left. \frac{d\phi}{dx} \right|_{x=L} = 0 \quad (5.5)$$




FIGURE 5.14: Damping ratios for the first three bending modes as a function of clamping force.




FIGURE 5.15: First three mode shapes of a steel beam of 550mm with no damage.

At each end of a clamped-clamped beam the displacement and slope are both zero. As damage increases, the slope boundary condition is progressively violated, consequently the eigenvalue problem is also changing which leads to the mode shapes changing.




FIGURE 5.16: Evolution of the first three mode shapes of a steel beam of 550mm as a function of clamping force.

5.4 Damage features

Structural damage typically manifests as subtle changes in a structure's dynamic response. While raw vibration measurements and frequency response functions (FRFs) contain rich information about the underlying system, their high dimensionality and sensitivity to noise make them unsuitable for direct use in learning-based models. As a result, the selection of informative and robust features is a critical step in the development of data-driven approaches for damage quantification.

An effective feature set must consider two important requirements. Features must be sensitive to damage-related changes in structural properties, such as stiffness reduction or increased dissipation. Additionally, they must remain sufficiently robust to measurement noise, operational variability, and modelling uncertainty to enable reliable learning. This is particularly important when features are used for regression tasks, where smooth and continuous relationships between features and damage are required.

5.4.1 Principal component analysis

Principal Component Analysis (PCA) [175, 176] is a data-driven dimensionality reduction technique used to identify dominant patterns of variability within a dataset. In structural health monitoring, data such as raw measurements, frequency response functions, or spectral descriptors—are often high-dimensional and strongly correlated. This

redundancy can obscure damage-related information and negatively impact the performance and interpretability of models trained on this data. PCA addresses this issue by projecting the original feature space onto a lower-dimensional subspace spanned by orthogonal basis vectors, known as principal components, which are ordered according to the amount of variance they explain in the data. PCA provides a compact representation that preserves the most significant variations and can help researchers spot patterns in data.

PCA has been used to extract damage sensitive features from FRFs extensively in the literature [177–179]. In this case, PCA identifies the frequency components at which the beam’s dynamic response exhibits the greatest variability across experimental conditions. Figure 5.17 shows the first three principal components from the analysis.




FIGURE 5.17: Evolution of the first three principal components of the FRF of a steel beam of 550mm as a function of clamping force.

The principal components shown in Figure 5.17 are highly correlated with damage and could make useful features for transfer learning and active learning tasks. The PCA

algorithm is finding a linear combination of frequencies in the FRF that show the most variation over the experiment. It is an indicator that the experiment has been a success that the variations found by the algorithm are related to damage rather than noise.

5.4.2 Limitations of the experiment

The most impactful limitation of this set-up was the shaker-beam interaction, which impeded the quality of the data and limited the usefulness of the modal properties extracted from the system. The shaker was bolted in firmly to the rest of the structure to allow the position of the shaker to be adjusted with the size of the beams as to maintain the shaker position of $1/3$ the length of the beams. In future a new design will be required to isolate the shaker from the rest of the structure, either some padding could be used to reduce the transmission of vibrations to the rest of the structure (other than through the shaker-beam interface) or the shaker should be suspended in some other way. Despite best efforts, the structural frame that supported the beams in the clamped-clamped position might have had small deflections during testing. During moments of high deflection of the beams, the beams would “pull” the towers inwards, affecting the results and moving the set-up further from an ideal clamped-clamped boundary. To limit the effect of the frame collapsing, more bracing should have been used on the support beam to avoid it caving in.

Additionally, large amplitudes of excitation would cause other issues such as slipping in the clamps, leading to large non-linearities that break the linear assumptions made during the analysis. Additionally, due to only 5 accelerometers being used to capture beam response (space limitations) the set-up would struggle to capture more than two modeshapes to an acceptable level of resolution.

5.5 Conclusion

The experiment presented in this chapter was designed to address a specific and consequential gap in the PBSHM literature: the absence of publicly available datasets capturing continuous structural degradation across a population of structures. The dataset collected here — comprising frequency response measurements from a population of beams of varying material and geometry, subjected to progressive boundary condition degradation — represents a departure from the discrete, artificially induced damage states that characterise the majority of existing SHM benchmarks.

The modal analysis results presented in this chapter confirm that the experimental design achieved its primary objective. Natural frequencies decreased monotonically with

clamping force across the population, providing a clear and physically interpretable continuous damage variable. Crucially, this trends was observed consistently across beams of different materials and lengths, suggesting that the underlying degradation signature is transferable across the population despite geometric and material heterogeneity. This is an important result for the transfer learning work presented in Chapter 6, as it provides empirical evidence that the population shares meaningful statistical structure that can be exploited.

The choice of boundary condition loosening as the damage mechanism warrants reflection. Compared to mass addition or section removal — the most common approaches in existing datasets — progressive clamping force reduction offers several advantages. It is reversible, continuously controllable, and physically representative of a class of real degradation mechanisms including bolt loosening, foundation deterioration, and joint degradation. The resulting damage variable, the clamping force, is also directly measurable, providing a ground truth label for regression that is free from the ambiguity associated with discrete damage classifications. At the same time, this mechanism is a simplification of real structural degradation, and the limitations this imposes on the generalisability of methods developed on this dataset are discussed in Chapter 7.

Several limitations of the experimental programme warrant acknowledgement. The population was restricted to beams sharing a common rectangular cross-section and a single degradation mechanism, which does not reflect the full diversity of structural variation encountered in real asset populations. The experiment was conducted under controlled laboratory conditions, excluding the environmental and operational variability — temperature fluctuations, varying loading conditions, and measurement noise from operational sources — that routinely confound damage detection in field-deployed systems. Furthermore, the sensor configuration was fixed across all beams, whereas in practice the number and placement of sensors will differ between population members, introducing a further source of heterogeneity that the current dataset does not represent.

Future experimental campaigns should address these limitations by considering a broader range of cross-sectional geometries, boundary condition types, and damage mechanisms — including scenarios in which degradation manifests differently across population members, better reflecting the heterogeneity that population-based methods must ultimately contend with. Expanding the range of clamping force regimes and degradation rates would allow the sensitivity of the developed methods to different damage trajectories to be more thoroughly characterised. The introduction of controlled environmental variability — for instance, through temperature cycling — would provide a more realistic and challenging benchmark for methods that must distinguish damage from confounding influences. Finally, varying the sensor configuration across population members

would allow the heterogeneous transfer problem to be studied in a more experimentally grounded setting, providing a natural extension of the transfer learning methodology presented in the following chapter.

This chapter fulfils **Contribution 3** outlined in Section 1.4: providing the community with one of the first benchmarks capturing continuous structural degradation across a heterogeneous population of structures — providing the community with a foundation for the development and evaluation of regression-based PBSHM technologies beyond those presented in this thesis. Chapter 6 builds directly on this dataset, developing a transfer learning framework capable of mapping continuous damage states across structurally distinct population members.

Chapter 6

Transfer with continuous health states

In this chapter transfer learning techniques are applied to the data collected from the experiment in Chapter 5. The goal is to improve the predictive performance in a data-scarce domain by leveraging data from a data-rich domain. While this has been explored for classification problems with discrete health states, in SHM, transfer learning for regression problems remains under-explored. In Chapter 5 it was shown that the principal components of the FRF data can be used as damage sensitive features, this chapter explores two main avenues for employing PCA within a transfer learning framework:

- **Subspace Alignment:** Constructing separate principal axes for the source and target domains independently, and subsequently learning a mapping function to align the target subspace with the source. This approach assumes that while the domains differ, the underlying features degrade in a similar manner.
- **Shared Feature Space:** Deriving a single set of principal axes that identify damage-sensitive features in both the source and target domains simultaneously. The objective is to create a domain-invariant feature space such that damage labels defined in the source domain can be applied directly to the target domain without further modification.

6.1 Principal Component Analysis

Data collected by SHM systems are often high dimensional, for example, there might be dense sensor arrays generating correlated time-series data. High dimensional data lead to large computational costs and risks over fitting in many regression and classification

tasks. Principal Component Analysis (PCA) is a statistical procedure that converts observations of correlated variables into a set of linearly uncorrelated variables called *Principal Components*. The core concept of PCA is that variance of a variable relates to information, if a variable changes significantly across the dataset, it is providing information about the state of the structure. If a variable rarely changes or fluctuates only slightly around a mean it is likely noise [176].

Mathematically, PCA performs an orthogonal linear transformation that transforms the data to a new coordinate system. The first coordinate (the first principal component) is aligned with the direction of greatest variance in the data. Each succeeding component is aligned with the direction of the greatest remaining variance, under the constraint that it be orthogonal to the preceding components. By retaining only the first k components that capture some majority of the variance (e.g., 95%), the dimensionality is reduced while preserving the dominant structural behaviours and discarding the dimensions dominated by stochastic noise [180].

Consider an input dataset $\mathbf{X}$ consisting of n observations and m variables. The data matrix is defined as:

$$\mathbf{X} = \begin{bmatrix} x_{1,1} & x_{1,2} & \cdots & x_{1,m} \\ x_{2,1} & x_{2,2} & \cdots & x_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n,1} & x_{n,2} & \cdots & x_{n,m} \end{bmatrix} \in \mathbb{R}^{n \times m} \quad (6.1)$$

Before analysis, the data must be mean-centered to ensure the principal components describe the variance rather than the mean. The mean vector $\boldsymbol{\mu}$ is calculated as:

$$\mu_j = \frac{1}{n} \sum_{i=1}^n x_{i,j} \quad \text{for } j = 1, \dots, m \quad (6.2)$$

The mean-centered data matrix $\bar{\mathbf{X}}$ is obtained by subtracting the mean from each column:

$$\bar{x}_{i,j} = x_{i,j} - \mu_j \quad (6.3)$$

PCA relies on the covariance matrix to quantify the correlation between measurements. The covariance matrix $\mathbf{C}$ is calculated as:

$$\mathbf{C} = \frac{1}{n-1} \bar{\mathbf{X}}^T \bar{\mathbf{X}} \in \mathbb{R}^{m \times m} \quad (6.4)$$

The principal components are determined by solving the eigenvalue problem for the covariance matrix $\mathbf{C}$:

$$\mathbf{C}\mathbf{v}_k = \lambda_k \mathbf{v}_k \quad (6.5)$$

Where λ_k is the eigenvalue representing the variance explained by the k -th component. $\mathbf{v}_k$ is the k -th eigenvector, subject to $\|\mathbf{v}_k\| = 1$. The eigenvalues are sorted in descending order to identify the most significant components:

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_m \geq 0 \quad (6.6)$$

To reduce dimensionality, a transformation matrix $\mathbf{W}_k$ is constructed using the first k eigenvectors:

$$\mathbf{W}_k = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k] \in \mathbb{R}^{m \times k} \quad (6.7)$$

The projected data, denoted as $\mathbf{Z}$, is obtained by projecting the mean-centered data onto this subspace:

$$\mathbf{Z} = \bar{\mathbf{X}}\mathbf{W}_k \in \mathbb{R}^{n \times k} \quad (6.8)$$

6.1.1 PCA in PBSHM

While PCA effectively compresses high-dimensional data into informative features, its application in transfer learning introduces significant challenges related to domain shift. PCA is an unsupervised algorithm, meaning that it does not take damage labels into account. Additionally, PCA is a data-dependent transformation, the resulting feature space is intrinsically tied to the covariance structure of the specific dataset used for training. In an SHM context, a source structure and a target structure rarely share identical covariance properties due to differences in material parameters, boundary conditions, or sensor placement. Mathematically, this implies that the principal subspaces of the two domains are misaligned:

$$\text{span}(\mathbf{W}_{\text{source}}) \neq \text{span}(\mathbf{W}_{\text{target}}) \quad (6.9)$$

If the regression model trained on the Source domain is applied directly to the target domain using the source's projection matrix $\mathbf{W}_{\text{source}}$, the resulting feature values will

lack physical consistency. A feature vector $\mathbf{z}$ in the source domain may correspond to a completely different physical state in the target domain due to the rotation of the principal axes.

The literature has attempted to address this, Transfer Component Analysis (TCA) [181] minimises the minimum mean discrepancy (MMD) between the source and target in a reproducing kernel Hilbert space to find a projection matrix that maps source and target data into a shared latent space.

6.2 Naive PCA

Over the course of the experiment detailed in Chapter 5, as the boundary conditions change, the response at a given frequency may change, this can be seen in Figure 5.8. PCA identifies axes in a dataset with the maximum variance, as such, here, one can assume that the identified frequencies are those that are most effected by damage. Such that the data projected onto the principal components can be used as damage sensitive features.

This section examines the consequences of learning independent principal axes in the source and target domains and attempts to learn a mapping between them such that a model trained on the source domain can be used to make predictions on source domain data.

For this example, the source beam is stainless steel and 620mm long, the FRFs of this beam over the experiment can be seen in Figure 6.1. The target beam is 480mm and made of brass. The FRFs of the target beam can be seen in Figure 6.2.

A standard PCA implementation from Scikit-learn [182] can be used to learn separate principal axes for the source and target domains. The weights of the principal components can be seen in Figure 6.3 and 6.4, and are plotted over the source and target FRFs respectively to help visualise which frequencies have been identified by the PCA algorithm.

The projections of the source and target data on to their respective principal subspaces are shown in Figures 6.5 and 6.6.

6.2.1 Fitting a model to the source domain

To make predictions about the damage labels of target data, a model must first be trained on the source domain. In this case study, each domain consists of a dataset with




FIGURE 6.1: The FRF of the source beam, stainless steel - 620mm.




FIGURE 6.2: The FRF of the target beam, brass - 480mm.

200 points, generated by segmenting the 25-minute test from Chapter 5 into 200 distinct sections. This moderate dataset size lends itself well to Gaussian Processes (GPs)




FIGURE 6.3: The first principal component weights learnt on the source and target beams plotted over the respective FRFs.




FIGURE 6.4: The second principal component weights learnt on the source and target beams plotted over the respective FRFs.




FIGURE 6.5: The source data projected onto the principal axes learnt on the source domain.




FIGURE 6.6: The target data projected onto the principal axes learnt on the target domain.

[183], which provide robust estimates for regression tasks alongside rigorous uncertainty quantification. A GP is fitted to the source domain data after it has been projected onto its principal axes. To effectively capture the underlying structure of this data, the GP employs a composite covariance function. Specifically, an anisotropic Matérn kernel ($\nu = 1.5$) scaled by a constant variance kernel, combined with an additive white noise kernel. The Matérn kernel was selected because setting $\nu = 1.5$ models functions that are exactly once differentiable and captures the continuous spatial nature of the damage states. Furthermore, allowing the kernel to be anisotropic enables the model to learn independent length scales for the different variations along each principal component. Figure 6.7 shows the resulting mean prediction, and Figure 6.8 displays the variance prediction. The regression model is not the focus of this research, but is detailed here for clarity. Once a mapping is established from the target domain to the source domain, this model can be leveraged to predict the health states of the target data.




FIGURE 6.7: The mean prediction of the GP trained on the source data.

6.2.2 Finding a mapping

With PCA employed to extract damage-sensitive features and a Gaussian Process used to model damage in the source domain, a domain adaptation step is necessary. A mapping from the target to the source domain is required to allow the source-trained GP to predict damage labels for the target data. This approach is motivated by the




FIGURE 6.8: The covariance prediction of the GP trained on the source data.

observation of repeated geometrical patterns in the projected data across the beam population. Ideally, the mapping procedure would be suitable for mapping any target structure onto any source structure. As an initial methodology, the first two principal components will be used and it will be assumed that the mapping consists of a linear transformation; a translation, learnt through normal condition alignment, and a rotation learned through the log marginal likelihood of the observed target labels under the GP's predictive distribution.

6.2.2.1 Pre-processing

Normal condition alignment (NCA) has been shown to be an important pre-processing step for transfer learning in structural health monitoring [143] and can be a starting point when learning a map from target to source domains. It is usually safe to assume that, at the time of installation, structures are undamaged, and thus data collected early in the life of the structure can be used to align the domains. In this chapter, the same approach is applied to align the data from the different beams. The mean and standard deviation of the undamaged data are used to standardise all the data from that beam. In NCA, the source domain is first standardised via Equation 6.10 to centre the data and give features equal treatment.

$$\mathbf{z}_s^{(j)} = \frac{\mathbf{x}_s^{(j)} - \boldsymbol{\mu}_s}{\boldsymbol{\sigma}_s} \quad (6.10)$$

Where $\boldsymbol{\mu}_s$ is the means of the source data, and $\boldsymbol{\sigma}_s$ is the standard deviation. $\mathbf{x}_s^{(j)}$ is the source data before centring and $\mathbf{z}_s^{(j)}$ the centred data. The normal-condition of the target domain is then aligned with that of the source using Equation 6.11

$$\mathbf{z}_t^{(j)} = \left(\frac{\mathbf{x}_t^{(j)} - \boldsymbol{\mu}_{t,n}}{\boldsymbol{\sigma}_{t,n}} \right) \boldsymbol{\sigma}_{s,n} + \boldsymbol{\mu}_{s,n} \quad (6.11)$$

where $\mathbf{z}_t$ is the aligned target domain data and $\boldsymbol{\mu}_{t,n}, \boldsymbol{\mu}_{s,n}$ and $\boldsymbol{\sigma}_{t,n}, \boldsymbol{\sigma}_{s,n}$ are the means and standard deviations of the normal-condition data for the source and target respectively.

Figure 6.9 shows the initial alignment from NCA of the target data onto the source data.




FIGURE 6.9: The first and second principal components of the source and target. Showing the initial alignment provided by NCA.

6.2.3 Learning the rotation

After a translation is learnt by NCA, a rotation can be learnt by finding the rotation around the normal condition source data that maximises the log-likelihood of the target labels with respect to a model trained on the source domain. The rotation is learnt

about the centre of the normal condition source data. This is finding the rotation of the target domain such that the target data is explained best by the GP trained on the source data. In other words it is finding the rotation that makes the rotated data *agree* with the model trained on the source.

$$\ln f(x) = -\frac{1}{2} [k \ln(2\pi) + \ln |\Sigma| + (x - \mu)^T \Sigma^{-1} (x - \mu)] \quad (6.12)$$

where x is the target data vector, μ is the GP mean prediction of the damage at each of the target data points, Σ is the covariance matrix, and $|\Sigma|$ is the determinant. There is an underlying assumption here that there is some intrinsic similarity between the degradation of the source and target data, and that the mapping of the target data does not warp this similarity.

One can imagine this scenario as two structures being monitored, where the source domain has plentiful damage state information and the target domain has very little, this could be a new structure that currently only has normal condition data. As data arrives throughout the life of this new structure, a better mapping can be learnt such that the GP makes improved predictions about the source data's damage state. Table 6.1 collates the alignment angles learnt by maximising the likelihood with different amounts of target data. The normalised log likelihood column shows the likelihood per datapoint available to the model. The rotations of the target data can be seen in Figure 6.10 and the corresponding log likelihood landscapes can be seen in Figure 6.11 .

TABLE 6.1: Log likelihood and optimal rotation angle across varying data availability in the target domain.

% of Available Data	Learnt Rotation (deg)	Log Likelihood	Normalised Log Likelihood
20%	21.76	63.67	1.592
30%	-126.16	95.35	1.589
40%	-150.28	125.86	1.573
50%	-149.99	156.96	1.570
60%	-152.14	188.08	1.567
70%	-156.73	218.04	1.557
80%	-162.33	248.75	1.555
90%	-165.06	280.07	1.556
100%	-165.20	318.07	1.590

6.2.4 PCA stability

In the Section above, the PCA was calculated with access to 100% of the target data and then those principal axes were used to project progressive amounts of the target FRF data. The main problem with this approach is the stability of the principal axes




FIGURE 6.10: The progression of calculated angle as data becomes available in the target domain.

identified by the PCA algorithm. The axes found when data from the entire experiment is available can be very different to those identified when only the normal-condition data is available. This is expected as the PCA algorithm is looking for variance in the input




FIGURE 6.11: The progression of the log likelihood landscape as data becomes available in the target domain.

features (frequencies). When only normal condition data are available, the majority of the variance is a result of noise or other benign variations and none of it is as a results of damage - as damage has not been observed yet. It's common that as new data becomes available, the axes found through PCA can be mirrored or otherwise change dramatically via large rotations. This effect means that the mapping being calculated at each iteration is different. To address this, it would be useful to identify principal axes that produce damage sensitive scores without having access to the entire data-history of a target structure. The next section examines the consequences of using the principal axes learnt from the source domain and projecting the target data onto it directly, avoiding the need for damage state information in the target domain.

6.3 Shared Feature Space

The projection of the target data onto the principal axes learnt on the source domain data from Section 6.2 can be seen in Figure 6.12. It shows that the identified frequencies do not seem to be correlated with damage in the target domain. This is intuitive as the source and target beams are made of different materials and are different lengths, there is no reason for the beams to behave similarly at a given frequency (the principal axes being linear combinations of the response amplitude at the identified frequencies).




FIGURE 6.12: The target data projected onto the principal axes learnt on the source data.

As can be seen in Equation 5.1, the natural frequencies are proportional to the root of the Young's modulus and inversely proportional to the length of the beam squared. This relationship can be used to non-dimensionalise the frequency axis of the FRFs such that the 1st natural frequencies of each beam are aligned to the same non-dimensional frequency.

In a population of beams, the response at absolute frequency values are difficult to compare. The natural frequencies of a short beam may be very different to the natural frequencies of a different length beam. To a standard machine learning algorithm, these look like very different signals and it may struggle to identify shared features. However, the physics governing their behaviour is identical. The n -th natural frequency is always a specific multiple of the fundamental frequency. By scaling the entire frequency axis

such that the fundamental frequency measured experimentally during normal condition is 100 ($scale = 100/\omega_1$), we convert the data from absolute frequency (Hz) to frequency ratio (dimensionless). Via scaling to this new domain, the first peak for every beam can be aligned at $\Omega = 100$. The geometric and material differences are removed while preserving the overall structure of the frequency responses (the relative spacing and shape of the peaks).

6.4 Non-dimensionalised frequency

The frequency normalisation procedure described in the previous section was applied to align the frequency-domain structural responses. By scaling the frequency axis of each beam relative to its first natural frequency, the resonance peaks are brought into alignment, facilitating direct comparison between structures with differing stiffness or length properties. The non-dimensionalised FRFs of a source (Stainless Steel - 620mm) and target (Brass - 480mm) beam can be seen in Figure 6.13.




FIGURE 6.13: The non-dimensionalised FRFs, with the first natural frequencies normalised to 100. The top graph is the source data, and the bottom graph is the target data.

Again, with the natural frequencies aligned, principal axes can be learnt from the source data, the target data can be projected on to these axes meaning the responses at these frequencies can more easily be compared. Figure 6.14 shows the projected source and

target data. It can be seen that damage sensitive axes have been identified, because these axes are learnt in the source domain, they do not depend on the number of available target data. This helps mitigate some of the stability issues highlighted in Section 6.2. The component weights of the first and second principal components can be seen in Figures 6.15 and 6.16.




FIGURE 6.14: The non-dimensionalised source and target data projected onto the principal axes learnt in the source domain.

Through the methodology described in Section 6.2.2, a mapping can be learnt so that a model trained on the source data can be used to predict damage states in the target domain. Table 6.2 collates the alignment angles learnt by maximising the likelihood with different amounts of target data. The rotations of the target data can be seen in Figure 6.17 and the corresponding log likelihood landscapes can be seen in Figure 6.18. Again, The normalised log likelihood column shows the likelihood per datapoint available to the model.

Standard PCA can be sensitive to noise in the data and has no mathematical formulation to consider populations. The method seeks directions of maximum variance in the training data, meaning that noise-driven variance can hide variance associated with damage. Even with non-dimensionalised inputs, this method does not generalise reliably across all instances of transfer within the beam population documented in Chapter 5.

A natural extension is to move away from a single, population-agnostic projection and instead learn a latent space that explicitly accounts for inter-unit variability. Personalised




FIGURE 6.15: The 1st principal component weights plotted over the scaled source and target FRFs.




FIGURE 6.16: The 2nd principal component weights plotted over the scaled source and target FRFs

TABLE 6.2: Impact of data availability on the learnt rotation angle and log likelihood

% of Available Data	Learnt Rotation (deg)	Log Likelihood	Normalised Log Likelihood
20%	-2.79	64.84	1.621
30%	179.58	96.02	1.600
40%	-181.09	71.37	0.892
50%	-178.83	-70.86	-0.709
60%	-181.82	-378.16	-3.151
70%	-182.49	-850.65	-6.076
80%	175.08	-1122.53	-7.016
90%	175.24	-1294.41	-7.191
100%	175.51	-1478.43	-7.392

PCA addresses this by incorporating population-level information into the dimensionality reduction step, allowing each structure’s projection to reflect both the shared and individual characteristics of the dataset. The following section introduces this approach and examines whether it produces a more robust basis for the subsequent rotation-based alignment.

6.5 Population-based PCA

A population view of PCA is presented in [184]. Personalised PCA (perPCA) takes a hierarchical view of PCA, decoupling data into two distinct subspaces.

- **Global Subspace:** captures patterns common across the population.
- **Local Subspace:** captures patterns specific to a single member of the population.

A key aspect of this work is that these two subspaces are mutually orthogonal, meaning the shared features and unique features do not overlap. This means that variations between members in the population can be captured while identifying the common physics. For the k -th member of the population, the data matrix $X_k \in \mathbb{R}^{d \times n_k}$ is decomposed into shared, unique, and noise components:

$$X_k = UR_k + V_k S_k + E_k \quad (6.13)$$

Where $U \in \mathbb{R}^{d \times r_g}$ is the Global (Shared) Basis. $V_k \in \mathbb{R}^{d \times r_k}$ is the Local (Unique) Basis for member k . $R_k = U^\top X_k$ represents the shared loadings (coefficients). $S_k = V_k^\top X_k$ represents the unique loadings (coefficients) and E_k is the noise matrix. A strict orthogonality is enforced within and between global and local basis:




FIGURE 6.17: The progression of calculated angle as data becomes available in the target domain.

$$U^{\top}U = I_{r_g} \quad (\text{Global basis is orthonormal})$$

$$V_k^{\top}V_k = I_{r_k} \quad (\text{Local basis is orthonormal})$$

$$U^{\top}V_k = 0 \quad (\text{Global and Local subspaces are disjoint})$$




FIGURE 6.18: The progression of the log likelihood landscape as data becomes available in the target domain.

The constraint $U^\top V_k = 0$ is the most critical; it prevents the unique features of member k from “leaking” into the shared global model.

The goal is to reconstruct the data matrices X_k for K different clients while separating the data into a shared global subspace and client-specific (local) subspaces, and maintaining the orthogonality constraints highlighted above. The loss function $\mathcal{L}$ is defined as:

$$\min_{U, \{V_k\}_{k=1}^K} \sum_{k=1}^K \left\| X_k - UU^\top X_k - V_k V_k^\top X_k \right\|_F^2$$

Subject to the Stiefel Manifold constraints:

$$U \in \text{St}(d, r_g), \quad V_k \in \text{St}(d, r_k), \quad \text{and} \quad U^\top V_k = 0$$

The algorithm is based on Stiefel manifold gradient descent that can be proved to converge linearly into the global optimum under mild conditions.

6.5.1 Personalised PCA in PBSHM

Personalised PCA is a natural fit for PBSHM, where methods that separate population level features and local features have already shown to be effective [68, 131, 136, 185].

Again, input features should be normalised before applying the perPCA algorithm. The frequency inputs have been non-dimensionalised as shown in Section 6.4. A population is formed from the FRF data of beams from the experiment in Chapter 5. The perPCA algorithm is trained on the three mild steel beams (source), and the rest of the population (target) will be projected onto the global components. This can be seen in Figure 6.19. The global principal components weights are plotted over beam FRFs in Figures 6.20.

The effect damage has on projected data in Figure 6.19 is similar across the population, meaning the mapping scheme detailed in this chapter can be used to improve damage prediction for individual assets within the population with limited labels. Importantly, the global components learnt on the mild steel beams produce damage sensitive projections for beam materials outside of the training data. Additionally, because perPCA does not require any data from the target domain, the perPCA components are invariant to target data availability. This solves the issue of axes stability from the standard PCA results shown in Section 6.2.

However, this method has not yielded perfect results. The brass 480mm and 620mm beams do not have a structured projection. The exact reason is unclear, however, the brass beams are perhaps the most heterogeneous members of the population, which may cause difficulties during transfer. Additionally, the brass beams are the most rigid of the population making them more strongly coupled with the shaker. Further exploration of the effects of perPCA in engineering structures provides scope for future work.

6.6 Discussions and Conclusions

A critical limitation of conventional data-driven SHM is that supervised regression methods require dense ground truth labels spanning the entire continuous degradation trajectory; obtaining such comprehensive data is often cost-prohibitive or otherwise infeasible. Transfer learning offers a mechanism to mitigate these data requirements in a target domain. However, when target data is sparse unsupervised transfer methods tend to fail, resulting in negative transfer where the predicted trends deviate significantly from




FIGURE 6.19: The FRF data from all the beams in the population projected onto the global principal components learnt on mild steel beams.

the physical reality. This chapter investigated the challenge of transfer learning for continuous variable prediction in Structural Health Monitoring (SHM), specifically within the context of a heterogeneous population of beams undergoing boundary condition degradation.




FIGURE 6.20: 1st and 2nd Global Principal Components plotted over the FRF of a representative beam.

6.6.1 The Instability of Data-Driven Alignment in Sparse Regimes

The initial approach employed standard Principal Component Analysis to extract features from both the source and target domains independently, followed by learning a linear mapping between the subspaces. While this method theoretically aligns the domain distributions, the results highlighted a critical vulnerability: sensitivity to data scarcity.

In the source domain, the principal axes can capture the variance associated with the stiffness degradation. However, in the data-sparse target domain, the PCA calculation was dominated by noise before later stage damage data is available. Consequently, the principal axes learned on the target were not merely rotated versions of the source axes; they represented fundamentally different features. Attempting to map the robust source space onto this unstable target space can result in negative transfer, as the mapping forced the model to align physical trends with statistical artifacts.

6.6.2 Non-dimensionalisation as a stabilisation mechanism

To overcome the instability of the target principal axes, this study introduced a physics-based preprocessing step: non-dimensionalisation of the frequency axis to align the

first natural frequency (ω_1) across the population. This step effectively removed the geometric differences (variance caused by differing beam lengths) from the dataset. By aligning the fundamental resonances, the structural variance was reduced to two primary components: material damping properties and the evolving boundary condition. The results demonstrated that after non-dimensionalisation, the principal axes learned on the source domain could be directly applicable to the target domain.

6.6.3 Decoupling shared and personal behaviours

As a hierarchical extension of PCA, personalized PCA is a natural fit for population modelling for SHM and explicitly decouples the dataset into shared and unique subspaces. The global basis (U) successfully captured the “universal” physics of the loosening clamp—a trajectory of stiffness loss common to all beams regardless of their material or length. Crucially, the orthogonality constraint ($U^\top V_k = 0$) ensured that the unique features of a specific beam did not corrupt the shared model of the degradation process. This allowed for the construction of a generalized degradation model. The global basis facilitates transfer learning through the following procedure:

1. Data Acquisition: Acquire input-output data from the target structure.
2. Spectral Estimation: Compute the Frequency Response Functions (FRFs) from the raw data.
3. Global Projection: Project the target FRFs onto the global basis (U) learnt from a population.
4. Domain Adaptation: Derive a mapping that aligns the projected target data with the source domain, enabling the source-trained model to predict health states for the target structure.

6.7 Conclusion

This chapter has developed and demonstrated a physics-guided transfer learning methodology for continuous structural health monitoring across a structural population, addressing the challenges that arise when the quantity to be transferred is a continuous damage variable rather than a discrete label. The methodology was applied to the beam population dataset introduced in Chapter 5, where the goal was to transfer a regression model trained on a source beam to target beams of different material and geometry. Together, the works of this chapter fulfil Contribution 3: extending transfer learning technologies

for PBSHM to regression tasks. The results establish that, with appropriate preprocessing and population-aware dimensionality reduction, the domain alignment method highlighted in this chapter is capable of transferring continuous health state information across a population.

Chapter 7

Conclusions and Future Work

Enhanced decision support for asset management provides one of the primary motivations for the development and implementation of SHM technologies; however, actual uptake of SHM has been limited by the (lack of) availability of cost-benefit analyses. Framing SHM systems in the context of decision-support allows the value of the technology to be quantified, thereby giving impetus to the broader adoption of SHM system into industry. The work presented in the current thesis focusses on the data scarcity problem in SHM, specifically on how PBSHM techniques can be applied to regression problems.

7.1 Summary and Conclusions

Chapter 4 introduced a risk-based active learning framework for populations of structures, demonstrated on a population of machining tools undergoing health degradation. A Bayesian hierarchical model was used to capture shared degradation behaviour across the population, partially pooling information across assets such that predictions for data-scarce individuals could be informed by observations from the wider population. Inspection decisions were then planned using a risk-based decision process, in which the expected utility of each possible action — inspect, replace, or continue monitoring — was derived analytically and evaluated at each decision point. This formulation allowed the framework to select the most economically justified action in real time, directing inspections toward tools where the consequence of an undetected failure was greatest rather than following a fixed periodic schedule.

The case study demonstrated that this approach yields a significant reduction in monitoring costs, achieved by eliminating unnecessary inspections without meaningfully increasing the rate at which workpieces are damaged — a result that compares favourably

against conventional periodic inspection strategies. This work demonstrated that active learning and transfer learning are a powerful combination of technologies: by directing data collection according to risk rather than information-theoretic measures alone, the framework improved population-level predictions while remaining grounded in the practical economics of the inspection problem. To the author's knowledge, this represents the first application of active learning across a population of structures in the SHM literature, and the first work to address inspection budget allocation at the population level in a formally principled way.

Chapter 5 addressed a significant gap in the PBSHM literature: the lack of experimental data capturing continuous damage progression across a population of structures. Existing SHM datasets largely consist of discrete, artificially induced damage states — typically the addition of masses or removal of material sections — which limits the development and evaluation of regression technologies that must operate across a continuous damage spectrum. To address this, an innovative experimental campaign was conducted in which continuous degradation was recorded across a population of beams varying in both length and material. Damage was induced progressively by reducing the clamping force at the beam boundary, simulating a gradual loss of boundary stiffness rather than an instantaneous, discrete change in structural state. Modal analysis of the resulting data revealed consistent and physically interpretable trends across the population. Natural frequencies decreased progressively with clamping force reduction, confirming that the degradation mechanism produced a measurable and monotonic change in structural dynamic behaviour. These findings are consistent across beams of different lengths and materials, suggesting that the underlying degradation mechanism produces transferable signatures across the population despite geometric and material heterogeneity. Together, these results provide a detailed characterisation of the dynamic behaviour of the beam population and form a robust experimental basis for the development and evaluation of the continuous, population-based transfer learning methodology presented in Chapter 6. Moreover, this dataset will allow further development of population-based prognosis models in the future.

Chapter 6 built upon the dataset introduced in Chapter 5 to develop a transfer learning framework capable of mapping continuous damage states across structurally distinct assets. PCA was shown to identify damage sensitive features from FRF data, but the identified principal axes were unstable in data scarce domains. To address this, a physics-based feature alignment technique was first applied to non-dimensionalise the frequency response functions of the source and target beams, reducing the influence of geometric and material differences on the learned representations. This alignment allows principal axes learned in the source domain to be directly applied to the target domain identifying damage sensitive features suitable for mapping, where the damage

state degradation presents in similar ways in both domains. The transfer mapping itself combined a translation learned via Normal Condition Alignment with a rotation optimised through maximising the log marginal likelihood of the observed target labels under the GP’s predictive distribution, providing a structured and interpretable approach to domain adaptation for regression in heterogeneous structural populations. Finally, a population-based formulation of PCA (personalised PCA) was then used to extract damage-sensitive features suitable for the transfer mapping, and was shown to preserve damage information across variations in beam length and material.

7.2 Limitations, Challenges and Future Works

There are several limitations of the works presented in the current thesis, alongside broader unaddressed challenges associated with data scarcity in SHM that must be overcome in order to achieve robust, cost-optimal operation and maintenance of structural assets.

Budget Allocation and Inspection Planning Across Populations

The risk-based active learning framework presented in Chapter 4 carries several practical limitations that warrant attention, alongside a number of natural extensions that would strengthen its theoretical foundations and broaden its applicability. The derivation of expected utility equations requires careful, problem-specific formulation and a thorough understanding of the underlying decision context. While this is achievable in well-characterised industrial settings, it may present a barrier to adoption in applications where the cost structure or failure consequences are difficult to quantify. The Bayesian hierarchical model, while well-suited to capturing population-level structure, introduces non-trivial computational considerations: partial pooling increases the dimensionality of the posterior that a Monte Carlo sampler must explore, which can restrict prior specification and lead to complex posterior geometries that are difficult to navigate efficiently. In online monitoring scenarios where decisions must be made in real time alongside continuous data acquisition, these costs may be prohibitive without further approximation. Future work should therefore explore scalable inference strategies — such as Gaussian processes with structured priors or amortised variational approaches — that could preserve the risk-based decision structure while making the framework tractable for real-time and large-scale deployment. Perhaps most significantly, a full Value of Information analysis was omitted from the algorithm as presented. A complete VoI calculation — quantifying the benefit of new observations on the predictive capability of

the population model — would provide a more theoretically complete basis for the inspection decisions, but its inclusion would add considerable computational burden. This limitation points toward a broader and more consequential open problem: the current framework selects inspections based on the risk associated with individual assets, but does not formally account for the downstream benefit that a new observation on one structure confers upon predictions for the rest of the population. A unified framework that quantifies this population-level value of information and allocates a fixed inspection budget accordingly would represent a significant theoretical advance, connecting the risk-based active learning approach to the broader VoI literature in a principled way. Incorporating such an analysis within a scalable approximation scheme represents one of the most natural and valuable directions for future work. Related to this, the question of optimal sensor placement across a population has received relatively little attention. Extending active learning principles to the experimental design problem — determining not just when to inspect, but where to instrument — would allow monitoring resources to be deployed more efficiently across a fleet of structures. This is particularly relevant in scenarios where sensor installation is costly and the population exhibits significant structural heterogeneity, such that the most informative sensor locations differ meaningfully between assets.

Experimental Development

The experimental dataset introduced in Chapter 5 represents a meaningful advance over existing discrete-damage SHM benchmarks, but several limitations remain. The population was restricted to beams of varying length and material, all sharing the same rectangular cross-section and a single degradation mechanism — a progressive reduction in clamping force at one boundary. While this produced a well-controlled and physically interpretable dataset, it does not capture the full diversity of structural variation encountered in real asset populations. Notably, the experiment was conducted under controlled laboratory conditions, excluding the environmental and operational variability — such as temperature fluctuations and varying loading conditions — that routinely confound damage detection in field-deployed systems. Additionally, the sensor configuration was fixed across all beams, whereas in practice the number and placement of sensors will differ between population members, introducing a further source of heterogeneity that the dataset does not represent. Future experimental campaigns should consider a broader range of cross-sectional geometries, boundary condition types, and damage mechanisms — including scenarios in which degradation manifests differently across population members, better reflecting the heterogeneity that population-based methods must ultimately contend with. Expanding the range of clamping force regimes

would further allow the sensitivity of the framework to different rates and extents of degradation to be more thoroughly characterised, and would provide a richer basis for evaluating the limits of population-based transfer when the source and target damage trajectories are no longer aligned.

Extensions to the Transfer Learning Mapping

The transfer learning framework developed in Chapter 6 relies on the assumption that damage-sensitive features can be related across source and target domains through a rigid rotation and translation in a shared latent space. This assumption holds well when the degradation mechanism is consistent across the population, producing feature space geometries that are related by location and orientation — as demonstrated on the beam population. However, it represents a meaningful restriction in the general case. If the damage patterns of the source and target are highly disparate — due to fundamentally different failure modes, non-linear damage effects at higher severities, or significant geometric dissimilarity — the rigid mapping may be insufficiently expressive and negative transfer becomes a real risk. Furthermore, the method assumes that the relationship between feature spaces is approximately linear; real structures subjected to severe or multi-site damage may exhibit non-linear changes in their dynamic response that cannot be adequately captured by the current formulation. A practically valuable direction for future work would be the development of a principled criterion for assessing when the rigid mapping assumption is likely to fail — for instance, based on a measure of geometric discrepancy between the source and target feature distributions prior to alignment — allowing the framework to automatically identify cases where a more flexible approach is warranted. Such approaches might include non-linear manifold alignment methods or deep feature alignment strategies, which could accommodate greater structural diversity while ideally retaining some degree of interpretability. Related to this, the current framework does not explicitly account for uncertainty in the learned mapping itself — propagating this uncertainty through to the final predictions would provide a more complete picture of when the transfer can be trusted and when it should be treated with caution. This connects naturally to the probabilistic foundations of the thesis and represents a meaningful avenue for future development.

Broader Challenges

Beyond the specific limitations of each chapter, several broader challenges remain for the field. The works presented here are validated on relatively controlled laboratory populations; scaling these methods to large, heterogeneous in-service populations — where

data quality, sensor configurations, and operational conditions vary unpredictably — remains an open and important problem. Furthermore, while the thesis addresses data scarcity through transfer learning and active learning separately, a unified framework that jointly optimises information transfer and data collection across a population has yet to be fully realised. Developing such a framework, grounded in principled decision theory and scalable to realistic engineering populations, represents one of the most consequential directions for future research in PBSHM.

7.3 Closing remarks

A recurring motivation across the works presented in this thesis is that the field of SHM has historically been constrained by a tendency to frame damage as a discrete, categorical phenomenon amenable to classification. While classification remains a valuable capability, it provides an incomplete picture of structural health — one that is insufficient for the prognostic and decision-support applications that industry increasingly demands. Knowing that a structure is damaged is useful; knowing how severely, how rapidly, and how that severity will evolve is necessary to make informed decisions for O&M. Regression is the capability that bridges this gap, and its development for population-based settings is the central concern of this thesis.

The three contributions presented in Section 1.4 address a different facet of this challenge. The risk-based active learning framework of Chapter 4 demonstrated that regression models for continuously degrading systems can be combined with decision theory to produce inspection strategies that are both data-efficient and economically grounded — showing that the value of regression for decision making and resource allocation. The experimental dataset introduced in Chapter 5 provided the community with a benchmark that reflects the continuous nature of real structural degradation, moving beyond the discrete damage states that have historically limited the development of regression technologies in SHM. The transfer learning framework of Chapter 6 demonstrated that continuous damage states can be mapped across structurally distinct population members, establishing that population-based regression is feasible even in the presence of geometric and material heterogeneity.

Together, these contributions reflect a broader conviction that the maturation of PBSHM as an engineering discipline requires an expansion in the available technologies include state-of-the-art regression tools. Although further work is necessary to translate these methods to the full complexity of in-service structural populations, the results presented here suggest that a regression-focused, population-based approach to SHM

offers a principled and practically consequential path toward the reliable, quantitative monitoring of safety-critical and high-value structural assets.

Appendix A

Chapter 5 - additional material

This appendix contains the FRFs for the population of beams recorded in Chapter 5.




FIGURE A.1: The FRFs for the stainless steel beams.




FIGURE A.2: The FRFs for the mild steel beams.

Appendices/A/.figures/appendix_b.png

FIGURE A.3: The FRFs for the brass beams.




FIGURE A.4: The FRFs for the aluminium beams.

Bibliography

- [1] Impact Report: Powering The Future With Wind. Technical report, Global Wind Energy Council, 2024.
- [2] Renewables 2023. Technical report, International Energy Agency, 2024.
- [3] A. Rytter. Vibrational Based Inspection of Civil Engineering Structures. 1993. Place: Aalborg.
- [4] K. Worden and J. M. Dulieu-Barton. An Overview of Intelligent Fault Detection in Systems and Structures. *Structural Health Monitoring*, 2004.
- [5] C. R. Farrar and K. Worden. *Structural health monitoring: a machine learning perspective*. John Wiley & Sons, 2012.
- [6] O. Zienkiewicz and R. Taylor. *The Finite Element Method for Solid and Structural Mechanics*. Elsevier, 2005.
- [7] A. A. Shabana. *Dynamics of multibody systems*. Cambridge university press, 2020.
- [8] J. D. Anderson. *Computational fluid dynamics: the basics with applications*. New York: McGraw-Hill, 2002.
- [9] O. Abdeljaber, O. Avci, S. Kiranyaz, M. Gabbouj, and D. J. Inman. Real-time vibration-based structural damage detection using one-dimensional convolutional neural networks. *Journal of Sound and Vibration*, 2017.
- [10] Y.-J. Cha, W. Choi, and O. Büyüköztürk. Deep Learning-Based Crack Damage Detection Using Convolutional Neural Networks. *Computer-Aided Civil and Infrastructure Engineering*, 2017.
- [11] O. Janssens, V. Slavkovikj, B. Vervisch, K. Stockman, M. Loccufer, S. Verstockt, R. Van de Walle, and S. Van Hoecke. Convolutional Neural Network Based Fault Detection for Rotating Machinery. *Journal of Sound and Vibration*, 2016.

- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012.
- [13] S. Xu and H. Y. Noh. PhyMDAN: Physics-informed knowledge transfer between buildings for seismic damage diagnosis through adversarial learning. *Mechanical Systems and Signal Processing*, 2021.
- [14] J. Poole, P. Gardner, A. J. Hughes, N. Dervilis, T. Dardeno, R. S. Mills, and K. Worden. Physics-Informed Transfer Learning for Shm Via Feature Selection, 2025.
- [15] E. J. Cross, S. J. Gibson, M. R. Jones, D. J. Pitchforth, S. Zhang, and T. J. Rogers. Physics-Informed Machine Learning for Structural Health Monitoring. In A. Cury, D. Ribeiro, F. Ubertini, and M. D. Todd, editors, *Structural Health Monitoring Based on Data Science Techniques*. Springer International Publishing, Cham, 2022.
- [16] E. J. Cross and T. J. Rogers. Physics-derived covariance functions for machine learning in structural dynamics. *IFAC-PapersOnLine*, 2021.
- [17] E. J. Cross, T. J. Rogers, D. J. Pitchforth, S. J. Gibson, S. Zhang, and M. R. Jones. A spectrum of physics-informed Gaussian processes for regression in engineering. *Data-Centric Engineering*, 2024.
- [18] M. R. Jones, T. J. Rogers, and E. J. Cross. Constraining Gaussian processes for physics-informed acoustic emission mapping. *Mechanical Systems and Signal Processing*, 2023.
- [19] B. Peeters and G. De Roeck. One-year monitoring of the Z24-Bridge: environmental effects versus damage events. *Earthquake Engineering & Structural Dynamics*, 2001.
- [20] J. Wilson, M. D. Champneys, M. Tipuric, R. Mills, D. J. Wagg, and T. J. Rogers. Multiple-input, multiple-output modal testing of a Hawk T1A aircraft: a new full-scale dataset for structural health monitoring. *Structural Health Monitoring*, 2024.
- [21] G. E. Hinton and R. R. Salakhutdinov. Reducing the Dimensionality of Data with Neural Networks. *Science*, 2006.
- [22] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.

- [23] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the Support of a High-Dimensional Distribution. *Neural Computation*, 2001.
- [24] K. Worden, G. Manson, and N. R. J. Fieller. Damage Detection Using Outlier Analysis. *Journal of Sound and Vibration*, 2000.
- [25] A. de Sousa, J. da Silva Coelho, M. R. Machado, and M. Dutkiewicz. Multiclass Supervised Machine Learning Algorithms Applied to Damage and Assessment Using Beam Dynamic Response. *Journal of Vibration Engineering & Technologies*, 2023.
- [26] L. A. Bull, P. A. Gardner, J. Gosliga, T. J. Rogers, N. Dervilis, E. J. Cross, E. Papatheou, A. E. Maguire, C. Campos, and K. Worden. Foundations of population-based SHM, Part I: Homogeneous populations and forms. *Mechanical Systems and Signal Processing*, 2021.
- [27] J. Gosliga, P. A. Gardner, L. A. Bull, N. Dervilis, and K. Worden. Foundations of Population-based SHM, Part II: Heterogeneous populations—Graphs, networks, and communities. *Mechanical systems and signal processing*, 2021.
- [28] P. Gardner, L. A. Bull, J. Gosliga, N. Dervilis, and K. Worden. Foundations of population-based SHM, Part III: Heterogeneous populations—Mapping and transfer. *Mechanical Systems and Signal Processing*, 2021.
- [29] G. Tsialiamanis, C. Mylonas, E. Chatzi, N. Dervilis, D. J. Wagg, and K. Worden. Foundations of population-based SHM, Part IV: The geometry of spaces of structures and their feature spaces. *Mechanical Systems and Signal Processing*, 2021.
- [30] S. J. Pan. Transfer Learning. In *Data Classification*. Chapman and Hall/CRC, 2014.
- [31] L. Bull, K. Worden, G. Manson, and N. Dervilis. Active learning for semi-supervised structural health monitoring. *Journal of Sound and Vibration*, 2018.
- [32] X. Zhu, P. Zhang, X. Lin, and Y. Shi. Active learning from data streams. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*. IEEE, 2007.
- [33] R. E. Moore. *Interval analysis*. Prentice-Hall, 1996.
- [34] R. L. Muhanna and R. L. Mullen. Uncertainty in Mechanics Problems—Interval-Based Approach. *Journal of Engineering Mechanics*, 2001.

- [35] L. A. Zadeh. Fuzzy sets. *Information and Control*, 1965.
- [36] R. Ganguli. A Fuzzy Logic System for Ground Based Structural Health Monitoring of a Helicopter Rotor Using Modal Data. *Journal of Intelligent Material Systems and Structures*, 2001.
- [37] A. N. Kolmogorov. *Foundations of the theory of probability: Second English Edition*. 2018.
- [38] K. P. Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [39] A. J. Hughes. *On risk-based decision-making for structural health monitoring*. phd, University of Sheffield, 2022.
- [40] C. R. Farrar, S. W. Doebling, and D. A. Nix. Vibration-based structural damage identification. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 2001.
- [41] C. Wickramarachchi. *Automated testing of advanced cutting tool materials*. PhD thesis, University of Sheffield, 2019.
- [42] E. A. Johnson, H. F. Lam, L. S. Katafygiotis, and J. L. Beck. Phase I IASC-ASCE Structural Health Monitoring Benchmark Problem Using Simulated Data. *Journal of Engineering Mechanics*, 2004.
- [43] Case Western Reserve University Bearing Data Center Website.
- [44] L. H. Yam, Y. J. Yan, and J. S. Jiang. Vibration-based damage detection for composite structures using wavelet transform and neural network identification. *Composite Structures*, 2003.
- [45] 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*, 2006.
- [46] A. Mita and H. Hagiwara. Damage diagnosis of a building structure using support vector machine and modal frequency patterns. In *Smart Structures and Materials 2003: Smart Systems and Nondestructive Evaluation for Civil Infrastructures*, volume 5057. SPIE, 2003.
- [47] F. Schwenker and E. Trentin. Pattern classification and clustering: A review of partially supervised learning approaches. *Pattern Recognition Letters*, 2014.
- [48] L. A. Bull, K. Worden, T. J. Rogers, C. Wickramarachchi, E. J. Cross, T. McLeay, W. Leahy, and N. Dervilis. A probabilistic framework for online structural health

- monitoring: active learning from machining data streams. In *Journal of Physics: Conference Series*, volume 1264. IOP Publishing, 2019.
- [49] 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*, 2022.
- [50] H. S. Seung, M. Oppor, and H. Sompolinsky. Query by committee. In *Proceedings of the fifth annual workshop on Computational learning theory, COLT '92*, New York, NY, USA, 1992. Association for Computing Machinery.
- [51] M. E. Ramirez-Loaiza, M. Sharma, G. Kumar, and M. Bilgic. Active learning: an empirical study of common baselines. *Data Mining and Knowledge Discovery*, 2017.
- [52] B. Settles. Active Learning Literature Survey. Technical Report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- [53] L. A. Bull, P. Gardner, T. J. Rogers, E. J. Cross, N. Dervilis, and K. Worden. Probabilistic Inference for Structural Health Monitoring: New Modes of Learning from Data. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 2021.
- [54] D. Barber. *Bayesian reasoning and machine learning*. cambridge university press, 2012.
- [55] NSS. Power of Bayesian Statistics & Probability | Data Analysis (Updated 2025), 2016.
- [56] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC, New York, 1995.
- [57] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 1953.
- [58] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 1970.
- [59] M. D. Hoffman and A. Gelman. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.*, 2014.
- [60] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An Introduction to Variational Methods for Graphical Models. *Machine Learning*, 1999.

- [61] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 2017.
- [62] L. Tierney and J. B. Kadane. Accurate Approximations for Posterior Moments and Marginal Densities. *Journal of the American Statistical Association*, 1986.
- [63] M. A. Beaumont. Approximate Bayesian Computation. *Annual Review of Statistics and Its Application*, 2019.
- [64] J. Skilling. Nested sampling for general Bayesian computation. *Bayesian Analysis*, 2006.
- [65] R. van de Schoot, S. Depaoli, R. King, B. Kramer, K. Märtens, M. G. Tadesse, M. Vannucci, A. Gelman, D. Veen, J. Willemsen, and C. Yau. Bayesian statistics and modelling. *Nature Reviews Methods Primers*, 2021.
- [66] R. McElreath. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. Chapman and Hall/CRC, New York, 2018.
- [67] A. Gelman. Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). 2006.
- [68] L. A. Bull, D. Di Francesco, M. Dhada, O. Steinert, T. Lindgren, A. K. Parlikad, A. B. Duncan, and M. Girolami. Hierarchical Bayesian modeling for knowledge transfer across engineering fleets via multitask learning. *Computer-Aided Civil and Infrastructure Engineering*, 2023.
- [69] K. Worden, L. A. Bull, P. Gardner, J. Gosliga, T. J. Rogers, E. J. Cross, E. Papatheou, W. Lin, and N. Dervilis. A brief introduction to recent developments in population-based structural health monitoring. *Frontiers in Built Environment*, 2020.
- [70] D. S. Brennan, J. Gosliga, E. J. Cross, and K. Worden. Foundations of population-based SHM, Part V: Network, framework and database. *Mechanical Systems and Signal Processing*, 2025.
- [71] P. Gardner, X. Liu, and K. Worden. On the application of domain adaptation in structural health monitoring. *Mechanical Systems and Signal Processing*, 2020.
- [72] R. J. Allemang. The modal assurance criterion—twenty years of use and abuse. *Sound and Vibration*, 2003.
- [73] S. J. Pan and Q. Yang. A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 2010.

- [74] W. Dai, Q. Yang, G.-R. Xue, and Y. Yu. Self-taught clustering. In *Proceedings of the 25th international conference on Machine learning, ICML '08*, New York, NY, USA, 2008. Association for Computing Machinery.
- [75] Z. Wang, Y. Song, and C. Zhang. Transferred Dimensionality Reduction. In W. Daelemans, B. Goethals, and K. Morik, editors, *Machine Learning and Knowledge Discovery in Databases*, Berlin, Heidelberg, 2008. Springer.
- [76] R. Caruana. Multitask Learning. *Machine Learning*, 1997.
- [77] L. A. Bull, P. A. Gardner, J. Gosliga, N. Dervilis, E. Papatheou, A. E. Maguire, C. Campos, T. J. Rogers, E. J. Cross, and K. Worden. Towards population-based structural health monitoring, Part I: Homogeneous populations and forms. In *Model Validation and Uncertainty Quantification, Volume 3: Proceedings of the 38th IMAC, A Conference and Exposition on Structural Dynamics 2020*. Springer, 2020.
- [78] S. Targ, D. Almeida, and K. Lyman. Resnet in Resnet: Generalizing Residual Architectures, 2016. arXiv:1603.08029 [cs].
- [79] BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding - ACL Anthology.
- [80] A. Radford and K. Narasimhan. Improving Language Understanding by Generative Pre-Training. 2018.
- [81] S. Ruder. An Overview of Multi-Task Learning in Deep Neural Networks, 2017.
- [82] Domain randomization for transferring deep neural networks from simulation to the real world | IEEE Conference Publication | IEEE Xplore.
- [83] J. Liu, Q. Zhang, X. Li, G. Li, Z. Liu, Y. Xie, K. Li, and B. Liu. Transfer learning-based strategies for fault diagnosis in building energy systems. *Energy and Buildings*, 2021.
- [84] Y. Lei, B. Yang, X. Jiang, F. Jia, N. Li, and A. K. Nandi. Applications of machine learning to machine fault diagnosis: A review and roadmap. *Mechanical Systems and Signal Processing*, 2020.
- [85] P. Wang, J. Lu, B. Zhang, and Z. Tang. A review on transfer learning for brain-computer interface classification. In *2015 5th International Conference on Information Science and Technology (ICIST)*, 2015.

- [86] G. Ma, S. Xu, T. Yang, Z. Du, L. Zhu, H. Ding, and Y. Yuan. A Transfer Learning-Based Method for Personalized State of Health Estimation of Lithium-Ion Batteries. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.
- [87] C. Zhang, L. Tu, Z. Yang, B. Du, Z. Zhou, J. Wu, and L. Chen. A CMMOG-based lithium-battery SOH estimation method using multi-task learning framework. *Journal of Energy Storage*, 2025.
- [88] W. Li, H. Zhang, B. van Vlijmen, P. Dechent, and D. U. Sauer. Forecasting battery capacity and power degradation with multi-task learning. *Energy Storage Materials*, 2022.
- [89] Y. Che, Y. Zheng, Y. Wu, X. Lin, J. Li, X. Hu, and R. Teodorescu. Battery States Monitoring for Electric Vehicles Based on Transferred Multi-Task Learning. *IEEE Transactions on Vehicular Technology*, 2023.
- [90] A. S. Qureshi, A. Khan, A. Zameer, and A. Usman. Wind power prediction using deep neural network based meta regression and transfer learning. *Applied Soft Computing*, 2017.
- [91] C. Song, H. Yang, J. Cai, P. Yang, H. Bao, K. Xu, and X.-B. Meng. Multi-energy load forecasting via hierarchical multi-task learning and spatiotemporal attention. *Applied Energy*, 2024.
- [92] D. Jha, L. Ward, A. Paul, W.-k. Liao, A. Choudhary, C. Wolverton, and A. Agrawal. ElemNet: Deep Learning the Chemistry of Materials From Only Elemental Composition. *Scientific Reports*, 2018.
- [93] T. A. Dardeno, L. A. Bull, N. Dervilis, and K. Worden. When does a bridge become an aeroplane?, 2024. arXiv:2411.18406 [cs].
- [94] P. Gardner, L. A. Bull, N. Dervilis, and K. Worden. Overcoming the problem of repair in structural health monitoring: Metric-informed transfer learning. *Journal of Sound and Vibration*, 2021.
- [95] J. Poole, K. Worden, N. Dervilis, V. Giglioni, R. S. Mills, P. Gardner, and A. J. Hughes. Active Transfer Learning for Structural Health Monitoring, 2025.
- [96] V. Giglioni, J. Poole, R. Mills, I. Venanzi, F. Ubertini, and K. Worden. On the application of domain adaptation for knowledge transfer and damage detection across bridge spans: an experimental case study. *Procedia Structural Integrity*, 2024.

- [97] H.-P. Wan, G.-S. Dong, Y. Luo, and Y.-Q. Ni. An improved complex multi-task Bayesian compressive sensing approach for compression and reconstruction of SHM data. *Mechanical Systems and Signal Processing*, 2022.
- [98] Y. Li, T. Bao, Z. Chen, Z. Gao, X. Shu, and K. Zhang. A missing sensor measurement data reconstruction framework powered by multi-task Gaussian process regression for dam structural health monitoring systems. *Measurement*, 2021.
- [99] D. D. Lewis and W. A. Gale. A Sequential Algorithm for Training Text Classifiers. In B. W. Croft and C. J. van Rijsbergen, editors, *SIGIR '94*, London, 1994. Springer.
- [100] D. Lewis and J. Catlett. Heterogeneous Uncertainty Sampling for Supervised Learning. In *Machine Learning Proceedings 1994*. Morgan Kaufmann, 1994.
- [101] D. Cohn, L. Atlas, and R. Ladner. Improving generalization with active learning. *Machine Learning*, 1994.
- [102] B. Settles, M. Craven, and S. Ray. Multiple-Instance Active Learning. In *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- [103] D. J. C. MacKay. Information-Based Objective Functions for Active Data Selection. *Neural Computation*, 1992.
- [104] D. A. Cohn. Neural Network Exploration Using Optimal Experiment Design. *Neural Networks*, 1996.
- [105] Y. Guo and D. Schuurmans. Discriminative Batch Mode Active Learning. In *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- [106] A. Kirsch, J. van Amersfoort, and Y. Gal. BatchBALD: Efficient and Diverse Batch Acquisition for Deep Bayesian Active Learning. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [107] W.-N. Hsu and H.-T. Lin. Active Learning by Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2015.
- [108] S. Budd, E. C. Robinson, and B. Kainz. A survey on active learning and human-in-the-loop deep learning for medical image analysis. *Medical Image Analysis*, 2021.
- [109] C. Feng, M.-Y. Liu, C.-C. Kao, and T.-Y. Lee. Deep Active Learning for Civil Infrastructure Defect Detection and Classification. 2017.

- [110] J. Smailović, M. Grčar, N. Lavrač, and M. Žnidaršič. Stream-based active learning for sentiment analysis in the financial domain. *Information Sciences*, 2014.
- [111] D. Sculley. Online active learning methods for fast label-efficient spam filtering. 2007.
- [112] Y. Gal, R. Islam, and Z. Ghahramani. Deep Bayesian Active Learning with Image Data. In *Proceedings of the 34th International Conference on Machine Learning*. PMLR, 2017.
- [113] L. Bassman Oftelie, P. Rajak, R. K. Kalia, A. Nakano, F. Sha, J. Sun, D. J. Singh, M. Aykol, P. Huck, K. Persson, and P. Vashishta. Active learning for accelerated design of layered materials. *npj Computational Materials*, 2018.
- [114] T. Lookman, P. V. Balachandran, D. Xue, and R. Yuan. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. *npj Computational Materials*, 2019.
- [115] A. T. Taylor, T. A. Berrueta, and T. D. Murphey. Active learning in robotics: A review of control principles. *Mechatronics*, 2021.
- [116] D. A. Cohn, Z. Ghahramani, and M. I. Jordan. Active learning with statistical models. *Journal of artificial intelligence research*, 1996.
- [117] V. V. Fedorov. *Theory Of Optimal Experiments*. Elsevier, 2013.
- [118] A. Krause, A. Singh, and C. Guestrin. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *J. Mach. Learn. Res.*, 2008.
- [119] H. Mania, M. I. Jordan, and B. Recht. Active Learning for Nonlinear System Identification with Guarantees, 2020.
- [120] C. Papadimitriou. Optimal sensor placement methodology for parametric identification of structural systems. *Journal of Sound and Vibration*, 2004.
- [121] L. A. Bull, T. J. Rogers, C. Wickramarachchi, E. J. Cross, K. Worden, and N. Dervilis. Probabilistic active learning: An online framework for structural health monitoring. *Mechanical Systems and Signal Processing*, 2019. ISBN: 0888-3270.
- [122] A. Keane, A. Forrester, and A. Sobester. *Engineering Design via Surrogate Modelling: A Practical Guide*. American Institute of Aeronautics and Astronautics, Inc., Washington, DC, 2008.

- [123] S. Seo, M. Wallat, T. Graepel, and K. Obermayer. Gaussian process regression: active data selection and test point rejection. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, volume 3, 2000.
- [124] E. Snelson and Z. Ghahramani. Sparse Gaussian Processes using Pseudo-inputs. In *Advances in Neural Information Processing Systems*, volume 18. MIT Press, 2005.
- [125] Y. Yang and M. Loog. A variance maximization criterion for active learning. *Pattern Recognition*, 2018.
- [126] C. Ling, Z. Lu, and X. Zhu. Efficient methods by active learning Kriging coupled with variance reduction based sampling methods for time-dependent failure probability. *Reliability Engineering & System Safety*, 2019.
- [127] T. Rainforth, A. Foster, D. R. Ivanova, and F. B. Smith. Modern Bayesian Experimental Design. *Statistical Science*, 2024.
- [128] Q. Wei, Y. Chen, M. Salimi, J. C. Denny, Q. Mei, T. A. Lasko, Q. Chen, S. Wu, A. Franklin, T. Cohen, and H. Xu. Cost-aware active learning for named entity recognition in clinical text. *Journal of the American Medical Informatics Association*, 2019.
- [129] Y. Huang and J. Beck. Hierarchical sparse bayesian learning for structural health monitoring with incomplete modal data. *International Journal for Uncertainty Quantification*, 2015.
- [130] S. M. Smith, A. J. Hughes, T. A. Dardeno, L. A. Bull, N. Dervilis, and K. Worden. Anomaly Detection in Offshore Wind Turbine Structures using Hierarchical Bayesian Modelling. *arXiv preprint arXiv:2402.19295*, 2024.
- [131] T. A. Dardeno, L. A. Bull, N. Dervilis, and K. Worden. A population form via hierarchical Bayesian modelling of the FRF. In *Society for Experimental Mechanics Annual Conference and Exposition*. Springer, 2023.
- [132] G. Y. He, Y. X. Zhao, and C. L. Yan. Parameter estimation in multiaxial fatigue short crack growth model using hierarchical Bayesian linear regression. *Fatigue & Fracture of Engineering Materials & Structures*, 2023.
- [133] Y. Huang, J. Beck, and H. Li. Multitask sparse Bayesian learning with applications in structural health monitoring. *Computer-Aided Civil and Infrastructure Engineering*, 2019.

- [134] M. Dhada, L. Bull, M. Girolami, and A. Parlikad. Collaborative prognosis using a Weibull statistical hierarchical model. *Reliability Engineering & System Safety*, 2025.
- [135] M. Dhada, M. Girolami, and A. K. Parlikad. Anomaly detection in a fleet of industrial assets with hierarchical statistical modeling. *Data-Centric Engineering*, 2020.
- [136] T. A. Dardeno, K. Worden, N. Dervilis, R. S. Mills, and L. A. Bull. On the hierarchical Bayesian modelling of frequency response functions. *Mechanical Systems and Signal Processing*, 2024. ISBN: 0888-3270.
- [137] S. M. Brealy, L. A. Bull, P. Beltrando, A. Sommer, N. Dervilis, and K. Worden. Inter-turbine Modelling of Wind-Farm Power using Multi-task Learning. *Mechanical Systems and Signal Processing*, 2025.
- [138] G. Tsialiamanis, C. Sbarufatti, N. Dervilis, and K. Worden. On a meta-learning population-based approach to damage prognosis. *Mechanical Systems and Signal Processing*, 2024.
- [139] O. Sedehi, A. M. Kosikova, C. Papadimitriou, and L. S. Katafygiotis. On the integration of Physics-Based Machine Learning with hierarchical Bayesian modeling techniques. *Mechanical Systems and Signal Processing*, 2024.
- [140] D. Di Francesco, M. Chryssanthopoulos, M. H. Faber, and U. Bharadwaj. Decision-theoretic inspection planning using imperfect and incomplete data. *Data-Centric Engineering*, 2021.
- [141] N. Papadimas and T. Dodwell. A hierarchical Bayesian approach for calibration of stochastic material models. *Data-Centric Engineering*, 2021.
- [142] P. Gardner, L. A. Bull, N. Dervilis, and K. Worden. On the application of kernelised Bayesian transfer learning to population-based structural health monitoring. *Mechanical Systems and Signal Processing*, 2022.
- [143] J. Poole, P. Gardner, N. Dervilis, L. Bull, and K. Worden. On statistic alignment for domain adaptation in structural health monitoring. *Structural Health Monitoring*, 2023.
- [144] C. T. Wickramarachchi. Statistical alignment in transfer learning to address the repair problem: An experimental case study. 2023.
- [145] N. Rezazadeh, A. D. Luca, D. Perfetto, G. Lamanna, F. Annaz, M. D. Oliveira, N. Rezazadeh, A. D. Luca, D. Perfetto, G. Lamanna, F. Annaz, and M. D. Oliveira.

- Domain-Adaptive Graph Attention Semi-Supervised Network for Temperature-Resilient SHM of Composite Plates. *Sensors*, 2025.
- [146] G. Tsialiamanis, C. Mylonas, E. N. Chatzi, D. J. Wagg, N. Dervilis, and K. Worden. On an Application of Graph Neural Networks in Population-Based SHM. In R. Madarshahian and F. Hemez, editors, *Data Science in Engineering, Volume 9*, Cham, 2022. Springer International Publishing.
- [147] J. Gosliga, P. Gardner, L. A. Bull, N. Dervilis, and K. Worden. Towards population-based structural health monitoring, Part III: Graphs, networks and communities. In *Topics in Modal Analysis & Testing, Volume 8: Proceedings of the 38th IMAC, A Conference and Exposition on Structural Dynamics 2020*. Springer, 2021.
- [148] X. Sun and P. Chang. Geometry-Consistent Alignment for Time-Lapse Crack Monitoring in Field-Based SHM, 2025.
- [149] V. Giglioni, I. Venanzi, and F. Ubertini. An Adversarial-Based Method for Knowledge Transfer Across Bridges. *International Journal of Bridge Engineering, Management and Research*, 2025.
- [150] X. Chen, S. Wang, J. Wang, and M. Long. Representation Subspace Distance for Domain Adaptation Regression. In *Proceedings of the 38th International Conference on Machine Learning*. PMLR, 2021.
- [151] B. Zhang, H. Dong, H. A. A. M. Qaid, Y. Wang, B. Zhang, H. Dong, H. A. A. M. Qaid, and Y. Wang. Deep Domain Adaptation with Correlation Alignment and Supervised Contrastive Learning for Intelligent Fault Diagnosis in Bearings and Gears of Rotating Machinery. *Actuators*, 2024.
- [152] H. Cheng, X. Kong, Q. Wang, H. Ma, S. Yang, and G. Chen. Deep transfer learning based on dynamic domain adaptation for remaining useful life prediction under different working conditions. *Journal of Intelligent Manufacturing*, 2023.
- [153] I. Nejjar, G. Frusque, F. Forest, and O. Fink. Uncertainty-guided alignment for unsupervised domain adaptation in regression. *Reliability Engineering & System Safety*, 2026.
- [154] J. Li, X. Li, and D. He. Domain Adaptation Remaining Useful Life Prediction Method Based on AdaBN-DCNN. In *2019 Prognostics and System Health Management Conference (PHM-Qingdao)*, 2019.
- [155] A. J. Hughes, L. A. Bull, P. Gardner, N. Dervilis, and K. Worden. On robust risk-based active-learning algorithms for enhanced decision support. *Mechanical Systems and Signal Processing*, 2022.

- [156] A. J. Hughes, R. J. Barthorpe, N. Dervilis, C. R. Farrar, and K. Worden. A probabilistic risk-based decision framework for structural health monitoring. *Mechanical Systems and Signal Processing*, 2021.
- [157] C. Papadimitriou, J. L. Beck, and S.-K. Au. Entropy-Based Optimal Sensor Location for Structural Model Updating. *Journal of Vibration and Control*, 2000.
- [158] M. Tod, F. Mentré, Y. Merlé, and A. Mallet. Robust Optimal Design for the Estimation of Hyperparameters in Population Pharmacokinetics. *Journal of Pharmacokinetics and Biopharmaceutics*, 1998.
- [159] F. MENTRÉ, A. MALLET, and D. BACCAR. Optimal design in random-effects regression models. *Biometrika*, 1997.
- [160] M. de Angelis, E. Patelli, and M. Beer. Advanced Line Sampling for efficient robust reliability analysis. *Structural Safety*, 2015.
- [161] J. Behrendorf, T.-E. Regenhardt, M. Broggi, and M. Beer. Numerically efficient computation of the survival signature for the reliability analysis of large networks. *Reliability Engineering & System Safety*, 2021.
- [162] B. Echard, N. Gayton, and M. Lemaire. AK-MCS: An active learning reliability method combining Kriging and Monte Carlo Simulation. *Structural Safety*, 2011.
- [163] B. Sudret. Meta-models for structural reliability and uncertainty quantification, 2012. arXiv:1203.2062 [stat].
- [164] A. Gelman and J. Hill. *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press, 2006.
- [165] I. G. Kreft and J. De Leeuw. *Introducing multilevel modeling*. Sage, 1998.
- [166] O. Morgenstern, J. Von Neumann, H. W. Kuhn, and A. Rubinstein. *Theory of games and economic behavior*. J. Wiley and Sons, 1964.
- [167] D. Phan, N. Pradhan, and M. Jankowiak. Composable effects for flexible and accelerated probabilistic programming in NumPyro. *arXiv preprint arXiv:1912.11554*, 2019.
- [168] A. B. Sharkawy, M. A. El-Sharief, and M.-E. S. Soliman. Surface roughness prediction in end milling process using intelligent systems. *International journal of machine learning and cybernetics*, 2014. ISBN: 1868-8071.
- [169] Fusion 360, 2025.

- [170] R. D. Blevins. *Formulas for dynamics, acoustics and vibration*. John Wiley & Sons, 2015.
- [171] J. Poole, P. Gardner, N. Dervilis, L. Bull, and K. Worden. On statistic alignment for domain adaptation in structural health monitoring. *Structural Health Monitoring*, 2023.
- [172] J. S. Bendat and A. G. Piersol. *Random Data: Analysis and Measurement Procedures*. John Wiley & Sons, Inc., USA, 2nd edition, 1990.
- [173] B. Peeters, H. Van der Auweraer, P. Guillaume, and J. Leuridan. The PolyMAX Frequency-Domain Method: A New Standard for Modal Parameter Estimation? *Shock and Vibration*, 2004. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1155/2004/523692>.
- [174] D. J. Ewins. *Modal testing: theory, practice, and application*. Number 10 in Mechanical engineering research studies. Research Studies Press, Baldock, 2nd ed. edition, 2000.
- [175] K. P. F.R.S. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 1901.
- [176] H. Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 1933.
- [177] P. Y. Siow, Z. C. Ong, S. Y. Khoo, and K.-S. Lim. Damage Sensitive PCA-FRF Feature in Unsupervised Machine Learning for Damage Detection of Plate-Like Structures. *International Journal of Structural Stability and Dynamics*, 2021.
- [178] F. Khoshnoudian and S. Talaei. A New Damage Index Using FRF Data, 2D-PCA Method and Pattern Recognition Techniques. *International Journal of Structural Stability and Dynamics*, 2017.
- [179] Y. Q. Ni, X. T. Zhou, and J. M. Ko. Experimental investigation of seismic damage identification using PCA-compressed frequency response functions and neural networks. *Journal of Sound and Vibration*, 2006.
- [180] I. Jolliffe. Principal Component Analysis. In *Wiley StatsRef: Statistics Reference Online*. John Wiley & Sons, Ltd, 2014. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118445112.stat06472>.
- [181] S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang. Domain Adaptation via Transfer Component Analysis. *IEEE Transactions on Neural Networks*, 2011.

- [182] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 2011.
- [183] C. Rasmussen and C. Williams. *Gaussian processes for machine learning*.
- [184] N. Shi and R. A. Kontar. Personalized PCA: Decoupling Shared and Unique Features, 2024.
- [185] N. Papadimas and T. Dodwell. A hierarchical Bayesian approach for calibration of stochastic material models. *Data-Centric Engineering*, 2021.