Data-driven Virtual Monitoring and Life-cycle Management of Offshore Wind Support Structures

Nandar Hlaing

Supervisor: Prof. Philippe Rigo

Co-supervisors: Dr. Pablo G. Morato, Prof. Christof Devriendt

Applied Sciences
University of Liège

This dissertation is submitted for the degree of
Doctor of Philosophy in Engineering Sciences

April 2024

PhD Committee

Vincent Denoël (chairman)
Professor
University of Liège, Belgium

Philippe Rigo (supervisor)
Professor
University of Liège, Belgium

Christof Devriendt (co-supervisor)
Professor
Vrije Universiteit Brussel, Belgium

Maarten Arnst
Professor
University of Liège, Belgium

John D. Sørensen
Professor
Aalborg University, Denmark

Eleni Chatzi
Professor
ETH Zurich, Switzerland

Charalampos P. Andriotis
Professor
Delft University of Technology, The Netherlands

I dedicate this thesis to my mother, who has raised me up to become who I am today.

Acknowledgements

Throughout this PhD journey, I have received the support and assistance from many individuals. I would first like to thank my supervisor, Prof. Philippe Rigo, and co-supervisors, Dr. Pablo G. Morato (Delft University of Technology) and Prof. Christof Devriendt (Vrije Universiteit Brussel) for their guidance and continuous support along this research journey.

I gratefully recognize the knowledge transfer from Prof. John D. Sørensen (Aalborg University), Prof. Daniel Straub (Technische Universität München), Prof. Konstantinos G. Papakonstantinou (The Pennsylvania State University), Prof. Eleni Chatzi (ETH Zurich), Prof. Charalampos P. Andriotis (Delft University of Technology), Prof. Maarten Arnst (University of Liège) and Prof. Gilles Louppe (University of Liège). Their inspirational talks and lectures have notably increased my motivation in this field of research.

I would like to acknowledge the coordinators and fellow researchers of PhairywinD project for inspiring me with their research and letting me grow as a team player. The colleagues at ANAST should also take the credit for creating a vibrant atmosphere at the office.

More personally, I would like to express my deep appreciation to my family for respecting and supporting my choices during past years. I also appreciate *my other half* for the persistent encouragement, and especially for being with me in all ups and downs.

Finally, I would like to acknowledge the financial support granted by the Belgian Energy Transition Fund (FPS Economy) through PhairywinD (https://www.phairywind.be) project.

Abstract

With the rapid advancement of offshore wind energy, efficiently managing inspection and maintenance (I&M) of wind turbine support structures has become increasingly important. Various deterioration mechanisms in the harsh marine environment accelerate the structural degradation, inducing a risk of failure which might result in substantial economic losses. Estimation of such deterioration processes involves uncertainties which often hurdle decision-making in life-cycle management planning. Collecting additional data, e.g., through structural health monitoring, can reduce uncertainties in the estimation of deterioration mechanisms, enabling more rational and informed maintenance decisions. However, collecting continuous information through monitoring systems also incurs significant sensor installation and maintenance costs. Addressing these concerns, this thesis is dedicated towards the development of a life-cycle management framework for offshore wind structures by leveraging digital twin technology with the objective of optimally allocating inspection, monitoring, and maintenance actions. The life-cycle management planning is formally formulated as a decentralized partially observable Markov decision process (POMDP), which is a principled framework for decision-making under uncertainty. In this work, maintenance decisions are informed not only by inspection and monitoring data, but also by a probabilistic digital twin. Particularly, the concept of probabilistic digital twins for virtual monitoring is presented relying on Bayesian neural networks and state-of-the-art learning algorithms, and implemented in both numerical and real-world case studies. Featuring high-dimensional state, action, and observation spaces, the formulated POMDP is solved via multi-agent reinforcement learning (MARL) algorithms, advising decisions on when and where to inspect, monitor, or maintain. The outcomes of this research not only showcase the potential of digital twins as virtual sensors, but also quantify their added benefit in life-cycle management planning. While the primary application of this research is in the context of offshore wind turbine support structures, the insights and developed methodologies can be adapted to a wide spectrum of engineering systems or infrastructures, marking a significant shift towards data-driven life-cycle management.

Table of contents

List of figures

List of tables

Chapter 1

Introduction

1.1 Rationale and motivation of the research

Structures and infrastructures are subject to various gradual deterioration mechanisms which can significantly degrade the performance or even trigger failure. For instance, offshore wind substructures experience fatigue deterioration, as a result of the combined action of cyclic wind and wave loads. Along with fatigue, offshore wind turbines are also subject to corrosion deterioration, leading not only to a thickness reduction of structural members, but also accelerating the propagation of existing fatigue cracks. To effectively control the risk of structural failure and mitigate its consequences, it becomes crucial to strategically allocate maintenance actions throughout the lifetime. However, estimating the failure risk of engineering structures involves significant uncertainties, such as aleatory uncertainties arising from stochastic nature of loads and dynamic behaviour of structures themselves, and epistemic uncertainties of deployed analytical/numerical models, therefore hurdling decision-making.

In order to reduce uncertainties, information can be collected through inspections and/or monitoring and used to update the knowledge of structural deterioration. Whereas non-destructive inspection methods previously stood as a primary source of information, the concept of structural health monitoring (SHM) has recently gained importance along with the rapid advancement of sensor technologies. For example, SCADA systems, motion sensors and fiber-optic load cells are increasingly installed in offshore wind farms nowadays. Although the additional collected information facilitates the decision-maker to make more informed decisions, inspections and monitoring activities however incur significant costs. Especially, continuous monitoring for several years might be difficult due to sensor installation and maintenance costs, and limited lifetime of sensors. The concept of digital twin has therefore been introduced in order to continuously simulate the physical assets, retrieve information and make decisions.

A digital twin is a virtual model designed to reflect a physical object. The object being studied, e.g., a wind turbine, is equipped with sensors in critical areas of functionality. These sensors produce data about different aspects of the physical object's performance, such as energy output, load information, environmental conditions, etc., which is then relayed and applied to

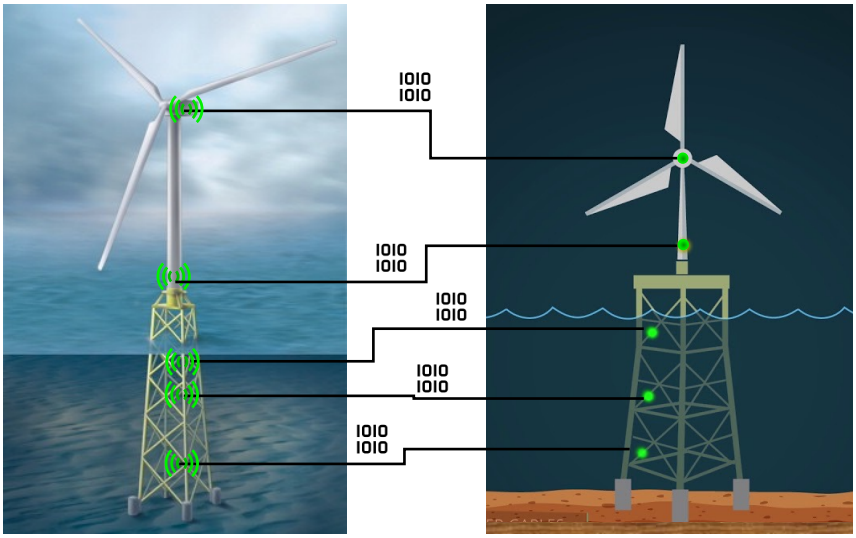

Fig. 1.1. A jacket-supported offshore wind turbine and its digital representation.

the digital copy. Once informed with such data, the virtual model can be used to run simulations and identify possible improvements, which can then be applied back to the original physical object. However, in this loop of physical-virtual interaction, there underlie some questions to be resolved- from the modeling complexity of the digital twin to the efficiency/frequency of updating the digital twin. How the information retrieved from a digital twin can be useful to improve life-cycle decisions on the physical asset also becomes a crucial aspect.

In summary, there is a need for integrated life-cycle management planning that already takes into account the possibility of developing and deploying digital twins. The current life-cycle management methodologies should be extended to be able to answer the following questions in addition to classical inspection and maintenance (I&M) planning questions of when/where/how to inspect or maintain:

- When to install sensors?
- Where to install sensors?
- When a digital twin should be updated or re-calibrated?

This thesis is devoted to propose solutions to these questions, ultimately leading towards efficient and cost-effective digitalization and life-cycle management of infrastructures such as offshore wind turbines.

1.2 Decision-making under uncertainty: State-of-the-art

Life-cycle management planning for offshore wind structures forms a complex decision-making problem under uncertainty:

- The stochastic offshore environments combined with the dynamic behaviour of wind turbines induce aleatory uncertainties.
- Due to limited statistical data available for offshore wind turbines, especially as technology and designs evolve rapidly, the estimation of deterioration mechanisms and failure probabilities relies on analytical models or numerical simulations which are often associated with epistemic uncertainties.
- Although collecting additional data can reduce statistical/epistemic uncertainties, inspections and monitoring systems are also subject to measurement uncertainties.

Under these uncertain circumstances, decision-makers need not only to identify optimal maintenance actions but also to determine when and where to gather information in order to minimize the expected life-cycle cost. The optimal policy in a decision-making problem can be theoretically identified by conducting a pre-posterior analysis, taking into account the probabilities and consequences of all the potential events [5]. However, in practice, the number of branches in a decision tree exponentially grows with respect to the time horizon (curse of history), number of actions and observations (curse of dimensionality), and becomes computationally intractable. State-of-the-art I&M planning methods often introduce simplifications in order to alleviate this challenge. Table 1.1 provides a systematic overview of decision-making approaches available in the literature, classified in terms of optimization method and practical applications.

Many I&M planning problems rely on pre-defined heuristic decision rules to control the computational complexity [6]. For instance, inspections are often planned at equidistant intervals or upon exceeding a certain failure probability threshold, and maintenance actions are conducted when a problem is detected during inspections, e.g., detection of cracks [7]. By imposing heuristic decision rules, the decision-making problem is simplified and I&M strategies can be identified within a reasonable computational time. Heuristics-based methods have often been implemented in the I&M planning of offshore wind structures [8, 9, 10]. Dynamic Bayesian Networks (DBNs) are integrated with heuristics [11, 12, 13, 14, 15] in recent I&M planning approaches, offering an efficient and robust inference for Bayesian updating with observed information. Even though heuristic decision rules alleviate the computational complexity in I&M decision-making problems, the obtained I&M strategies are constrained by the number of evaluated pre-defined rules out of the vast available policy space. Moreover, when I&M optimization is extended to the system level, defining the decision rules is not as simple as in the component level due to the increased number of action/observation combinations among the components [12, 16, 17].

On the other hand, the I&M planning problems are formulated as a partially observable Markov decision process (POMDP) [18, 19, 20], a principled mathematical framework for decision-making under partial observability. Solving POMDPs also demands computational expense, and value iteration algorithms were able to solve only small state space problems. However, with the development of point-based solvers [21, 22, 23, 24], the inherent complexities of the solution process have been alleviated, by focusing the computation on a set of belief points

Table 1.1. Survey of inspection and maintenance planning methods available in the literature.

Ref	Optimization method			Application	
	Heuristics	POMDP	DRL (MARL)	OWT	Others
[30]	✓			✓	
[10]	✓			✓	
[26]		✓			✓
[27]		✓			✓
[31]	✓			✓	
[25]		✓		✓	
[32]	✓				✓
[33]	✓				✓
[34]	✓			✓	
[14]	✓				✓
[35]		✓			✓
[12]	✓				✓
[36]		✓		✓	
[37]		✓			✓
[38]			✓		✓
[39]		✓		✓	
[40]			✓		✓
[41]			✓		✓
[29]		✓		✓	✓
[42]	✓				✓
[43]			✓		✓
[44]			✓	✓	
[45]			✓		✓
[46]			✓		✓
[47]	✓				✓
[48]	✓			✓	
[49]			✓	✓	✓
[50]			✓		✓
[51]			✓		✓

that are representative of reachable beliefs. POMDP-based approaches have been successfully used in I&M planning of deteriorating systems [25, 26, 27, 28]. Moreover, POMDP-based policies are demonstrated to be adaptive and cost-effective than traditional heuristics-based strategies, and point-based algorithms can efficiently provide optimal I&M policies within a reasonable computational time [29].

When I&M planning tackles very large engineering systems such as bridge networks, offshore wind farms, POMDPs also suffer the curse of dimensionality. Addressing the computational complexities of managing large engineering systems, some life-cycle management planning methods rely on deep reinforcement learning (DRL) or multi-agent reinforcement learning (MARL) algorithms in order to identify I&M strategies [38, 43]. Several state-of-the-art MARL

algorithms are available, mainly categorized either as value-based or policy-based. Value-based methods focus on optimizing the value function, i.e., the expected reward starting from a state or state-action pair following a particular policy, to guide action selection. Some examples include deep Q-networks (DQN) [52], independent Q-learning (IQL) [53], and QMIX [54], etc. On the other hand, policy-based methods directly learn the optimal policy through neural networks, by estimating the gradient of the policy's performance and updating the policy in the direction that increases the expected reward. Actor-critic methods also expand upon policy-based algorithms, using two network models: one to find the optimal policy (actor) and one to compute the value function (critic), e.g., multi-agent deep deterministic policy gradient (MADDPG) [55], multi-agent proximal policy optimization [56, 57], among others. Actor-critic methods have been adopted for I&M planning of engineering systems in complex scenarios, such as under budget constraints [40], or when there are various forms of correlation among the system's components [43]. A benchmark of state-of-the-art multi-agent reinforcement learning algorithms with the specific application of managing the life-cycle of large scale infrastructures can be found in [49].

1.3 Virtual monitoring: State-of-the-art

Virtual monitoring methods play a pivotal role in addressing the challenges associated with obtaining direct measurements in various engineering applications [58]. These methods are particularly valuable when accessibility issues prevent the deployment of physical sensors [59, 60], enabling indirect estimation of critical information that is otherwise difficult to obtain. Table 1.2 summarizes the overview of the reported virtual monitoring models in engineering community, categorized under modeling approach, probabilistic reasoning and their applications.

Virtual monitoring can be modeled either as physics-based [68, 69, 80] or data-based [4, 62, 72, 81, 82], or a combination thereof [79]. Physics-driven virtual sensors, often referred to as white-box models, rely on fundamental physics laws and governing equations to estimate unmeasured quantities. Physics-based models may be periodically updated with in-situ observation data [66, 83]. On the other hand, data-driven or black-box models use machine learning techniques such as regression models or neural networks to establish empirical relationships between input variables and target responses. Data-based models are flexible and can efficiently identify meaningful patterns in large datasets and have been widely used as virtual sensing models in various engineering applications [84, 85, 86]. Physics-informed grey-box models combined the strengths of conventional physics-based models and data-based machine learning models. Uniquely, these models integrate domain-specific insights by embedding the fundamental physics laws into the learning framework [87, 88, 89]. This integration ensures adherence to the established rules of the system, being particularly beneficial in scenarios where only limited training data is available.

In the specific case of offshore wind turbines, virtual monitoring methods are also required to tackle additional challenges on top of maximizing prediction accuracy: (i) the quantification

Table 1.2. Survey of virtual monitoring methods available in the literature.

Ref	Modeling approach		Probabilistic reasoning		Application		
	Physics	Data	Deterministic	Probabilistic	OWT (SS)	OWT (B & RNA)	Others
[61]	✓		✓		✓		
[62]	✓		✓		✓		
[63]	✓		✓		✓		
[64]	✓		✓		✓	✓	
[65]		✓	✓				✓
[66]	✓		✓		✓		
[67]		✓	✓		✓		
[68]	✓		✓		✓		
[69]	✓			✓	✓		
[70]		✓	✓			✓	
[71]		✓		✓	✓		
[72]		✓	✓		✓		
[73]		✓	✓		✓		
[74]	✓		✓			✓	
[75]	✓	✓	✓			✓	
[76]	✓	✓	✓		✓		✓
[77]	✓			✓			✓
[78]	✓			✓		✓	
[79]	✓	✓		✓	✓		

SS = Support structures; B & RNA = Blade and rotor nacelle assembly

of uncertainties in model predictions, and (ii) implementation across an entire wind farm. As mentioned before, the prediction of complex structural responses is associated with significant uncertainties. The aleatory uncertainty arises from natural random phenomena such that the turbine response to a single set of input parameters is not unique. Since conventional numerical simulators typically rely on deterministic modelling, this aleatory uncertainty is often addressed by conducting multiple simulations with different random seeds for each set of inputs, e.g., wind, wave, yaw angle, etc., as recommended by offshore wind design practices [90, 91]. Deployment of virtual monitoring models also induces additional epistemic uncertainty, which is often quantified by comparing model predictions with respect to sensor data [72, 92]. However, sensors are also prone to damage/non-functionality in harsh marine environment, and there is often a lack of ground truth labels. Since the dynamics of offshore wind structures changes over time [93], a virtual monitoring model becomes less accurate after a certain number of years as well as the previously quantified model uncertainty is no longer representative.

Another important challenge is the farm-wide applicability of virtual monitoring approaches. Due to the financial impracticality of fully instrumenting all wind turbines, virtual monitoring models often rely on data from only one or a handful of 'fleet-leader' turbine(s). This approach

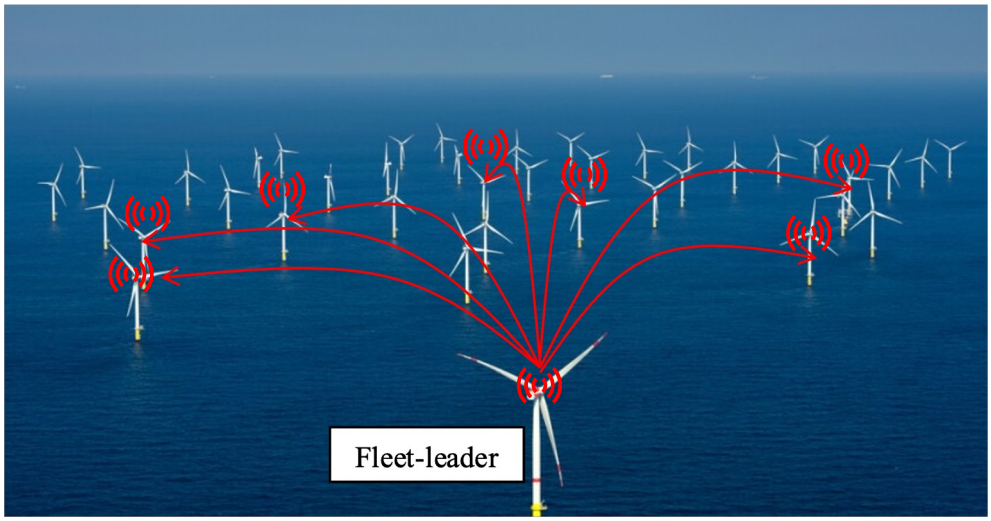

Fig. 1.2. Farm-wide deployment of fleet-leader based virtual monitoring model.

may lead to inaccuracies when models, based on a limited turbine set, are deployed to differently designed turbines within the same farm [72, 94], thus hindering their farm-wide application.

Digital twin concept

A digital twin is a virtual replica of real-world assets, processes or systems on which simulations can be run to predict the behavior of the physical twin. In this context, virtual monitoring models also have potential to be deployed as digital twins, in addition to the former objective of replacing sensors. The 'digital twin' concept proposed by [95] consists of three distinct parts: the physical object or process and its physical environment, the digital representation of the object or process, and the communication between the physical and virtual representations. The definition of the physical object and its physical environment will principally govern the complexity of a digital twin. Even if a structural system is of interest, one may focus only on a single sub-component of the system and its specific deterioration mechanisms.

In the context of virtual representation, a digital twin is often built with updatable parameters which are calibrated with respect to the collected data. Physics-based models are especially useful when specific knowledge of model parameters is available. The whole physical system can be twinned in this approach, thus allowing to infer measurements all over the geometry of the object. In complex engineering systems, retrieving certain model parameters can however be challenging, limiting the practicality of physics-driven digital twins. On the other hand, data-based models do not require such knowledge and are capable of capturing intricate patterns in the data, making them suitable for digital twinning of complex engineering structures. However, not having a physical model may cause lack of interpretation capabilities and extrapolation to previously unmeasured locations in the object, and therefore may require component-wise representation.

The essential characteristic of a digital twin is the communication with its physical object. The connection between the physical and digital versions is a closed loop, i.e., not only data from the physical subject will be used by the virtual twin in order to improve its representation, but also the predictions from the digital twin should also be used for the physical twin/environment. In physical-to-virtual communication, when the digital twin should be updated with the newly collected data is an important aspect to be considered since it will also significantly influence the modeling complexity. Real-time calibration can be computationally expensive, and might only be necessary in cases where prompt decisions are to be made, e.g., turbine control systems. However, in the specific case of offshore wind support structure I&M management, the deterioration mechanisms are gradual and also, decisions are made on a yearly basis. In such case, the digital twin might not need to be calibrated every minute or even day, and longer periods are tolerant as long as the model accuracy is enough for supporting decisions. However, the virtual-to-physical communication can be still efficient as it only depends on the computational time to simulate the digital twin. To complete a closed loop digital twin framework, the information from the digital twin should also be re-used to improve decisions on physical assets, e.g., I&M planning, lifetime extension or design improvement.

1.4 Objectives of the research

Most existing I&M planning frameworks do not consider digital twin and/or virtual monitoring in the life-cycle management. Although sensor installation is sometimes considered in the action space [44, 96], the monitoring data is only used to update the deterioration, disregarding the potential of developing a digital twin for further information retrieval. On the other hand, several approaches have been proposed for offshore wind turbine/farm virtual monitoring, only with the focus of improving prediction accuracy, farm-wide application etc., but not specifying how they can be used to improve I&M or lifetime extension decisions. Few digital twin based decision-making methods can be found in the literature [97, 98, 99], presuming the existence of sensors and/or digital twin over the decision horizon. In principle, there is a need for decision-making frameworks which take into account the possibility of digital twin development and deployment in pre-posterior reasoning.

The overarching aim of this thesis is to contribute to the development of probabilistic virtual monitoring models for offshore wind turbines and a life-cycle management framework using digital twins. The proposed life-cycle management planning methodology should be able to identify when to (re)train a digital twin, as well as to control the risk of failure events by optimally allocating inspections and maintenance interventions. Specifically, this investigation targets to achieve the following objectives:

- Objective 1: To apply state-of-the-art decision-making methods in life-cycle management planning for offshore wind support structures;

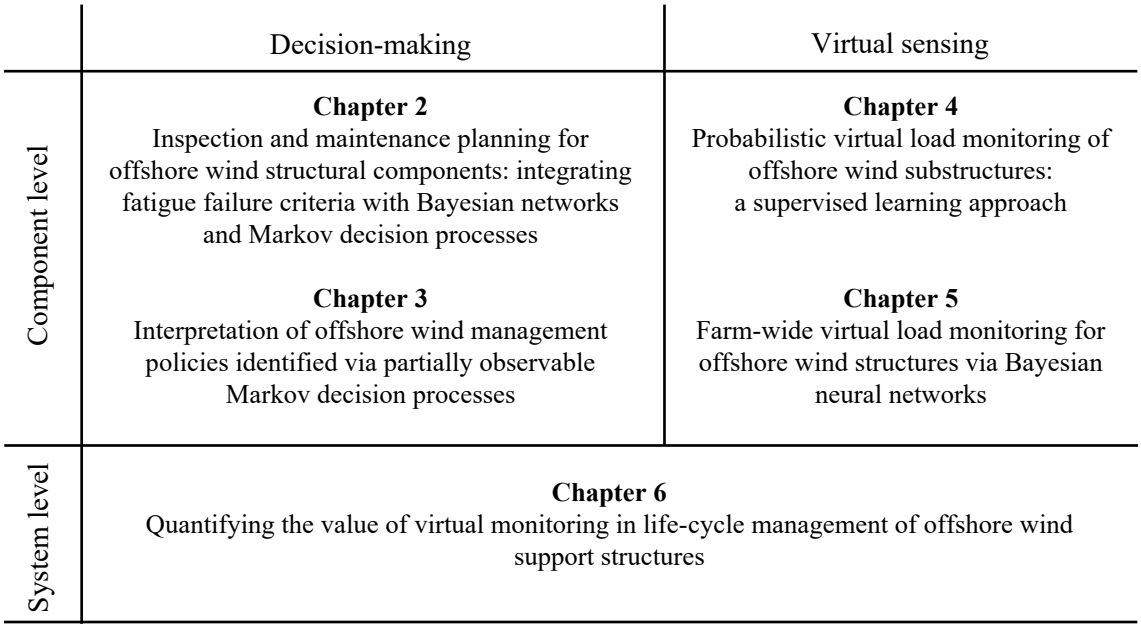

Fig. 1.3. Organization of the PhD dissertation.

- Objective 2: To develop a probabilistic virtual monitoring framework for offshore wind farms, enabling the quantification of aleatory and epistemic uncertainties in the estimated predictions;
- Objective 3: To integrate virtual monitoring (digital twin) into life-cycle management of offshore wind substructures;
- Objective 4: To quantify the value of virtual monitoring in life-cycle management and to identify optimal model re-calibration intervals.

1.5 Outline of the thesis

Addressing the aforementioned objectives, the developments in this work are organized as illustrated in Figure 1.3, first presenting the contributions with regard to decision-making for I&M planning of offshore wind structural components, followed by the development of virtual sensing models. At last, the integration of the developed virtual monitoring scheme with the life-cycle management framework is presented. Each following chapter of this thesis corresponds to a published, submitted or under review paper, reaching the concluding remarks and further research directions in Chapter 7. In all the papers, the author of this thesis holds a primary responsibility for the conceptualization of the study, formal analysis, investigation, interpretation of the results, and writing. Additionally, the contributions of each co-author are listed at the end of each chapter.

First, I&M planning is performed for an offshore wind structural component in Chapter 2, relying on partially observable Markov decision processes (POMDPs). A deterioration rate

dynamic Bayesian network (DBN) is adopted to model various fatigue/fracture mechanics models and failure definitions, which are commonly used for offshore wind substructures. The DBN-POMDP based I&M planning is performed for an offshore wind turbine's fatigue hotspot, considering multiple levels of redundancy in the failure definition. The results demonstrate that POMDPs are flexible and adaptive, providing optimal I&M strategies according the uncertainty of the deterioration model, failure redundancy and the measurement uncertainty of inspection techniques.

Thereafter, interpretation and analysis of POMDP-based I&M strategies are presented in Chapter 3. The objective is to accelerate the practical implementation of POMDPs by offshore wind asset management decision-makers, who may only be accustomed to calendar- and/or condition-based approaches. Adaptive I&M policies and expected economic "regret" are also analyzed in case a decision-maker opts for an action other than the POMDP-suggested optimal policy.

In Chapter 4, a virtual load monitoring framework is introduced for fatigue load prediction in offshore wind structures. Relying on supervised learning algorithms, artificial neural networks are used to map either high or low-frequency SCADA data to the concurrent structural response data, encoding aleatory uncertainty in the output as a probability distribution. The proposed method is tested on a simulated dataset of NREL 5-MW offshore wind turbine and the results show a strong agreement between predicted and actual values of stress range distributions and fatigue damage estimates.

The development of virtual monitoring is further extended in Chapter 5, proposing a fully probabilistic scheme along with the uncertainty quantification framework using Bayesian neural networks (BNNs). Not only the aleatory uncertainty is modeled in the output distribution but also the model (epistemic) uncertainty is characterized by random weights. The proposed virtual monitoring framework is tested on an in-situ dataset from three wind turbines in a Belgian offshore farm. The results demonstrate that BNNs can automatically inform if the provided predictions are inaccurate, thus allowing deployment to non-fully instrumented turbines in farm-wide monitoring.

The integration of the developed probabilistic virtual monitoring into life-cycle management framework is presented in Chapter 6. Once the sensors are installed to collect training data, BNNs are used as virtual sensors in the following years, providing load prediction and its uncertainty information. Applying state-of-the-art multi-agent reinforcement learning algorithms, life-cycle management strategies are then planned for a whole offshore support structure consisting of several components. The expected benefit of using virtual monitoring system is quantified, relying on the value of information principle.

1.6 List of papers

Papers included in this thesis

The following journal and conference papers are included in this thesis, selected as the most direct and relevant contributions to address the objectives stated in Section 1.4:

- **Paper 1:** Hlaing, N., Morato, P. G., Nielsen, J. S., Amirafshari, P., Kolios, A. and Rigo, P. (2022). Inspection and maintenance planning for offshore wind structural components: Integrating fatigue failure criteria with Bayesian networks and Markov decision processes. *Structure and Infrastructure Engineering, 18*(7), 983-1001.

- **Paper 2:** Hlaing, N., Morato, P. G., Papakonstantinou, K. G., Andriotis, C. P. and Rigo, P. (2022). Interpretation of offshore wind management policies identified via partially observable Markov decision processes. In *The 18th European Academy of Wind Energy (EAWE) PhD seminar*.

- **Paper 3:** Hlaing, N., Morato, P. G. and Rigo, P. (2022). Probabilistic virtual load monitoring of offshore wind substructures: A supervised learning approach. In *The Proceedings of the 32nd International Ocean and Polar Engineering Conference (ISOPE 2022)*, volume 4, 3137–3144.

- **Paper 4:** Hlaing, N., Morato, P. G., de Nolasco Santos, F., Weijtjens, W., Devriendt, C. and Rigo, P. (2023). Farm-wide virtual load monitoring for offshore wind structures via Bayesian neural networks. *Structural Health Monitoring*.

- **Paper 5:** Hlaing, N., Morato, P. G., and Rigo, P. Quantifying the value of virtual monitoring in life-cycle management of offshore wind support structures. Under internal review.

Other contributions

The following additional contributions, also developed during the PhD research, are not included in this thesis:

- **Paper 6:** Hlaing, N., Morato, P. G., Nielsen, J. S., Amirafshari, P., Kolios, A. and Rigo, P. (2021). The effect of failure criteria on risk-based inspection planning of offshore wind support structures. In *The Proceedings of the 7th International Symposium on Life-Cycle Civil Engineering (IALCCE 2021)*.

- **Paper 7:** Morato, P. G., Papakonstantinou, K. G., Andriotis, C. P., Hlaing, N. and Kolios, A. (2023). Interpretation and analysis of deep reinforcement learning driven inspection and maintenance policies for engineering systems. In *The Proceedings of*

the 14th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP14).

- **Paper 8:** Morato, P. G., Andriotis, C. P., Papakonstantinou, K. G., Hlaing, N. and Rigo, P. Optimal management of offshore wind structural systems via deep reinforcement learning and Bayesian networks. Under internal review.

- **Dataset 1:** Hlaing, N. and Morato, P. G. (2022). Post-processed dataset from 50000 numerical simulations of monopile-supported NREL 5MW wind turbine in OpenFAST. Zenodo. URL: https://doi.org/10.5281/zenodo.5957394

- **Open-source code 1:** Hlaing, N. and Morato, P. G. BayesianNN: Implementation of Bayesian neural networks for virtual monitoring (offshore wind farms) and uncertainty quantification. URL: https://github.com/Nandarhline/BayesianNN

- **Open-source code 2:** Hlaing, N. IMP-MARL (Digital twin): Integration of virtual sensing in the life-cycle management. URL: https://github.com/Nandarhline/imp_marl/tree/digital_twin, forked from https://github.com/moratodpg/imp_marl.

Chapter 2

Inspection and maintenance planning for offshore wind structural components: Integrating fatigue failure criteria with Bayesian networks and Markov decision processes

---

Paper Hlaing, N., Morato, P. G., Nielsen, J. S., Amirafshari, P., Kolios, A. and Rigo, P. (2022). Inspection and maintenance planning for offshore wind structural components: Integrating fatigue failure criteria with Bayesian networks and Markov decision processes. *Structure and Infrastructure Engineering, 18*(7), 983-1001. DOI:10.1080/15732479.2022.2037667.

Abstract Exposed to the cyclic action of wind and waves, offshore wind structures are subject to fatigue deterioration processes throughout their operational life, therefore constituting a structural failure risk. In order to control the risk of adverse events, physics-based deterioration models, which often contain significant uncertainties, can be updated with information collected from inspections, thus enabling decision-makers to dictate more optimal and informed maintenance interventions. The identified decision rules are, however, influenced by the deterioration model and failure criterion specified in the formulation of the pre-posterior decision-making problem. In this paper, fatigue failure criteria are integrated with Bayesian networks and Markov decision processes. The proposed methodology is implemented in the numerical experiments, specified with various crack growth models and failure criteria, for the optimal management of an offshore wind structural detail under fatigue deterioration. Within the experiments, the crack propagation, structural reliability estimates, and the optimal policies derived through heuristics and partially observable Markov decision processes (POMDPs) are thoroughly analyzed, demonstrating the capability of failure assessment diagram to model the structural redundancy in offshore wind substructures, as well as the adaptability of POMDP policies.

2.1 Introduction

An optimal and rational management of offshore wind substructures is becoming increasingly important due to the growth of offshore wind installations, with a trend towards larger wind turbines, often located far offshore. Exposed to harsh marine environmental conditions, the degradation of offshore wind substructures is accentuated, thus inducing a risk of structural failure. Additionally, inspection and maintenance interventions may become more complex and expensive far offshore. In this context, inspection and maintenance (I&M) planning methods offer a framework for minimizing life-cycle costs, controlling structural failure risks by optimally allocating inspections and maintenance actions.

Already in the 1990s, early I&M planning methods address the decision-making problem by exploiting Bayesian decision analysis with the objective to identify optimal strategies for structures subjected to fatigue deterioration [6, 100], with many applications focused on I&M planning for offshore structures [101, 102]. By defining the I&M policies based on a set of predefined decision rules, the computational complications associated with solving an extensive pre-posterior analysis were alleviated, enabling the identification of rational strategies within a reasonable computational time [103, 104]. Heuristic-based I&M planning methods have been widely applied to the management of fatigue-sensitive structures, planning inspections at periodic intervals or immediately after a specified failure probability threshold is exceeded [105, 106]. More recently, the integration of discrete dynamic Bayesian networks (DBNs) into I&M methodologies has enabled the efficient evaluation of more sophisticated heuristic decision rules [34]. For instance, [12] has proposed an I&M planning approach for structural systems, evaluating system-level heuristic decision rules in a DBN simulation environment. Other existing I&M research works consider multiple conflicting objectives within the policy optimization [107, 108], planning maintenance actions, in some cases, based on specified thresholds [109]. Relying also on dynamic Bayesian networks, single- and multi-objective optimization methods provide robust Bayesian inference and enable the evaluation of advanced decision rules, e.g., adaptive repair thresholds [14]. Through multi-objective policy optimization methods, decision-makers can operate under budget constraints and/or control maintenance delays.

Even if heuristic decision rules alleviate the computational complexity of the I&M decision-making problem, as mentioned before, the obtained I&M strategies are constrained by the number of evaluated pre-defined rules out of the vast available policy space. Instead, I&M planning methods that rely on Markov decision processes determine adaptive policies, providing a mapping from the dynamically updated deterioration state to the optimal actions. Various I&M planning applications for deteriorating structures modeled the decision-making problem via Markov decision processes, e.g., [18, 19, 20]. The benefits offered by adaptive I&M policies are substantiated by [28], and [29] have demonstrated that partially observable Markov decision process (POMDP) solved via point-based algorithms can efficiently determine optimal I&M poli-

cies. An overview of state-of-the-art point-based solvers and their applicability to infrastructure management can be found in [110].

As explained before, I&M planning methods aim at controlling arising structural failure risks by timely allocating inspection and maintenance actions. Essentially, the failure risk of a structural component represents both the probability of a failure event and its associated economic, societal, and environmental consequences. The estimation of the failure probability is governed by a failure criterion, which is specified by the decision-maker. Within the context of fatigue deteriorating structures, a through-thickness failure criterion is normally prescribed, e.g., asset management of offshore pipelines and containers. This conventional criterion might be over-conservative for redundant structures such as the tubular joints of jacket-type offshore wind turbines (OWTs), as tubular connections have the capacity to sustain through-thickness cracks until the loading exceeds the resistance of the cracked structure. In such applications, the fatigue failure limit state can be, instead, formulated via a failure assessment diagram (FAD).

The failure assessment diagram, originally proposed by [111] and [112], describes the interaction between brittle fracture and plastic failure through a two-parameter failure criterion. The specification of FAD as the governing failure criterion has recently gained attention in offshore wind applications. Among them, [113] and [114] showcased the evaluation of flaw acceptability in offshore wind support structures using a failure assessment diagram of BS7910 [1]. In parallel with the reported FAD research work, several probabilistic fatigue studies and I&M planning methods for offshore wind substructures have still specified fatigue limit states based on the conventional through-thickness failure criterion, potentially drawing over-conservative conclusions.

In the reported I&M planning methods that formulate the fatigue failure limit state via a failure assessment diagram [115, 116], the identified I&M policies are, however, based on heuristic decision rules. In this paper, we originally integrate the modeling of stochastic fatigue deterioration processes and the specification of FAD-based limit states via dynamic Bayesian networks, and we introduce the necessary formulation for modeling the overarching I&M decision-making problem as a partially observable Markov decision process (POMDP). The proposed method is flexible and can be easily adopted by other applications whose limiting failure criterion is also defined as a function of multiple failure parameters. The applicability and efficacy of the proposed approach is verified through numerical experiments, in which optimal I&M strategies are determined for the specific case of an offshore wind tubular joint. Within the numerical experiments, the fatigue deterioration of offshore wind structural details is modeled by one-dimensional and two-dimensional fracture mechanics methods as well as various failure criteria, thoroughly investigating the effect of model selection on the identified I&M strategies. The results reveal that the choice of failure criteria and the optimality of the implemented I&M planning methods significantly affect the resulting I&M policies. In particular, the benefits of adopting FAD criteria for offshore wind substructures and the cost savings provided by POMDP-based policies are both meticulously discussed.

2.2 Background: Risk-based inspection and maintenance planning

Risk-based inspection and maintenance planning is based on pre-posterior decision analysis integrated with deterioration modeling, inspection and repair modeling, and cost modeling. This section presents deterioration modeling including failure criteria and inspection modeling. Cost modeling along with policy optimization methods is discussed in Section 2.4. The probabilistic fatigue deterioration model is used as reference and fracture mechanics models are calibrated to the fatigue model. The through-thickness failure criterion has been commonly used in the I&M planning of offshore wind structures [17, 96]. In this paper, the failure assessment diagram criterion is also addressed and integrated into I&M planning. Procedures to obtain a FAD and failure assessment points are also presented.

2.2.1 Deterioration modeling

Probabilistic SN model

Offshore wind turbine support structures are subjected to a large number of environmental load cycles (e.g., waves) and other operational loading in their lifetime. For such structures, the long-term stress range distribution can be represented by a Weibull distribution, described by a scale parameter q and a shape parameter h. The shape parameter h for offshore wind (OW) substructures as recommended by DNV standards [117] is 0.8 and the scale parameter q is computed such that the cumulative lifetime damage of the structural component designed for T_d years is equivalent to the damage limit corrected by the fatigue design factor (FDF), i.e., $D_{SN}(t = T_d) = 1/FDF$ where the temporal fatigue evolution D_{SN} is defined as:

$$D_{SN}(t) = nt\left[\frac{q^{m_1}}{k_1}\gamma_1\left(1+\frac{m_1}{h};\left(\frac{S_1}{q}\right)^h\right) + \frac{q^{m_2}}{k_2}\gamma_2\left(1+\frac{m_2}{h};\left(\frac{S_1}{q}\right)^h\right)\right], t = 0, 1, ...T_d, \quad (2.1)$$

m_1, m_2, k_1, k_2, S_1 are parameters of the bi-linear SN curve. n is the number of stress cycles per year. γ_1 and γ_2 are the upper and lower incomplete gamma functions.

Probabilistic fatigue analysis can then be based on Palmgren-Miner's rule with the long-term stress range distribution. The limit state function applied in the fatigue deterioration model is:

$$g_{SN}(t) = \Delta - D_{SN}(t), \quad (2.2)$$

where Δ is the fatigue limit beyond which failure happens and is considered as a random variable owing to the uncertainty of Palmgren–Miner's rule. $D_{SN}(t)$ is the accumulated fatigue damage at year t calculated as in Equation (2.1).

The design of offshore structures is based on SN curves but the inspection information, i.e., the presence of crack or the crack size measurement, cannot be directly used to update the SN-based reliability. Inspection and maintenance planning therefore demands the use of fracture mechanics (FM) models. Different FM models are described in the following sections. It becomes necessary to relate the two deterioration models and such a relation can be attained by calibrating the FM models to the SN model which includes all the information from the design stage. The calibration is performed such that a similar fatigue life is calculated by fracture mechanics models as that of S-N test data. Whereas the FM models compute the crack growth, the only information that contains in the SN model is the failure or survival of the hotspot through fatigue damage. Therefore, the calibration between SN and FM models has been based on the probability of failure along the lifetime [117]. Typically, FM parameters with the largest influence on the crack growth and/or with the least available information are calibrated [103].

Two-dimensional crack growth model

In offshore wind support structures, fatigue cracks initiate from manufacturing imperfections and welding defects as illustrated in Figure 2.1. The severity increases over the time as the cracks grow under the cyclic loading of wind and waves. Paris-Erdogan's law has been widely used in linear elastic fracture mechanics (LEFM) to model the crack growth [118]:

$$\frac{d}{dn}d = C_d \left(\Delta K_d\right)^m, \tag{2.3}$$

$$\frac{d}{dn}c = C_c \left(\Delta K_c\right)^m, \tag{2.4}$$

where d is the crack depth growing through the thickness of the structural member and c is the crack length growing along the surface. n corresponds to the number of stress cycles. C_d, C_c and m are Paris law parameters, also called crack growth parameters and ΔK is the stress intensity factor range at the crack tip calculated as:

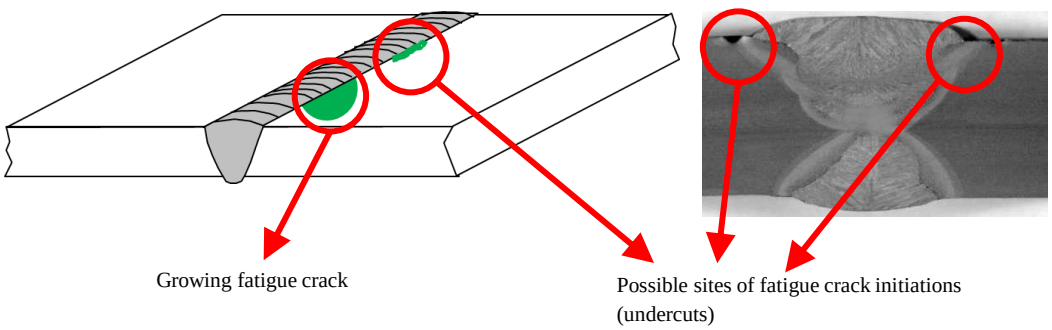

Fig. 2.1. Illustration of fatigue crack initiation.

$$\Delta K_d = \Delta\sigma Y_d(d,c)\sqrt{\pi d}, \tag{2.5}$$

$$\Delta K_c = \Delta\sigma Y_c(d,c)\sqrt{\pi d}, \tag{2.6}$$

where Y_d and Y_c are stress intensity correction factors and theoretically, are dependent on the geometry of the component, welded joint detail and time-varying two-dimensional crack size. The applied stress range $\Delta\sigma$ is assumed to be composed of membrane and bending stress components. The two components are quantified by the ratio of bending stress to total stress, denoted as the degree of bending D_b. Stress concentration due to weld geometry is also incorporated as the stress magnification factor M_k. The stress intensity factor ranges can then be described as:

$$\Delta K_d = \Delta\sigma \left[Y_{md}M_{kmd}(1-D_b) + Y_{bd}M_{kbd}\ D_b\right]\sqrt{\pi d}, \tag{2.7}$$

$$\Delta K_c = \Delta\sigma \left[Y_{mc}M_{kmc}(1-D_b) + Y_{bc}M_{kbc}\ D_b\right]\sqrt{\pi d}. \tag{2.8}$$

The subscripts d, c refer to crack depth and crack length and m, b refer to membrane and bending stress components respectively. Geometry functions $Y_{md}, Y_{bd}, Y_{mc}, Y_{bc}$ and stress magnification factors $M_{kmd}, M_{kbd}, M_{kmc}, M_{kbc}$ can be solved by using parametric equations, for instance, as in BS7910 [1]. Alternatively, one can perform finite element analysis of the cracked structure and directly compute the stress intensity factors K_d or K_c at each stress cycle [113].

One-dimensional crack growth model

In the one-dimensional fracture mechanics model, the crack propagation is considered only in the direction of the member's thickness. Therefore, the stress intensity correction factor Y_d simply becomes a function of crack depth only. Additionally, if it is further simplified such that the stress intensity correction factor Y_d does not depend on the time-varying crack depth and is approximated as a constant value over the lifetime, an explicit solution of the crack growth can then be obtained as follows:

$$d(t) = \left[\left(1-\frac{m}{2}\right)C_d Y_d^m \pi^{m/2}(\Delta\sigma)^m n + d_{t-1}^{1-m/2}\right]^{(1-m/2)^{-1}}. \tag{2.9}$$

Through this simplification, the crack propagation can be analytically computed attenuating the computation associated to solving the coupled equations. For small cracks situated far away from the boundaries of the structural member, Y_d can be taken as 1 with sufficient accuracy [119]. In probabilistic deterioration modeling, Y_d is assigned as a time-invariant random variable to introduce the model uncertainties associated to the simplifications of the stress intensity correction factor [2].

Through-thickness failure criterion

As illustrated in Figure 2.1, a crack grows both through the thickness and along the surface of the structural component taking a semi-elliptical shape. Assuming that the thickness is smaller than the length and the width of the member, the crack is likely to penetrate the whole thickness first. The failure criterion can be formulated depending on the capacity of the structure to further resist the applied load after through-thickness penetration.

In the through-thickness criterion, the failure happens when the crack depth reaches the thickness of the structural member which is also denoted as the critical crack size d_{crit}. This common criterion is particularly adopted for structures containing pressurized containment, e.g., pipelines, pressure vessels, etc. and is conservative for redundant structures such as OW jacket foundations. The following limit state function is employed for the through-thickness failure:

$$g_{FM}(t) = d_{crit} - d(t), \tag{2.10}$$

in which $d(t)$ is the crack propagation over time. The probability of failure $P_F(t)$ is then the probability of the limit state function being negative such that $P_F(t) = P(g_{FM}(t) \leq 0)$.

Failure assessment diagram

When a crack propagates through a structural member, ultimately the crack size may reach a critical size which corresponds to a critical stress intensity factor, usually taken as the characteristic value of the fracture toughness K_{mat} at which brittle fracture happens. Alternatively, if the applied load is substantially high compared to the material tensile strength, the member may reach its tensile capacity and fail by plastic collapse. In between brittle fracture and plastic collapse is an elastoplastic failure mode, where the failure occurs before reaching the plastic capacity or fracture toughness. The failure assessment diagram was therefore introduced to combine the two failure modes [111].

The most rigorous method to obtain a FAD for a particular application is to perform an elastic-plastic J-integral analysis [120]. Since it can be cumbersome, simplified approximations are available. For instance, BS7910 [1] provide three alternative FAD options which have been frequently used in offshore wind applications [113, 114]. They are of increasing complexity in terms of the required material properties and stress analysis data but also provide results of increasing accuracy with less conservatism. An example of the FAD is plotted in Figure 2.2.

In a FAD, the ordinate plots the fracture ratio K_r, also called the crack-driving parameter which represents the structure's susceptibility to brittle fracture. The abscissa plots the load ratio L_r which measures how close the structure containing the crack is to plastic collapse. The load ratio L_r is equal to 1 at the yield limit, however plastic collapse happens at a higher value which is equal to $L_{r,max}$. The failure of a structural component is then defined by means of a failure assessment line (FAL). If an assessment point lies inside the envelope below the assessment line, the component is assumed to be safe. If it falls on or outside the FAL, it is

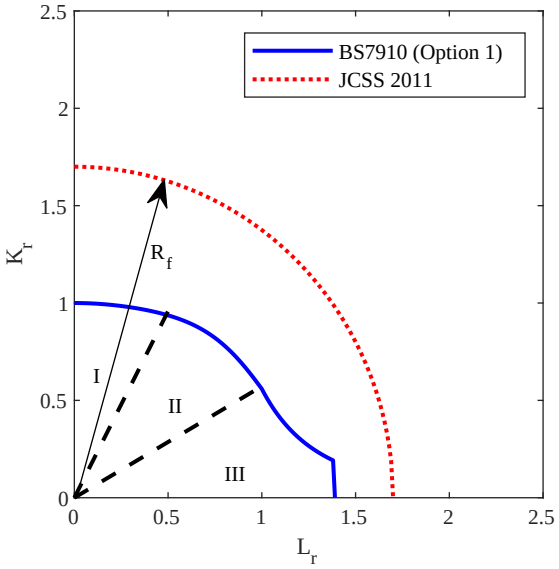

Fig. 2.2. Failure assessment diagram [1] and the simplified criterion [2].

assumed to be failed. The failure assessment line (FAL) is in fact a plot of the critical values of fracture ratio $K_{r,crit}$ for a range of load ratio, i.e., $0 \leq L_r \leq L_{r,max}$. The cut-off value for plastic collapse $L_{r,max}$ according to BS7910 [1] is:

$$L_{r,max} = \frac{\sigma_Y + \sigma_U}{2\sigma_Y}, \tag{2.11}$$

where σ_Y and σ_U are the design yield strength and ultimate strength of the material used. $K_{r,crit}$ is equal to 1 for fully brittle fracture and declines as the load ratio increases towards the collapse load as in Figure 2.2. In addition, as it is illustrated, the FAD can be partitioned into three different zones: Zone I is the fracture dominant zone, Zone II is the elastoplastic zone and Zone III is the plastic collapse dominant zone [115].

When the FAD is used as a limit state function, the failure occurs when the applied load exceeds the reduced capacity of the cracked structure. It becomes necessary to consider the combined influence of applied loads and non-monotonic strength deterioration of the cracked structure. Therefore, evaluation of the failure probability with a FAD requires to apply time-variant reliability methods which are extremely time-consuming. Instead, a simplified criterion proposed by [2] has been used in this work as an alternative to FAD. In this case, the failure is expected if the interaction of the crack-driving parameter K_r and the load ratio L_r exceeds a normalized resistance parameter R_f, see Figure 2.2. The concept of the normalized resistance parameter and the recommended values are described in [121] and [2]. Then, the limit state equation and assessment points K_r and L_r are reformulated as:

$$g_{FM}(t) = R_f - \sqrt{K_r^2(t) + L_r^2(t)}, \tag{2.12}$$

$$K_r = \frac{K_I}{K_{mat}} + \rho, \qquad L_r = \frac{\sigma_{ref}}{\sigma_Y}, \tag{2.13}$$

where R_f is the normalized resistance parameter. K_{mat} is the fracture toughness of the material. K_I is the stress intensity factor at the crack tip and can be computed for a particular crack size as follows:

$$K_I = \sigma Y_d(d, c)\sqrt{\pi d}, \tag{2.14}$$

where σ is the maximum applied stress. The plasticity correction factor $\rho \geq 0$ reflects the interaction between the applied primary loads and the secondary loads, e.g., residual stress R_s. Plasticity correction is important when the secondary loads are high which, for example, is the case of welded joints. In such case, ρ increases as the crack size becomes larger, representing the reduced load carrying capacity of the deteriorated structure driven by plasticity interaction effects. The plasticity correction can be evaluated according to the procedures in [2] or [1]. σ_Y is the material's yield stress and σ_{ref} is the net section stress or reference stress of the cracked structure. For a surface crack at the weld toe, σ_{ref} can be evaluated as in the following equation [1]:

$$\sigma_{ref} = \frac{(D_b \cdot \sigma) + \left((D_b \cdot \sigma)^2 + 9\left((1 - D_b) \cdot \sigma\right)^2 (1 - \mu'')^2\right)^{0.5}}{3(1 - \mu'')^2},$$

$$\text{where } \mu'' = \begin{cases} \frac{(d/t)}{1+(t/c)}, & \text{if } W \geq 2(c+t) \\ \frac{2d}{t}\frac{c}{W}, & \text{if } W < 2(c+t) \end{cases}, \tag{2.15}$$

where W and t are the width and the thickness of the structural member.

2.2.2 Inspection modeling

The information gathered during the operational lifetime can be used to update the uncertainties in the deterioration model. For instance, the probability distribution of the crack size can be updated after an inspection is performed. However, it is necessary to take into account the measurement quality of the observation model. The probability of detection (POD) curves have been adopted to characterize the quality of several non-destructive inspection techniques such as ultrasonic testing (UT), magnetic particle inspection (MPI) and eddy current (EC) inspection [122]. The POD depends on the crack size and the detection threshold, as illustrated in Figure 2.3. A signal response above the detection threshold will give an inspection outcome of crack detection and below the threshold results in a no-detection outcome. Accordingly, the POD is defined in detection theory [123] as follows:

$$POD(d) = \int_{\hat{d}_{s,th}}^{+\infty} f_{signal}(\hat{d}_s \mid d)\, d\hat{d}_s, \tag{2.16}$$

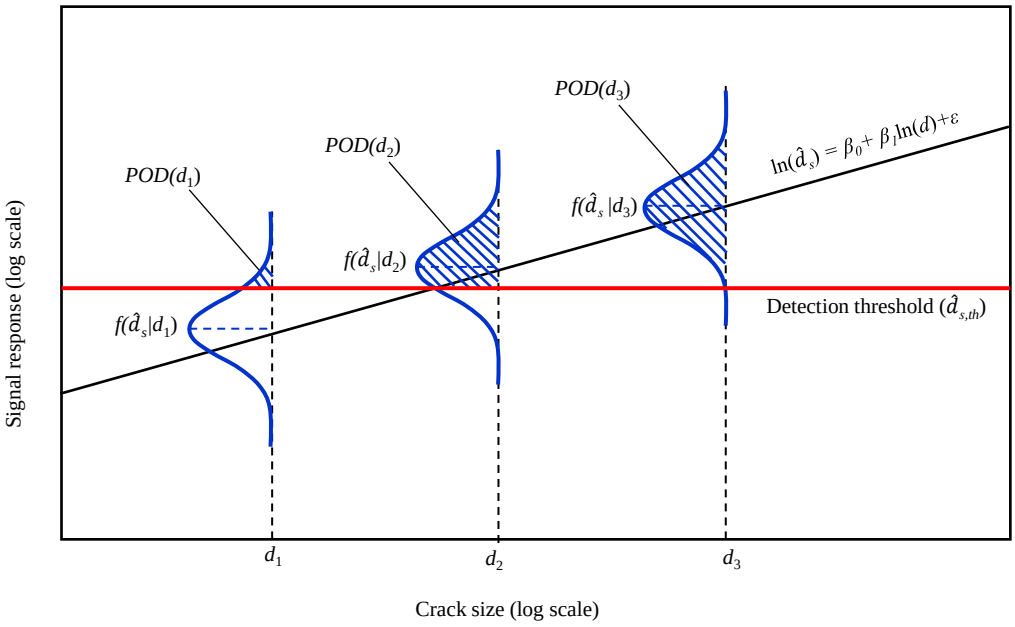

Fig. 2.3. Probability of detection, adapted from signal response method [3]. The mean values of the signal response fall on a regression line with the regression parameters β_0 and β_1. ε is associated with variability of imperfect inspection and is assumed normally distributed with a zero mean and a standard deviation σ_ε.

where d is the true crack size. $\hat{d}_{s,th}$ is the detection threshold and f_{signal} is the probability density function of the signal response $\hat{d}_s$. Given the regression parameters β_0, β_1 and the variability σ_ε of an inspection technique, Equation (2.16) can be derived as:

$$POD(d) = 1 - \Phi\left(\frac{\hat{d}_{s,th} - (\beta_0 + \beta_1 \cdot ln\,(d))}{\sigma_\varepsilon}\right), \tag{2.17}$$

where Φ is the cumulative distribution function of the standard normal distribution. For a particular detection threshold $\hat{d}_{s,th}$, the theoretical probability of detection curve according to Equation (2.17) is a monotonically increasing function of the crack size. Adjusting the detection threshold $\hat{d}_{s,th}$, the shape of a POD curve can be changed so that $POD = 1$ when $\hat{d}_{s,th} \to -\infty$ and $POD = 0$ when $\hat{d}_{s,th} \to \infty$ for any crack size. Based on [124] and [103], the following limit state function is used for the event of crack detection at time t:

$$g_D(t) = u - POD(d(t)), \tag{2.18}$$

where u is a uniformly distributed random variable in the interval $[0, 1]$. The POD of the crack depth at time t is computed according to Equation (2.17).

The additional information obtained from inspections can be used to update the reliability through conditional failure probability [125]. For instance, given no-detection outcome of

inspection at year t_{ins}, the updated failure probability for $t \geq t_{ins}$ is:

$$P_F(t) = P\left(g_{FM}(t) \leq 0 \mid g_D(t_{ins}) > 0\right) = \frac{P\left(g_{FM}(t) \leq 0 \ \cap \ g_D(t_{ins}) > 0\right)}{P(g_D(t_{ins}) > 0)}. \tag{2.19}$$

2.3 Stochastic deterioration modeling through dynamic Bayesian networks

Bayesian networks (BN), introduced by [126], is a graphical formalism to represent joint probability distributions of a set of random variables. Dynamic Bayesian networks (DBNs) are temporal repetitions of BNs which have been increasingly used in engineering structural reliability and risk analysis [11, 34, 127]. To implement DBNs in I&M planning, the continuous random variables involved in the deterioration model need to be discretized for prediction and exact inference tasks. This step is crucial since the accuracy of the results and the computational efficiency are influenced by the number of intervals and the discretized boundaries. Theoretically, the discretization error tends to 0 as the size of the intervals approaches 0. In practical applications, the discretization scheme is preferred to provide sufficient accuracy with maximum computational efficiency.

The state space $\mathcal{S}$ in DBNs is the domain of the discretized variables. In a stochastic deterioration process, the belief which is the probability distribution over the state space $P(s_t)$ transitions from one time step to the next one according the conditional probability $P(s_{t+1} \mid s_t)$, also denoted as the transition matrix. Markovian property is assumed here, i.e., the state at time $t+1$ depends only on the state at t and not on the past ones. Additionally, the transition matrix is time-invariant meaning that $P(s_{t+1} \mid s_t)$ is the same for any two consecutive time steps. Evidence from observations (e.g., inspections) can also be incorporated through Bayes' rule such that:

$$P(s_{t+1} \mid \boldsymbol{o_{t+1}}) \propto P(\boldsymbol{o_{t+1}} \mid s_{t+1}) P(s_{t+1}), \tag{2.20}$$

whereas the likelihood $P(\boldsymbol{o_{t+1}} \mid s_{t+1})$ quantifies the quality of the observation.

2.3.1 Deterioration rate DBNs adopting a through-thickness criterion

Dynamic Bayesian networks (DBNs) have been frequently used to model engineering deterioration processes in risk analysis, often through combination of random variables to reduce the dimension of the state space and the computation time [11]. When more complex deterioration model and failure criterion are used, the state space becomes high-dimensional and non-combinative, e.g., the Weibull scale parameter q cannot be combined with other variables since the crack size is conditional on it (through $\Delta\sigma$) and so are the failure assessment points K_r and L_r (through σ). In addition, the necessity of high-dimensional conditional probabilities for propagating the

belief and computing the failure probability, such as $P(d_{t+1} \mid d_t, c_t, q, C_a)$, $P(c_{t+1} \mid d_t, c_t, q, C_c)$, $P(K_{r,t} \mid d_t, c_t, q, K_{mat}, R_s)$, $P(L_{r,t} \mid d_t, c_t, q, \sigma_Y)$, $P(g_{FM} \mid L_r, K_r, R_f)$ increases computational complexity.

Another DBN representation, denoted here as "deterioration rate" DBN, represents a stochastic deterioration process as a function of the deterioration rate. The graphical representation of such DBNs adopting a through-thickness criterion is shown in Figure 2.4. The crack evolution is traced by the nodes d_t and is dependent on the deterioration rate τ_t. The node τ_t is a one-hot (one-zero) vector indicating the current deterioration rate. Unless any maintenance action is taken, it transitions one deterioration rate, i.e., τ_i to τ_{i+1}, at every time step. Note that the component may have the same deterioration rate for different time steps in the lifetime. For example, the crack may return to its initial deterioration rate τ_0 after a perfect repair or jump a number of deterioration rates back after an imperfect repair. The inspection model is considered within the observation nodes $\boldsymbol{o_t}$. The nodes F_t indicate the probability of failure. In the through-thickness failure criterion, F_t is dependent only on the crack size d_t. The failure subspace $\mathcal{S}_\mathcal{F} \subseteq \mathcal{S}$ is therefore defined based on the discretization scheme of d_t to compute the failure probability.

The initial belief $\mathbf{b_0}(s)$ corresponds to a joint probability distribution of the initial crack size and deterioration rate $P(d_0, \tau_0)$. The belief transitions from each time step t to the next $t+1$ according to the predefined conditional matrix as follows:

$$P(d_{t+1}, \tau_{t+1} \mid o_0, ..., o_t) = \sum_{d_t} \sum_{\tau_t} P(d_{t+1}, \tau_{t+1} \mid d_t, \tau_t) P(d_t, \tau_t \mid o_0, ..., o_t). \tag{2.21}$$

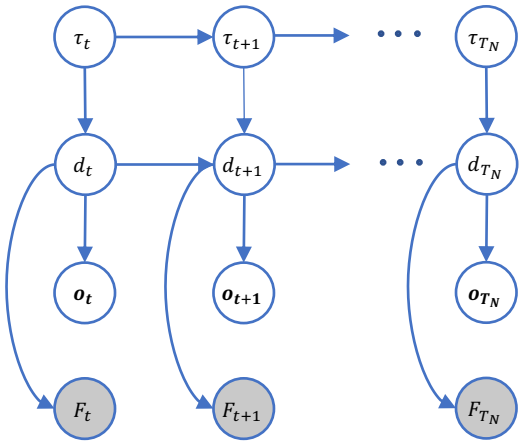

Fig. 2.4. Deterioration rate DBNs adopting a through-thickness criterion. The nodes d_t describe the crack evolution dependent on the deterioration rate τ_t. The nodes $\boldsymbol{o_t}$ represent the imperfect observations (inspections) conditional on d_t, and F_t indicates the probability of a failure event.

When the evidence is available, the estimation of the updated belief can be done through the normalization of:

$$P(d_{t+1}, \tau_{t+1} \mid o_0, ..., o_{t+1}) \propto P(o_{t+1} \mid d_{t+1}) P(d_{t+1}, \tau_{t+1} \mid o_0, ..., o_t). \tag{2.22}$$

2.3.2 Deterioration rate DBNs adopting a FAD criterion

A method to implement a FAD criterion in deterioration rate DBNs is presented here. The DBN model with the FAD criterion is illustrated in Figure 2.5. In this case, the probability of a failure event cannot be obtained only from the nodes d_t since it requires the crack length as well as other time-invariant variables $\sigma_Y, K_{mat}, R_s, R_f$ to evaluate the failure. Alternatively, one can include additional nodes g_{FMt}, denoted here as the limit state variable and computed from Equations (2.12-2.15) in the DBNs to allow the direct estimation of the failure probability F_t. In this I&M planning problem, the observation nodes $\boldsymbol{o_t}$ are conditional on the crack depth and the nodes d_t still need to be tracked. Therefore, the belief space $\mathbf{b_0}(s)$ becomes a joint distribution of the deterioration rate, the crack size and the limit state variable $P(d_t, \tau_t, g_{FMt})$.

The deterioration evolution over the subsequent time steps can be computed through transition and estimation steps. However, the computational complexity significantly increases due to the larger state space size with this failure criterion. The transition of the crack length has been implicitly considered in the DBNs through the nodes g_{FMt}. If the component returns to its initial belief $P(\tau_0, d_0, g_{FM0})$ after the perfect repair, both the crack depth and the crack length consistently return to the initial condition.

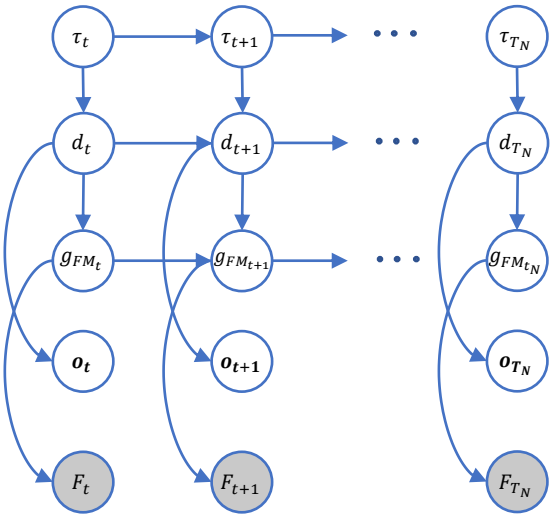

Fig. 2.5. Deterioration rate DBNs adopting a FAD criterion. The nodes d_t describe the crack evolution dependent on the deterioration rate τ_t. The nodes $\boldsymbol{o_t}$ represent the imperfect observations (inspections) conditional on d_t, and F_t indicates the probability of a failure event through the nodes g_{FMt}.

2.4 Policy optimization methods

2.4.1 I&M planning through heuristics

The objective of I&M planning is to identify the optimal strategy which provides the minimum total expected cost $\mathbb{E}[C_T]$. It is theoretically feasible to obtain optimal inspection and maintenance plans by means of the pre-posterior decision theory, however it becomes computationally intractable as the branches of the decision tree exponentially increase with time. One approach to circumvent this problem is to impose predefined decision rules in order to reduce the policies which have to be evaluated. Some of the decision rules that have been frequently applied in risk-based inspection planning of offshore structures are:

1. Inspections are planned either periodically or before an annual failure probability threshold ΔP_F is reached. The optimal interval and optimal annual failure probability threshold are then identified.

2. A perfect repair action is immediately performed if the inspection gives an outcome of crack detection ($g_D(t_{ins}) \leq 0$).

3. After the perfect repair, it is assumed that the component goes back to its initial state thus forming a new decision tree with a lifetime equal to $T_N - t_{ins}$.

I&M strategies are evaluated through Monte Carlo simulations to compute the total expected cost $\mathbb{E}[C_T]$. The expected failure cost $\mathbb{E}[C_F]$ is the sum of annual failure probabilities multiplied by the failure cost C_F. The expected cost of inspection $\mathbb{E}[C_I]$ is the product of inspection cost C_I and the number of inspections. The expected cost of repair $\mathbb{E}[C_R]$ is the product of the repair cost C_R and the number of repairs performed. All the costs are discounted by a factor $\gamma \in [0, 1]$ to take into account the time value of money. The total expected cost $\mathbb{E}[C_T]$ is the averaged sum of the failure, inspection and repair costs over N_{sim} simulations.

$$\mathbb{E}[C_T] = \frac{1}{N_{sim}} \sum_{n=1}^{N_{sim}} \left[\sum_{t=1}^{T_N} C_F \Delta P_F(t) \gamma^t + \sum_{t=t_{I_1}}^{t_{I_n}} C_I \gamma^t + \sum_{t=t_{R_1}}^{t_{R_n}} C_R \gamma^t \right], \tag{2.23}$$

where T_N is the planned lifetime of the structure. $\Delta P_F(t)$ is the annual failure probability for year t. I_n and R_n represent the number of inspections and repairs performed in each simulation.

2.4.2 I&M planning through POMDPs

In the following section, a brief description of partially observable Markov decision processes (POMDPs) is presented with its particular implementation in offshore wind I&M planning problem. A POMDP is a generalization of a Markov decision process (MDP) in which the agent takes probabilistic actions in a stochastic environment and imperfect observations. In

the 7-tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{O}, T, Z, R, \gamma \rangle$ process, the agent takes an action $a \in \mathcal{A}$ thereby transitioning the belief state $\mathbf{b}(s)$ according to the transition model $T(s', a, s) = P(s' \mid s, a)$. The agent then receives an imperfect observation $o \in \mathcal{O}$ with the probability $Z(o, s', a) = P(o \mid s', a)$ and also collects the reward $R(\mathbf{b}, a)$ for taking the action a.

An inspection and maintenance planning problem can be formulated as a POMDP through proper definition of its elements $\langle \mathcal{S}, \mathcal{A}, \mathcal{O}, T, Z, R, \gamma \rangle$. A concise explanation is provided below, and more details can be found in [26, 27, 29].

- States: As already described before, the implementation of DBNs/POMDPs requires efficient and effective discretization of continuous random variables. The first element of POMDP tuple $\mathcal{S}$ can be directly defined from the domain of the discretized intervals. For example, the through-thickness criterion POMDP consists of $|S| = |S_d| \cdot |S_\tau|$ states and that of FAD criterion POMDP is $|S| = |S_d| \cdot |S_\tau| \cdot |S_{g_{FM}}|$. The initial belief $\mathbf{b_0}(s)$ is the joint probability distribution of those random variables at $t = 0$.

- Action-Observation: Several maintenance actions $a \in \mathcal{A}$ can be defined herein such as "perfect-repair", "imperfect-repair" or "do-nothing". Observations $o \in \mathcal{O}$ refer to different types of inspection techniques described in Section 2.2.2. Note that monitoring can also be modeled as an observation through systematic post-processing of continuous data stream into discrete observations.

- Transition probabilities: A transition matrix $T(s', a, s)$ for each maintenance action $a \in \mathcal{A}$ is defined as the probability of the component changing from the state $s \in \mathcal{S}$ to the state $s' \in \mathcal{S}$.

  For the action "do-nothing", the transition matrix $T(s', a_{DN}, s)$ follows the stochastic deterioration process since no maintenance action is performed. Therefore, $P(d_{t+1}, \tau_{t+1} \mid d_t, \tau_t)$ and $P(d_{t+1}, \tau_{t+1}, g_{FMt+1} \mid d_t, \tau_t, g_{FMt})$ become the transition models for POMDPs with different failure criteria.

  The transition model for a "perfect-repair" action $T(s', a_{PR}, s)$ is constructed such that the component holding any belief $\mathbf{b}(s)$ returns to its initial condition $\mathbf{b_0}(s)$ [29]. Despite being briefed to only two actions in this paper, other repair transition matrices can also be defined for different types of maintenance actions [27].

- Observation probabilities: An observation matrix $Z(o, s', a)$ defines the probability of collecting an observation $o \in \mathcal{O}$ for the component being in state $s' \in \mathcal{S}$ after taking the action a. Frequently used ones in offshore I&M planning are "no-inspection", "binary-indication" and "continuous-indication", etc.

- Rewards: After taking an action $a \in \mathcal{A}$ every time step, the agent collects the reward $R(\mathbf{b}, a)$ which is a weighted sum of the belief $\mathbf{b}(s)$ and the state reward $R(s, a)$. One needs to define the state reward $R(s, a)$ for each action-observation combination of RBI planning.

The reward of "do-nothing/no-inspection" is the failure risk computed from the failure cost $-C_F$ assigned to the failure states within $\bar{R}(s, a_{DN-NI})$ and the transition probability as follows:

$$R(s, a_{DN-NI}) = P(s' \mid s, a)\bar{R}(s', a_{DN-NI}) - \bar{R}(s, a_{DN-NI}). \tag{2.24}$$

The reward of "do-nothing/inspection" is one inspection cost $-C_I$ additional to the reward of "Do-nothing/no-inspection" such that:

$$R(s, a_{DN-I}) = R(s, a_{DN-NI}) - C_I. \tag{2.25}$$

The reward of "perfect-repair/no-inspection" is simply equal to the repair cost $-C_R$ for any state:

$$R(s, a_{PR-NI}) = -C_R. \tag{2.26}$$

The objective of I&M planning being to identify the optimal policy which minimizes the total expected cost can be rephrased, within the POMDP framework, as to obtain a sequence of actions that maximizes the total expected reward. In an MDP policy ($\pi : \mathcal{S} \rightarrow \mathcal{A}$), the current state can prescribe which action to be taken. Since the agent cannot fully observe the current state in POMDPs, action decisions are planned based on the belief. A POMDP policy ($\pi : \mathbf{B} \rightarrow \mathcal{A}$) therefore maps a belief $\mathbf{b}$ to the prescribed action and the objective is to identify the optimal policy $\pi^*(\mathbf{b})$ which maximizes the expected sum of the rewards. The value of the optimal policy π^* is described by the value function:

$$V^*(\mathbf{b}) = \max_{a \in \mathcal{A}} \left[\sum_{s \in \mathcal{S}} b(s)R(s, a) + \gamma \sum_{o \in \mathcal{O}} P(o \mid \mathbf{b}, a)V^*(\mathbf{b}') \right]. \tag{2.27}$$

Recently, efficient point-based solvers have been developed which solve high-dimensional state space POMDPs based on a representative set of belief points [22, 24]. In the point-based solvers, the value function in Equation (2.27) is parametrised by a set of α-vectors each of which is associated to an action. For a certain belief $\mathbf{b}(s)$, the optimal action is the one corresponding to the α-vector which maximizes the value function:

$$V^*(\mathbf{b}) = \max_{\alpha \in \Gamma} \sum_{s \in \mathcal{S}} \alpha(s)b(s). \tag{2.28}$$

Since the belief is updated after every action and observation, as in Equations (2.21) and (2.22), the value function is therefore recomputed to choose sequential optimal actions over time.

2.5 Numerical experiments: Application to a tubular joint

With the objectives of implementing the presented I&M planning methods integrated with various failure criteria, as well as exploring the effects of failure criteria, deterioration and inspection models on the identified I&M strategies, a set of numerical experiments are conducted here for the particular case of an offshore tubular joint subjected to fatigue deterioration. Table 2.1 lists all the conducted experiments, classified by the implemented failure criterion and fracture mechanics model.

First, the I&M planning is performed with an inspection model in which the detection threshold $\hat{d}_{s,th}$ (Section 2.2.2) is fixed. The effects of fracture mechanics models and failure criteria on the crack propagation, reliability updating and optimal I&M plan are thoroughly analyzed. The optimal I&M strategies identified by different optimization methods for each case are also compared. Afterwards in Experiment 2, only one combination of 2-D FM model and FAD criterion is considered while the detection threshold $\hat{d}_{s,th}$ of the inspection technique is varied within a range to demonstrate how the I&M policies adapt with different inspection models. Detailed explanation of the deterioration models, inspection and cost models is provided in the following sections.

2.5.1 Deterioration models

SN Model

The fatigue deterioration is first estimated by computing the cumulative fatigue damage following the design recommendations provided by DNV standards [117, 128, 129]. Considering that the tubular joint is located just above the mean waterline, which is an accessible area for inspections, a fatigue design factor FDF of 2 is assigned in this case. Assuming the structural component is designed to the limit for a lifetime of 20 years, the scale parameter of the Weibull stress range distribution is found to be $q = 6.4839$ from Equation (2.1). The variables used in the SN

Table 2.1. List of analyzed cases for RBI planning.

Experiment 1 - Fixed detection threshold			
Option	Case name	Deterioration model	Failure criterion
1	1D-Thick-Fixed	1-D FM	Through-thickness
2	2D-Thick-Fixed	2-D FM	Through-thickness
3	2D-FAD-Fixed	2-D FM	FAD
Experiment 2 - Varied detection threshold			
Option	Case name	Deterioration model	Failure criterion
3	2D-FAD-Varied	2-D FM	FAD

Table 2.2. Variables used in SN model.

Variable	Distribution	Mean (Median)	Std (CoV)	References for Std (CoV)
m_1	Determ.	3		
m_2	Determ.	5		
$*log_{10}(k_1)$	Normal	12.48	0.2	[117, 129]
$*log_{10}(k_2)$	Normal	16.13	0.2	[117, 129]
S_1	Determ.	67.09		
n	Determ.	$3.5 \cdot 10^7$		
q	Normal	6.4839	(0.2)	[129]
h	Determ.	0.8		
Δ	Lognormal	(1)	(0.3)	[2, 117, 129]

Determ. = Deterministic
$*log_{10}(k_1)$ and $log_{10}(k_2)$ are fully correlated.

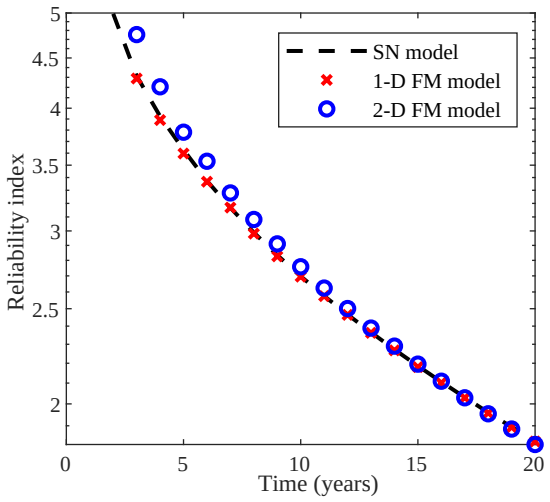

Fig. 2.6. Calibration between SN and FM approaches. The reliability index is the normal inverse cumulative distribution function of the failure probability $\beta(t) = -\Phi^{-1}(P_F(t))$.

approach are listed in Table 2.2. The reliability over the lifetime according to SN-Miner's rule is computed by crude Monte Carlo simulations with one million samples.

FM models

For each considered setting, the initial crack size d_0 and crack growth parameter C_d are calibrated to render a similar reliability in both SN and FM models. Calibration is performed by the least-square fitting of the normalized failure probability. In Option 3 cases, the through-thickness failure criterion is still used for the calibration since it is assumed that the cracks fail when they penetrate the thickness during SN tests. The FAD criterion is only used for reliability analysis and I&M planning. Figure 2.6 shows the goodness-of-fit of the calibrations. The calibration for the 2-D FM model shows some discrepancies in the high reliability region. Yet the probabilities

Table 2.3. Variables used in FM models.

Variable	Option	Distribution	Mean (Median)	Std (CoV)	References for Std (CoV)
d_0	1	Exponential	0.1235		*calibrated
	2,3	Exponential	0.1603		*calibrated
$log(C_d)$	1	Normal	-27.7903	0.3473	*calibrated
	2,3	Normal	-27.6302	0.4599	*calibrated
n	1,2,3	Determ.	$3.5 \cdot 10^7$		
h	1,2,3	Determ.	0.8		
m	1,2,3	Determ.	3		
q	1,2,3	Normal	6.4839	(0.2)	[129]
d_{crit}	1,2	Determ.	16		
Y_d	1	Lognormal	(1)	(0.1)	[2, 96]
d_0/c_0	2,3	Determ.	0.2		
D_b	2,3	Determ.	0.81		
C_d/C_c	2,3	Determ.	1		
R_s	3	Lognormal	(300)	(0.2)	[2]
σ_Y	3	Lognormal	(355)	(0.07)	[2, 130]
K_{mat}	3	3P-W			[2]
R_f	3	Lognormal	(1.7)	(0.18)	[2]

Determ. = Deterministic
3P-W = Three-parameter Weibull distribution

of failure in this region are very small so that they are assumed not to affect the optimal decision. The calibrated parameters together with all other parameters used in FM models are listed in Table 2.3. The normalized resistance parameter R_f for the FAD criterion is taken as recommended by [2].

Incorporation of residual stress

When the failure assessment diagram criterion is used for the case of welded joints, it is necessary to take into account the residual stress as a consequence of weld metal contraction being restrained by the base material [120]. The presence of residual stress in welded joints contributes as secondary stress component in the stress intensity factor such that $K_I = K_I^P + K_I^S$. However, secondary stress does not contribute in the plastic collapse since it has no significant effect on the tensile strength [1].

Realistic estimates of the residual stress are possible by finite element simulations of the considered welded detail. Alternatively, the residual stress can be conservatively assumed to be uniform. In the experiments, the values recommended in [2] are used for lognormal distribution of uniform residual stress R_s, see Table 2.3. The applied primary stress is considered to be fully reversed, i.e., the primary mean stress is zero and therefore, the value of stress amplitude is used when the primary stress intensity factor K_I^P and the load ratio L_r are computed.

Tensile strength and fracture toughness

Material properties are usually considered as uncertain variables due to production variability. The tensile strength of a structural material is often described by a lognormal distribution. Fracture toughness is a quantitative description of material's resistance to fracture failure beyond which the crack propagation becomes unstable. A three-parameter Weibull distribution is proposed to describe the fracture toughness K_{mat} as in the following equation:

$$F_{Kmat}(k) = 1 - exp\left[- \left(\frac{k - K_0}{A_k} \right)^{B_k} \right]. \tag{2.29}$$

The shape parameter B_k is 4 and the recommended value of the threshold parameter K_0 is 20 MPa$\sqrt{m}$ [2]. The scale parameter A_k is computed according to the following equation [1]. The resulting fracture toughness is in MPa$\sqrt{m}$.

$$A_k = \left[11 + 77 \; exp\left(\frac{T - T_0 - T_K}{52} \right) \right] \left(\frac{25}{t} \right)^{0.25} \left[ln\left(\frac{1}{1-p} \right) \right]^{0.25}, \tag{2.30}$$

where T is the temperature at which K_{mat} is to be determined (in °C). T_0 is the temperature for a median toughness of 100 MPa$\sqrt{m}$ in 25 mm thick specimens and calculated as $T_0 = T_{27J} - 18$°C. T_{27J} is the temperature for 27J measured in a standard Charpy V specimen. T_K is the temperature term that describes the scatter in the Charpy versus fracture toughness correlation. For $Std = 15$°C and 90% confidence, T_K is +25°C. t is the thickness of the material for which an estimate of K_{mat} is required (in mm), and p is the probability of K_{mat} being less than estimated and 5% is recommended without experimental evidence [131]. In this paper, the tubular joint is considered to be made of $EN10025 - S355 - JR$ structural steel and the required values for material properties are obtained as follows: $T = 10$°C, $T_{27J} = 20$°C and $\sigma_Y = 355$ MPa [130].

2.5.2 Inspection models

POD curves of different inspection methods frequently used for OWTs are provided in [117]. Eddy current (EC) inspection has become a common inspection method for offshore wind structures as it can be used to detect fatigue cracks without removing coating. The EC inspection in the normal working conditions is used as a reference inspection model in Experiment 1 - Fixed detection threshold. The parameters of signal response method β_0, β_1, σ_ε and $\hat{d}_{s,th}$ as in Equation (2.17), are therefore calibrated to provide an equivalent POD curve as the chosen inspection technique, see Figure 2.7a.

In Experiment 2, the risk-based I&M planning is conducted for a range of detection thresholds. The parameters of the inspection models used in the RBI experiments are shown in Table 2.4. The shape of the POD curve changing with the detection threshold is illustrated in Figure 2.7b.

Table 2.4. Parameters of inspection models.

	β_0	β_1	σ_ε	$\hat{d}_{s,th}$
Experiment 1 - Fixed detection threshold	7.3074	2.092	4.189	5.4898
Experiment 2 - Varied detection threshold	7.3074	2.092	4.189	$0 \leq \hat{d}_{s,th} \leq 10$

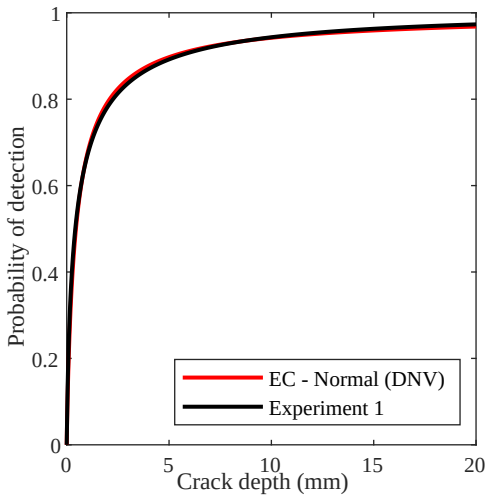

(a) Experiment 1 inspection model, $\hat{d}_{s,th} = 5.4898$.

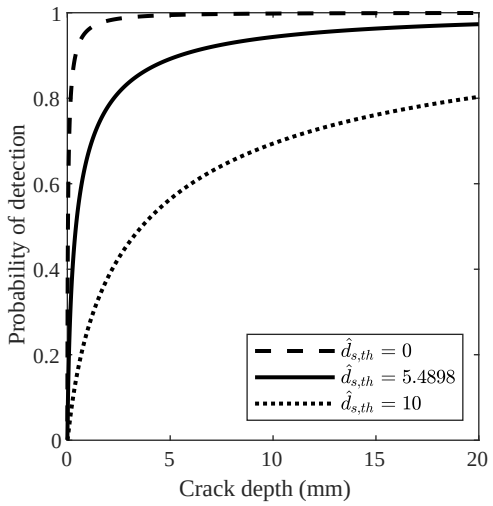

(b) Experiment 2 inspection model, $\hat{d}_{s,th} \in [0, 10]$.

Fig. 2.7. Illustration of inspection models.

2.5.3 Modeling I&M planning in POMDPs

The I&M planning experiments are formulated as POMDPs through the deterioration rate DBNs. As discussed in Section 2.3, the continuous random variables need to be discretized to implement the DBNs and the discretization scheme should be an optimal compromise between the accuracy of the results and the computational performance. Since the discretization is arbitrary and case-specific, several attempts are made on the selection of the number of intervals and boundary values, and the discretization schemes shown in Table 2.5 with the number of states $|S_d| = 40$, $|S_\tau| = 21$, $|S_{g_{FM}}| = 30$, are selected as the relevant ones. Accordingly, the deterioration rate DBNs with the through-thickness criterion has overall 840 states and that of the FAD criterion has 25,200 states with the additional limit state variable g_{FM}. Following the discretization schemes, the initial belief $\mathbf{b_0}$ and transition matrices $T(s', a, s)$ for each case are defined from one million simulations of crack size and FAD assessment points.

Since the existing point-based solvers are set up for the solution of infinite horizon POMDPs and the I&M planning is desired for 20 years (finite horizon), the state space of deterioration rate DBNs is augmented by encoding the time in the state space and adding a terminal state. For the detailed explanation about state augmentation, the reader is referred to [26, 27]. Finally, the state space of finite horizon POMDP with the through-thickness criterion has 9240 states and that of the FAD criterion has 277200 states.

Table 2.5. discretization schemes utilized in the numerical experiments.

Option	Case name	Variable	Interval boundaries
1	1D-Thick-Fixed	a	$0,\ d_{0,mean} : \frac{d_{crit} - d_{0,mean}}{\lvert S_d \rvert - 2} : d_{crit},\ \infty$
		τ	$0 : 1 : 20$
2	2D-Thick-Fixed	a	$0,\ d_{0,mean} : \frac{d_{crit} - d_{0,mean}}{\lvert S_d \rvert - 2} : d_{crit},\ \infty$
		τ	$0 : 1 : 20$
3	2D-FAD-Fixed/Varied	a	$0,\ d_{0,mean} : \frac{d_{crit} - d_{0,mean}}{\lvert S_d \rvert - 2} : d_{crit},\ \infty$
		τ	$0 : 1 : 20$
		g_{FM}	$-\infty,\ 0 : \frac{2}{\lvert S_{g_{FM}} \rvert - 2} : 2,\ \infty$

For all case studies, three action-observation pairs are considered: (1) do-nothing/no-inspection (DN-NI) (2) do-nothing/inspection (DN-I) and (3) perfect-repair/no-inspection (PR-NI). The consequence of a failure event is associated with a cost C_F of 10^6 monetary units. The cost of corrective repair and the risk of system failure conditional on component failure are taken into account in the failure cost C_F of the joint. The cost of inspection C_I independent of the detection threshold is 10^3 monetary units and the repair cost C_R is $1.2 \cdot 10^4$ monetary units. The discount factor $\gamma = 0.94$ is considered. SARSOP point-based solver is used for solving the POMDPs to obtain optimal I&M policies [22].

2.5.4 Results and discussion: Experiment 1 - Fixed detection threshold

Crack growth

As mentioned in the previous sections, both 1-D and 2-D fracture mechanics models are applied to estimate the crack deterioration. In both models, the crack propagation rate is influenced by the Paris law parameters and the stress intensity correction factor. To represent one-dimensional crack growth, Equation (2.9) is used where the stress intensity correction factor Y_d is assigned as a time-invariant random variable, see Table 2.3. In two-dimensional crack growth, the stress intensity correction factors Y_d and Y_c become functions of time-varying crack size and are recomputed at every time step. The geometry functions $Y_{md}, Y_{bd}, Y_{mc}, Y_{bc}$ and stress magnification factors $M_{kmd}, M_{kbd}, M_{kmc}, M_{kbc}$ are evaluated by parametric equations following the procedures of [132] and [117].

A crude Monte Carlo Simulation (MCS) containing 1 million samples was run to estimate the stochastic crack evolution. The comparison of mean crack propagation between 1-D and 2-D FM models can be seen in Figure 2.8a. The crack length is not comparable between the two models as the 1-D model only measures the crack depth. The crack depth grows faster in the 2-D model than in the 1-D model. In fact, this is due to different initial crack depths d_0 and crack growth

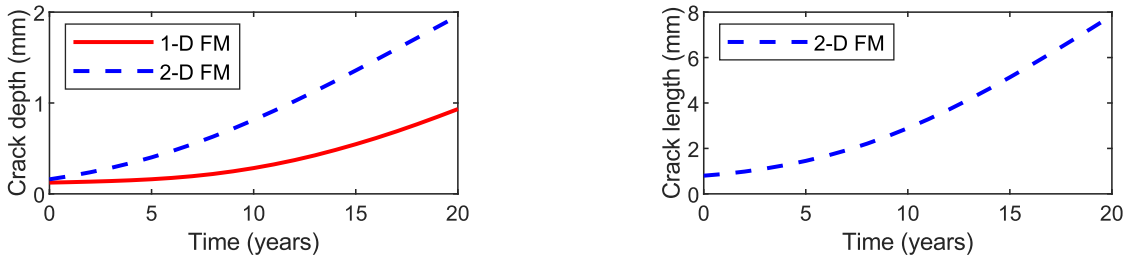

(a) Mean crack depth and crack length.

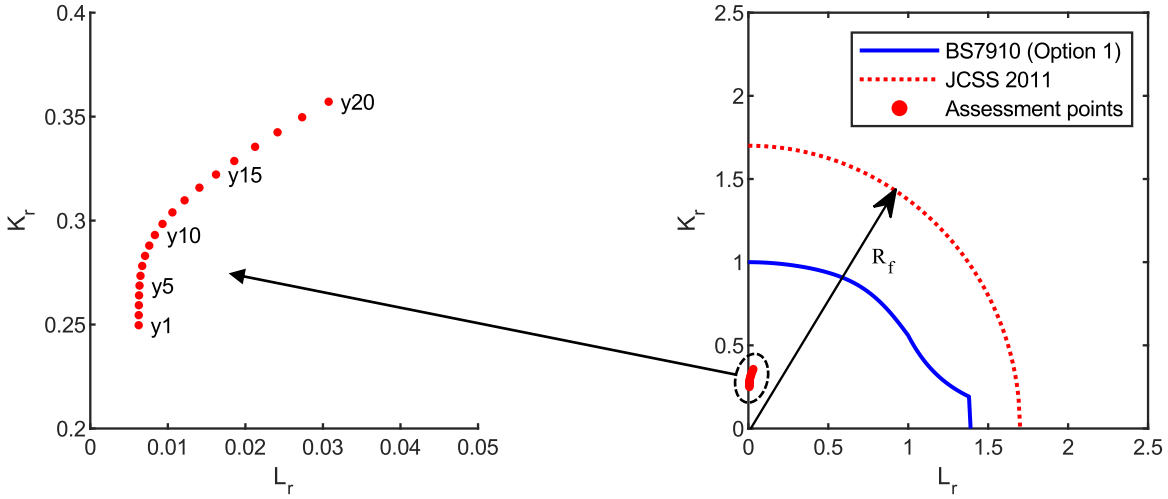

(b) Mean K_r and L_r (2-D FM).

Fig. 2.8. Illustration of crack deterioration.

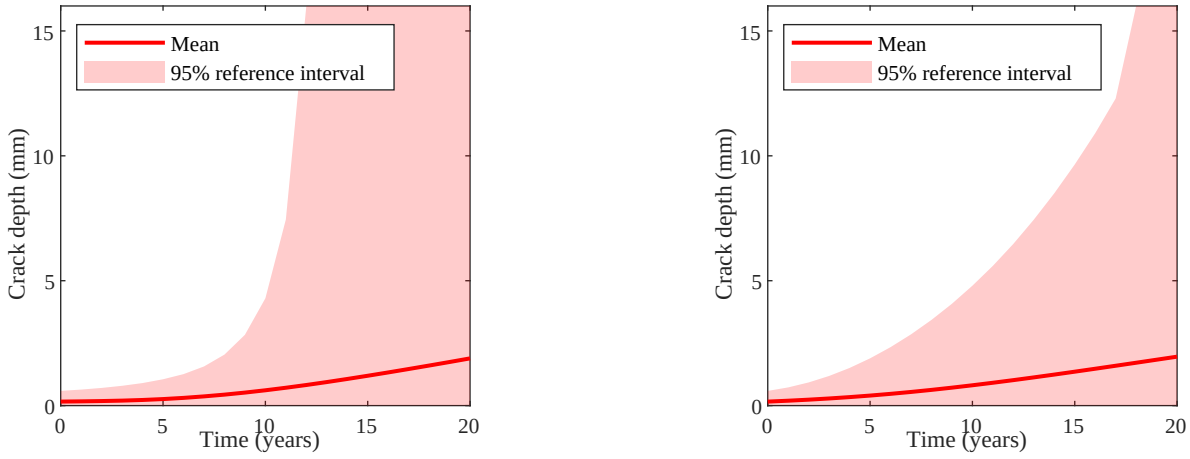

(a) Crack distribution by 1-D FM model.

(b) Crack distribution by 2-D FM model.

Fig. 2.9. Crack propagation by 1-D and 2-D FM models. To illustrate the difference between the two models, the same initial crack size $d_0 \sim Exp(0.1603)$ and crack growth parameter $log(C_d) \sim \mathcal{N}(-27.6302, 0.4599)$ are used.

parameters C_d of the two models since they are calibrated to the same target reliability, see Table 2.3. However, it successively implies different crack propagation by the two models. To examine this, crack propagation is computed by the two models using the same initial crack size $d_0 \sim Exp(0.1603)$ and the crack growth parameter $log(C_d) \sim \mathcal{N}(-27.6302, 0.4599)$. The mean crack depth over the lifetime and 95% reference interval (between 2.5% and 97.5% quantiles) are plotted in Figure 2.9. Note that a cut-off point is considered at $d_{crit} = 16$ mm and all bigger cracks remain at 16 mm in Monte Carlo simulations. It is observed that the two models give different variability in the crack distribution. The variability of the 1-D FM model rapidly increases compared to 2-D model. This effect is important in the structural reliability aspect such that the 1-D FM model with higher model uncertainty gives higher probability of failure than the 2-D FM model. And therefore, when calibrated to have the same reliability, the 2-D FM model results in a higher mean and standard deviation of C_d.

Using the 2-D FM model, the deterioration of the tubular joint can also be described by the fracture ratio K_r and the load ratio L_r computed from Equations (2.13-2.15). The mean values of K_r and L_r are plotted in Figure 2.8b. Before year 10, the FAD assessment point mainly increases in the K_r axis and the load ratio L_r is initially less sensitive to the crack size. With small values of crack size, the tubular joint has a sufficient intact area and therefore it is not subjected to high net section stress σ_{ref}. However, the load ratio starts increasing as the crack depth and crack length rapidly grow after year 10.

Updating reliability

The effect of failure criteria and FM models on the updated failure probability after an inspection is examined here. Assuming an inspection is performed at year $t = 11$ and no crack is detected during the inspection, the reliability is updated for different cases using Equation (2.19) through Monte Carlo simulations. In the deterioration rate DBNs, the updated failure probability is the probability of being in the failure states after performing the transition and estimation steps. As shown in Figure 2.10, consistent P_F values are obtained by the MCS and the DBNs, which verifies the proper discretization of the state space variables. On the other hand, the updated failure probabilities after the inspection are different among the analyzed cases. The 2D-FAD-Fixed case gives the smallest failure probabilities in all years. This is expected simply due to the assumption of the capacity to hold the through-thickness cracks when the FAD criterion is used.

The effect of fracture mechanics models on reliability updating can be analyzed through comparison of 1D-Thick-Fixed and 2D-Thick-Fixed cases. As discussed before, the two FM models provide different initial crack sizes d_0 and crack growth parameters C_d when calibrated to the SN failure probabilities. Consequently, the crack depth belongs to different probability distributions with respect to each FM model, as illustrated with the crack size histograms in Figure 2.11a. Firstly, the variation in prior distributions influences the updated failure probability given no crack detection. However, this is not obvious just after the no-detection

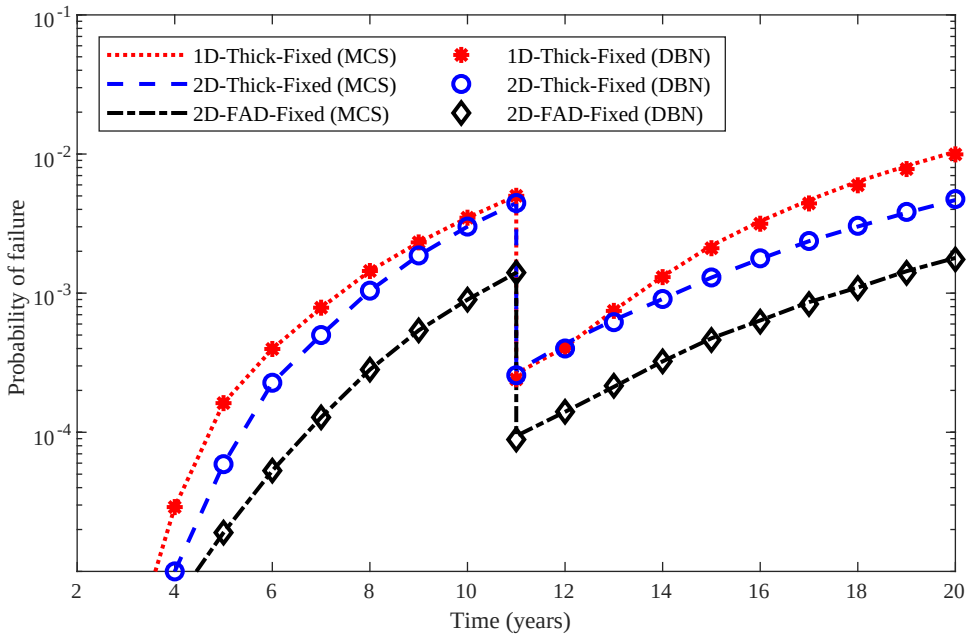

Fig. 2.10. The updated cumulative failure probability given no-detection in the inspection at year 11.

event, in Figure 2.11b. Since the probability of no-detection and probability of crack distribution are very low in the failure bin, i.e., $d \geq d_{crit}$, the difference in the updated P_F is extremely small when normalized by the overall no-detection probability according to Equation (2.19). As the posterior distribution further propagates accumulating and increasing the probability in the failure bin, the difference in the updated P_F can then be clearly observed as in Figure 2.11c and 2.11d. The more-detailed 2-D FM finally achieves higher reliability at the end of the lifetime. And vice versa, the decision-maker may use simple models with higher uncertainty, e.g., 1-D FM but may take a higher risk than using more precise ones or perform more inspections to remain at the same reliability and risk.

Secondly, the overall probability of detection with the same inspection model is different between the two cases. The 2-D FM model with its higher crack growth parameters results in a higher probability of detection. In heuristic-based inspection planning, maintenance decision rules are often prescribed based on inspection outcomes, e.g., repair is performed if a crack is detected. In this case, using the 2-D FM will result in higher maintenance costs than the 1-D FM for the same inspection strategy.

Optimal I&M strategies

I&M planning is performed through traditional heuristic-based methods as well as through the formulation as POMDPs. The SARSOP point-based POMDP solver is used for the computation of the optimal I&M policies [22]. Two sets of heuristic decision rules - equidistant inspections (EQ-INS) and inspections planned before an annual failure probability threshold is

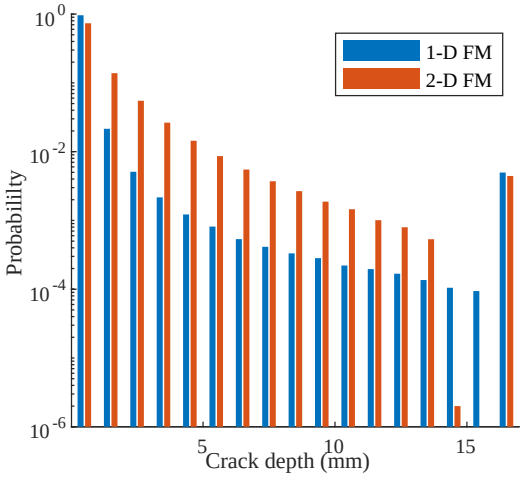

(a) Prior distribution of crack depth (t=11).

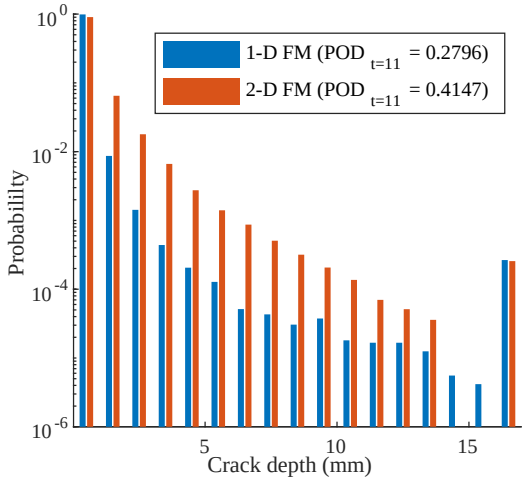

(b) Posterior distribution of crack depth (t=11).

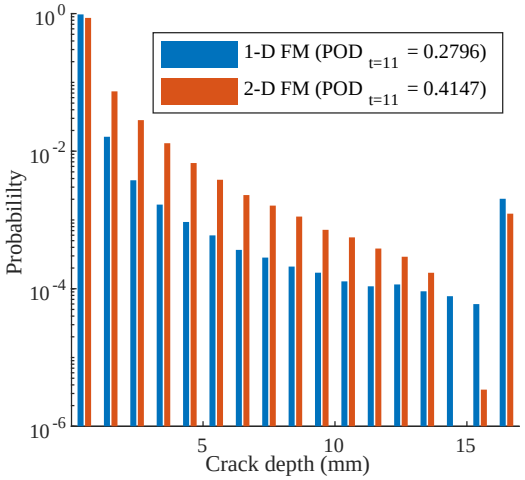

(c) Posterior distribution of crack depth (t=15).

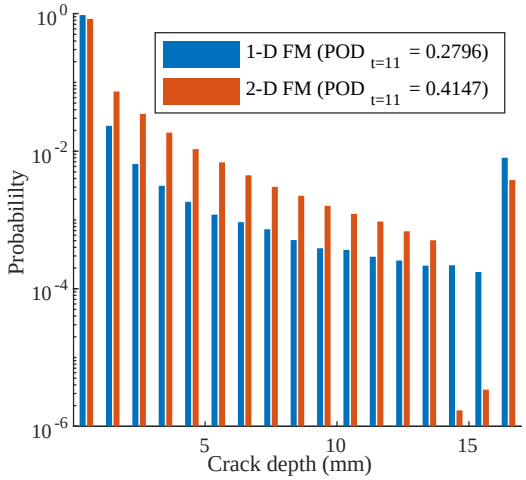

(d) Posterior distribution of crack depth (t=19).

Fig. 2.11. Propagation of crack distribution. (a) The prior crack distribution at year $t = 11$ is estimated by the two FM models. (b) The posterior P_F at the year of inspection is almost equal for 1-D and 2-D models whereas the overall detection probability is different. (c, d) The updated P_F w.r.t the two models becomes different as the posterior crack distribution propagates.

exceeded (THR-INS) - have been evaluated in the simulation environment through DBNs. If the inspection indicates the presence of crack, a perfect repair is immediately performed and after, the component goes back to its initial condition. The identified optimal I&M strategies from heuristics and POMDPs are evaluated in a simulation environment. The resulting total expected costs $\mathbb{E}[C_T]$ as well as the numerical confidence intervals over 10^5 simulations are reported in Table 2.6.

In all combinations of deterioration models and failure criteria, POMDP policies outperform traditional heuristics with significant differences in the total expected cost. Comparing among different options, 2D-FAD-Fixed case results in less expected cost than the other two options

Table 2.6. Experiment 1 - Fixed detection threshold: Comparison of the total expected cost.

	$\mathbb{E}[C_T]$	95% C.I
Option 1: 1D-Thick-Fixed		
Finite horizon POMDP - SARSOP	6267.20	±20.49
Heuristic EQ-INS ($\Delta_{ins} = 7$)	6892.72	±22.462
Heuristic THR-INS ($\Delta P_{Fth} = 9 \cdot 10^{-4}$)	6775.92	±20.045
Option 2: 2D-Thick-Fixed		
Finite horizon POMDP - SARSOP	5387.88	±22.43
Heuristic EQ-INS ($\Delta_{ins} = 11$)	6066.32	±15.681
Heuristic THR-INS ($\Delta P_{Fth} = 1.2 \cdot 10^{-3}$)	6066.32	±15.681
Option 3: 2D-FAD-Fixed		
Finite horizon POMDP - SARSOP	3599.45	±16.39
Heuristic EQ-INS ($\Delta_{ins} = 11$)	4118.25	±17.462
Heuristic THR-INS ($\Delta P_{Fth} = 5 \cdot 10^{-4}$)	4118.25	±17.462

which rely on the through-thickness criterion. When the FAD criterion is used, it is assumed that the through-thickness cracks can grow further in the length until the critical value of the stress intensity factor is reached. It is also worth mentioning that the fracture toughness of the material considered for the tubular joint is high enough so that the component does not fail before the crack reaches the thickness and can hold the through-thickness crack. Therefore, the failure probabilities over the lifetime with the FAD criterion are smaller than the other cases, as already discussed before. It results in a significant reduction of failure risk as well as less observations and maintenance actions can be generally expected. Random realizations are presented to visualize the policies prescribed by different approaches. Figure 2.12a, 2.12c and 2.12e represent the finite horizon POMDP policy realizations and Figure 2.12b, 2.12d and 2.12f show the realizations of the equidistant heuristic. As manifested in the POMDP policies, only one inspection is required in the 2D-FAD-Fixed case if the first inspection outcome is no-detection thanks to the low failure risk. Contrarily, more inspections are conducted in a number of years in the cases in which the through-thickness criterion is used.

In the heuristic-based policies, an immediate repair action is prescribed after detection, but the POMDP policies may opt to perform subsequent inspections in case of detection. A repair action is planned only if the subsequent inspections also give detection outcomes. In Figure 2.12a, POMDP policies plan one more inspection after a detection event and a repair is performed as the second inspection also gives detection. However, POMDP policies also consider more subsequent inspections when the deterioration model has lower uncertainty. For the 2D-Thick-Fixed case, POMDP plans up to three consecutive inspections as in Figure 2.12c. Since the last inspection declares no-detection, no maintenance action is taken. Therefore, POMDP solutions can provide adaptive policies for different scenarios such that, as in this example, it is plausible to take advantage of the more precise 2-D model to avoid an expensive repair.

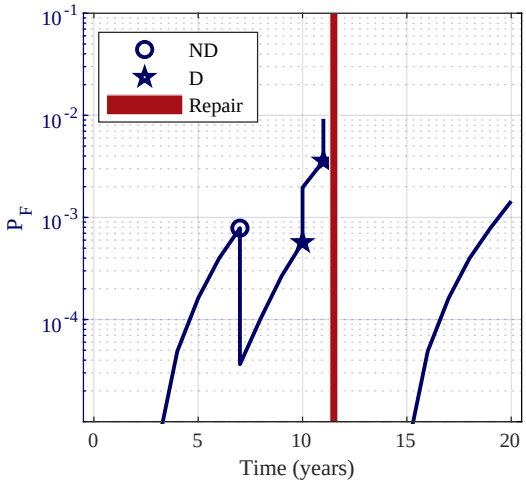

(a) 1D-Thick-Fixed (FH-POMDP).

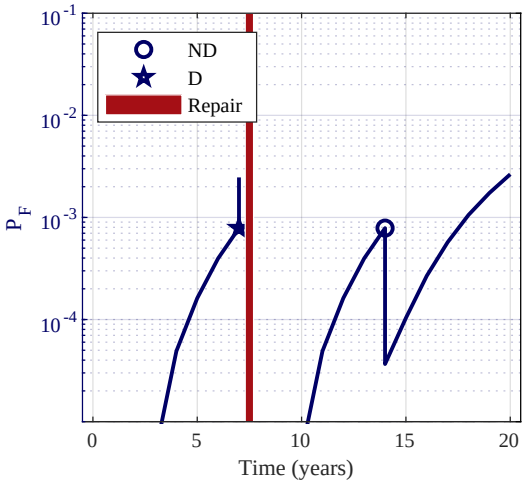

(b) 1D-Thick-Fixed (Heuristic EQ-INS).

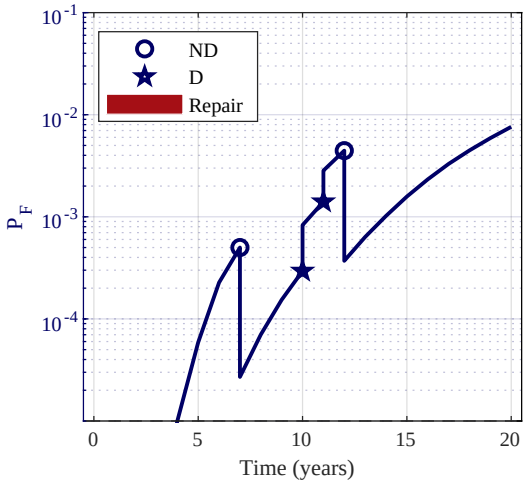

(c) 2D-Thick-Fixed (FH-POMDP).

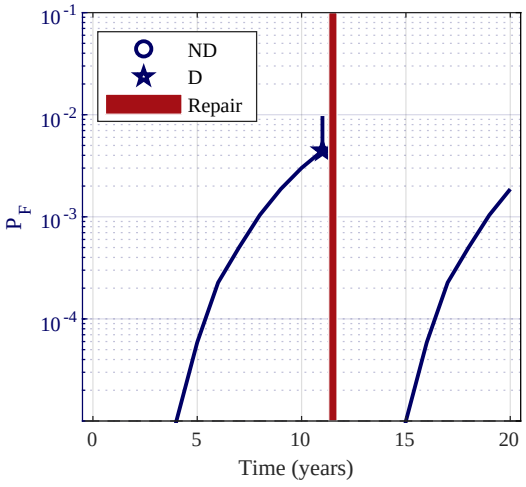

(d) 2D-Thick-Fixed (Heuristic EQ-INS).

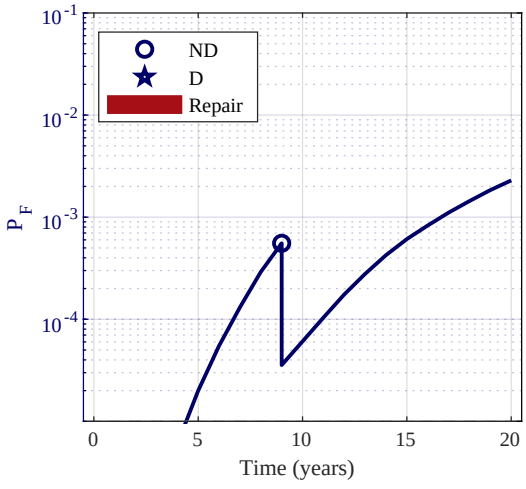

(e) 2D-FAD-Fixed (FH-POMDP).

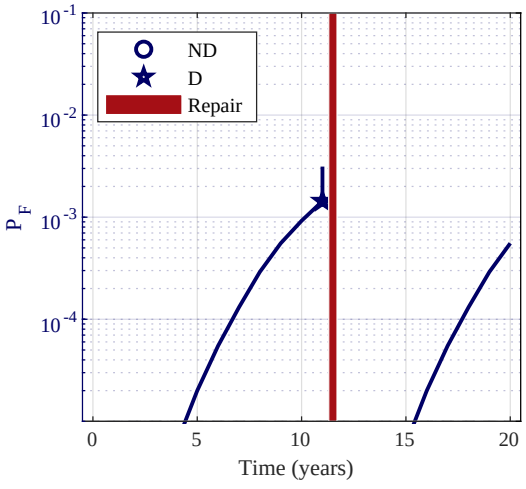

(f) 2D-FAD-Fixed (Heuristic EQ-INS).

Fig. 2.12. Policy realizations of Experiment 1 - Fixed detection threshold. Inspection outcomes are represented by a circle (for no-detection) or a five-pointed star (for detection). A red bar denotes that a perfect repair is performed.

In contrast, traditional heuristic approaches which only follow the predefined decision rules are not able to capture such aspects. Utilizing the 2-D FM instead of the 1-D FM affects the optimal heuristic policies such that the number of inspections is reduced, which opposes the pattern in the POMDP policies. In the equidistant inspection approach, two inspections ($\Delta_{ins} = 7$) are performed in the 1D-Thick-Fixed whereas only one inspection ($\Delta_{ins} = 11$) is conducted in the 2D-Thick-Fixed. As discussed before, using the 2-D FM model is a matter of reducing the risk but on the other hand increasing the maintenance cost. And in this case, the increased maintenance cost is more than the reduced risk, therefore one inspection has been reduced to adjust the optimal policy.

2.5.5 Results and discussion: Experiment 2 - Varied detection threshold

Optimal I&M strategies

Hereafter, only one combination with 2-D FM model and FAD criterion is considered. The optimization is performed only through POMDPs as it has been demonstrated in the previous experiment that POMDP outperforms traditional heuristic approaches for any combination of deterioration model and failure criteria. In this experiment, the RBI planning is repeatedly performed for several values of detection thresholds $0 \leq \hat{d}_{s,th} \leq 10$. The transition models and the rewards models of the POMDP remain the same as the 2D-FAD-Fixed case from the previous experiment. Only the observation model $Z(o, s', s)$ is modified for each threshold value.

The breakdown of the total expected cost evaluated over 10^5 simulations are reported in Figure 2.13. In general, moderate to high detection thresholds $5 \leq \hat{d}_{s,th} \leq 10$ provide the lowest expected costs. The policy realizations of some representative cases are also presented. Figure

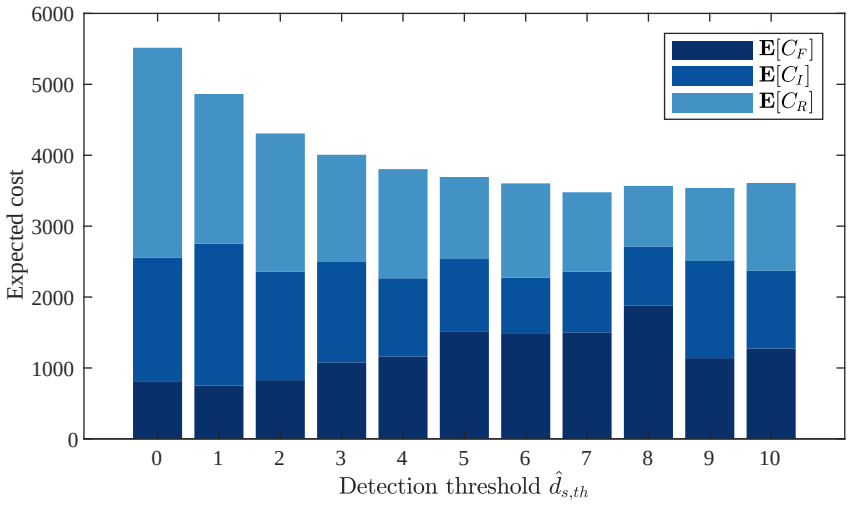

Fig. 2.13. Experiment 2 - Varied detection threshold: Comparison and breakdown of the total expected cost.

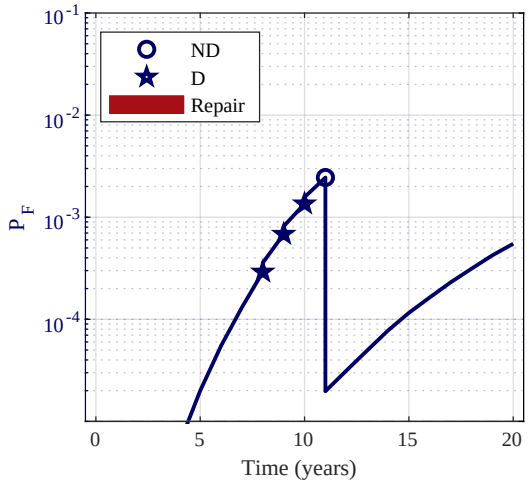

(a) 2D-FAD-$\hat{d}_{s,th} = 0$ (FH-POMDP).

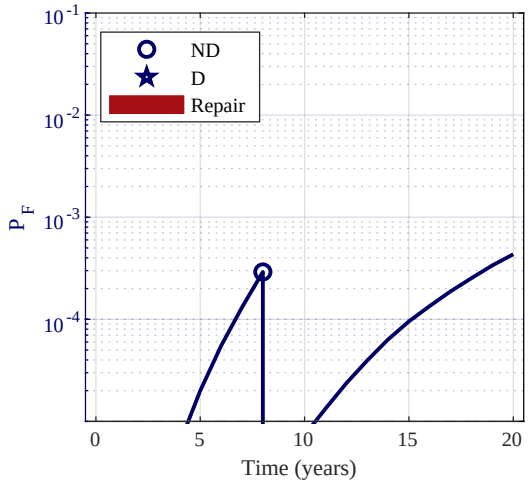

(b) 2D-FAD-$\hat{d}_{s,th} = 0$ (FH-POMDP).

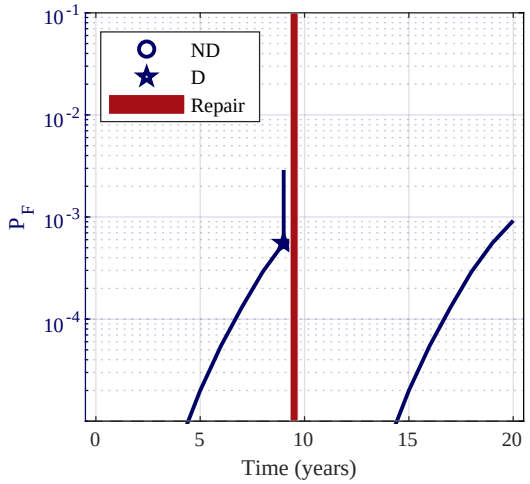

(c) 2D-FAD-$\hat{d}_{s,th} = 10$ (FH-POMDP).

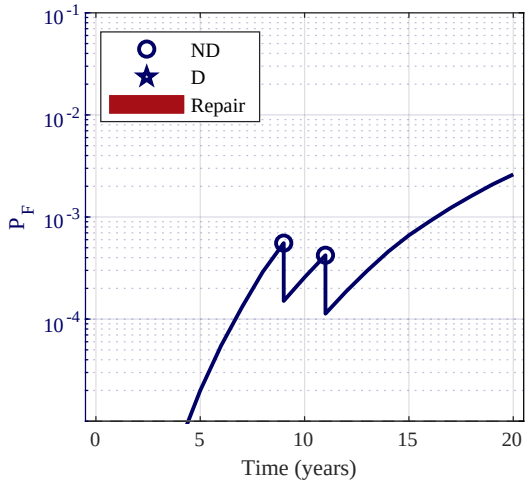

(d) 2D-FAD-$\hat{d}_{s,th} = 10$ (FH-POMDP).

Fig. 2.14. Policy realizations of Experiment 2 - Varied detection threshold. Inspection outcomes are represented by a circle (for no-detection) or a five-pointed star (for detection). A red bar denotes that a perfect repair is taken.

2.14a and 2.14b show random realizations of inspection planning with low detection threshold $\hat{d}_{s,th} = 0$ and Figure 2.14c and 2.14d show the realizations with high threshold $\hat{d}_{s,th} = 10$. An interesting pattern of inspection planning can be discovered in the POMDP prescribed policies. When the inspection model has a lower threshold, POMDP tends to plan several subsequent inspections in case of detection. The shape of the *POD* curve plateaus earlier with lower detection threshold, e.g., *POD* curve of $\hat{d}_{s,th} = 0$ becomes flat at around $d = 5$ mm in Figure 2.7b. If a detection outcome is obtained, one may not simply presume bigger crack size and/or higher probability failure, noting that even the small cracks which do not cause failure may also be detected. Therefore, the probability of failure is only slightly increased and successive inspections are planned instead of taking an expensive repair action, see Figure 2.14a. In the case of no-detection from the low threshold inspection, a small crack can be assured since

bigger cracks $d > 5$ mm have almost zero probability of no-detection in the POD curve. The probability of failure therefore drops drastically, see Figure 2.14b.

On the other hand, the POD curve with a high detection threshold slants up with the crack size, see $\hat{d}_{s,th} = 10$ in Figure 2.7b. Bigger cracks are more likely to be detected than smaller ones. If this inspection model gives a detection outcome, a severe crack can be expected with a higher confidence, thus resulting in a high jump of failure probability as in Figure 2.14c. The repair action is also performed immediately after one detection. Nevertheless, detection events are less frequent in this case and the expected repair cost is still remarkably lower compared to the low threshold inspection.

2.6 Conclusions

In this paper, the effects of failure criteria, deterioration and inspection models on the optimal inspection and maintenance (I&M) strategies are examined. Two-dimensional fracture mechanics model and failure assessment diagram (FAD) criterion have been successfully integrated with dynamic Bayesian networks (DBNs), therefore allowing the formulation and optimization of I&M planning via partially observable Markov decision processes (POMDPs). Various deterioration, inspection and failure criteria settings were tested for the optimal management of a structural detail subjected to fatigue deterioration, revealing the following findings:

- The failure assessment diagram might be preferred as the failure criterion to model redundant structures with capacity to sustain through-thickness cracks since it offers significant savings in the total expected cost, especially through advanced optimization methods such as POMDPs. However, one needs the knowledge of material properties, compulsorily yield strength and fracture toughness to implement the FAD criterion.

- 2-D fracture mechanics models are more robust than 1-D models since the effects of time-dependent crack size, geometry of the structure and welded detail are inherently considered in the stress intensity correction factors Y_d and Y_c. The shortcoming is that it requires FEM analysis or the use of parametric equations, both of which demand computational resources.

- The observation model, specified often through probability of detection (POD) curves, can be adjusted by varying the detection threshold of the inspection technique. Generally, very low detection thresholds are not recommended due to its flat and high POD curve, causing frequent detections, inspections and/or repairs.

Throughout the experiments, it is demonstrated that I&M polices provided by finite horizon POMDPs outperform heuristic-based polices for any combination of deterioration models and failure criteria. POMDPs also reveal adaptability in the policy patterns depending on the models specified in the I&M planning. The main limitation associated with implementing the 2-D

fracture mechanics model and FAD criterion in discrete DBNs/POMDPs is the high-dimensional state space. Computational intractability becomes a constraint to apply such complex models and failure criterion for the case of multiple components or longer horizon lengths. To overcome this curse of dimensionality, further research efforts are suggested towards the integration of the FAD criterion with POMDP-based deep reinforcement learning (DRL) approaches. The capability of POMDP-based DRL approaches to efficiently provide optimal I&M strategies for large state space problems has been demonstrated in [40].

Authorship contribution statement

Hlaing, N.: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing - Original draft, Writing - Review & editing, Visualization.
Morato, P. G.: Conceptualization, Software, Validation, Formal analysis, Writing - Review & editing, Supervision. **Nielsen, J. S.:** Validation, Formal analysis, Writing - Review & editing. **Amirafshari, P.:** Formal analysis, Writing - Review & editing. **Kolios, A.:** Supervision, Project administration. **Rigo, P.:** Supervision, Project administration, Funding acquisition.

Acknowledgments

This research has been developed at the University of Liège, Belgium, in collaboration with the University of Strathclyde, UK, and Aalborg University, Denmark. The financial support provided by the Belgian Energy Transition Fund (FPS Economy) through the projects PhairywinD and MaxWind is greatly appreciated. Dr. Morato would further like to acknowledge the support granted by the National Fund for Scientific Research in Belgium F.R.I.A. - F.N.R.S.

Chapter 3

Interpretation of offshore wind management policies identified via partially observable Markov decision processes

---

Paper Hlaing, N., Morato, P. G., Papakonstantinou, K. G., Andriotis, C. P. and Rigo, P. (2022). Interpretation of offshore wind management policies identified via partially observable Markov decision processes. In *The 18th European Academy of Wind Energy (EAWE) PhD seminar*.

3.1 Introduction

The installation of offshore wind turbines, profiting from available abundant and stable wind resources, has been steadily increasing in the last decade, yet preserving offshore wind structures in a good condition throughout their lifetime still remains a challenge. Structural components are exposed to deterioration mechanisms (e.g., fatigue, corrosion, among others), and far offshore, inspection and maintenance (I&M) operations can be complex and expensive. Hence the need for efficient optimal I&M planning methods has been increased in order to control the risk of structural failures by timely allocating inspection and maintenance interventions.

Identifying optimal I&M policies demands the solution of a complex sequential decision-making problem under uncertainty and imperfect information. Whereas time-, condition-, or heuristic-based strategies are conventionally followed in the offshore wind industry in order to alleviate the aforementioned computational difficulties, the resulting policies statically select inspection and maintenance actions and/or consist in predefined heuristic decision rules, e.g., equidistant inspections, repairs after detection inspection outcomes, which are optimized by exploring a subset out of the vast policy space. Instead, optimal management strategies can be identified via partially observable Markov decision processes (POMDPs), relying on mathematical principles conceived for planning under uncertainty [26]. POMDP policies, efficiently computed through point-based solvers, provide optimal adaptive I&M strategies that ultimately result in

substantial cost benefits compared to their state-of-the-art counterparts [29], also demonstrated in offshore wind inspection and maintenance planning settings [133].

Even if recently reported results demonstrate the benefits of implementing POMDP-based adaptive policies for the management of offshore wind assets, the interpretation and execution of POMDP-based strategies by decision-makers (e.g., designers, operators, etc.) accustomed to calendar- and/or condition-based conventional I&M approaches might be initially challenging. In this work, we analyze and interpret POMDP-based policies with the objective of accelerating their practical implementation by offshore wind asset management decision-makers. Also, we showcase the inherent flexibility and adaptability properties offered by POMDP-based policies in a typical offshore wind inspection and maintenance planning setting, in which a decision-maker opts for an action other than the one suggested in the optimal POMDP policy.

3.2 Optimal I&M planning for offshore wind structures through POMDPs

A Markov decision process (MDP) is a 5-tuple $\langle S, A, T, R, \gamma \rangle$ controlled stochastic process for optimal planning under uncertainty and perfect information. At every decision step, the agent observes state $s \in S$ and takes an action $a \in A$, then the state randomly transitions to state $s' \in S$ according to a stochastic transitional model $T(s, a, s') = P(s'|s, a)$, and finally the agent receives a reward $R(s, a)$. An MDP policy ($\pi : S \rightarrow A$) prescribes actions as a function of the current state, with the main objective of identifying the optimal policy $\pi^*(s)$, resulting in the maximum expected rewards (or minimum expected cost).

A POMDP is a generalization of an MDP in which the agent only receives partial information about the current state. In this case, the agent reasons according to the current belief $\mathbf{b}$, i.e., a probability distribution over states. A POMDP is defined as a 7-tuple $\langle S, A, O, T, Z, R, \gamma \rangle$ controlled stochastic process. While a POMDP transitional model corresponds to the underlying MDP, an observation model is additionally defined by specifying the probability $Z(o, s', a) = P(o \mid s', a)$ of collecting observation $o \in O$ after taking action a. After taking action a and collecting observation o, the belief $\mathbf{b}$ is updated via Bayes' rule:

$$b(s') \propto P(o \mid s', a) \sum_{s \in S} P(s' \mid s, a) b(s). \tag{3.1}$$

Since beliefs are dynamically updated, the current belief, $\mathbf{b}$ is a sufficient statistic of the past taken actions and collected observations. A POMDP policy therefore maps the current belief $\mathbf{b}$ to the action. As for an MDP, the goal is to identify the optimal policy $\pi^*(\mathbf{b})$ leading to the maximum expected reward.

The decision-making problem corresponding to the optimal inspection and maintenance planning for offshore wind structures can be adequately formulated as a POMDP, in which the agent reasons in a stochastic environment (i.e., probabilistic deterioration model) and under

imperfect information (i.e., measurement uncertainty associated with inspection techniques). Once the optimal POMDP policy $\pi^*(\mathbf{b})$ is identified, the decision-maker (e.g., operator, designer, etc.) selects inspection and/or maintenance actions according to the current belief state. As opposed to static decision rules, e.g., calendar- or condition-based maintenance approaches, POMDP policies are inherently adaptive since beliefs are dynamically updated, thus resulting in substantial cost benefits.

3.2.1 Solving POMDPs

The exact solution of a POMDP demands the identification of optimal actions for each belief state, which as mentioned before, is a continuous probability distribution over states, thus rendering the problem computationally challenging. Whereas value iteration algorithms or grid-based interpolation techniques might work well for solving very low-dimensional state space POMDPs, their application to higher dimensional state space POMDPs remains limited, also due to computational tractability problems. However, the recently developed point-based solvers, by executing Bellman backups only for a set of reachable belief points, have enabled the solution of medium to high dimensional state space POMDPs [35]. Since the value function is generally piece-wise linear and convex, it can be parametrized through a finite set Γ of α-vectors, each of them associated with an action [29]. At a certain belief state, the optimal action is, therefore, indicated by the α-vector that maximizes the value function. Point-based solvers are usually developed for the solution of infinite horizon settings, yet practical applications normally correspond to finite horizon problems, e.g., the operational lifetime of offshore wind structural components is often considered as 20 or 30 years. In that case, the infinite horizon POMDP can be transformed into a finite horizon POMDP through state augmentation techniques [26, 29].

3.3 Interpretation of POMDP-based management policies

With the objective of facilitating the interpretation of POMDP-based offshore wind management policies, we conduct hereafter an I&M planning case study for a fatigue-sensitive offshore wind structural component, inspired by [133]. The I&M decision problem is formulated as a POMDP, adequately defining the elements of the POMDP tuple, as follows:

- States: The structural component deterioration states correspond to the discretized fatigue crack size. In this study, the crack size is discretized into 40 deterioration states, with the last one indicating a failure state.
- Actions: Three action-observation combinations are considered, (i) Do-nothing/No-inspection (DN-NI), (ii) Do-nothing/Inspection (DN-I), and (iii) Perfect-repair/No-inspection (PR-NI).

- Observations: Inspections provide binary indications, resulting in either 'crack detection' or 'no crack detection'. If an inspection is not performed, no additional information is collected.
- Transition probabilities: The transitional model associated with a Do-nothing (DN) action is estimated through crack propagation Monte Carlo simulations, where the crack growth is computed according to Paris law. If a Perfect-repair (PR) action is undertaken, the structural component deterioration transitions to its initial belief condition state.
- Observation probabilities: The observation model is defined according to the detection probability curve that corresponds to eddy current inspection techniques [117].
- Rewards: At every decision step, the agent collects a reward, $R(\mathbf{b}, a)$, which is a weighted sum of the belief probability $b(s)$ and state reward $R(s, a)$. A penalization of one million monetary units is charged at the last state, i.e., failure condition, whereas 1,000 and 10,000 thousand monetary units are assigned as inspection and repair costs, respectively.

In this case study, the structural component lifetime is defined as 20 years and the corresponding finite horizon POMDP is computed via SARSOP point-based solver [22]. The resulting optimal policies are parametrized by a set of α-vectors, and as mentioned before, each α-vector is associated with a specific action. At each decision point, the decision-maker selects the α-vector (and corresponding action) that maximizes the value function $V^*(\mathbf{b})$ (minimizes the total expected cost):

$$V^*(\mathbf{b}) = \max_{\alpha \in \Gamma} \sum_{s \in S} b(s)\alpha(s). \tag{3.2}$$

The expected total cost associated with each α-vector can be simply computed as the weighted sum of the expected total cost corresponding to a specific deterioration state $\alpha(s)$ and the probability of being in that state $b(s)$. Figure 3.1a illustrates the expected total cost resulting from three α-vectors, indicating both the corresponding deterioration state values along with the representation of the initial belief $\mathbf{b}_0$. The key observation is that the actions recommended in POMDP-based policies are selected according to the current belief state $\mathbf{b}$, which is dynamically updated after each taken action and collected observation, as mentioned in Section 3.2. The expected costs associated with all α-vectors available at the initial decision step (i.e., $\mathbf{b}_0$) are additionally represented in Figure 3.1b. Logically, the optimal decision at this point is DN-NI, and its corresponding value function indicates the total expected cost $\mathbb{E}[C_T]$ for the considered 20-year decision horizon.

3.3.1 What if the optimal policy is not strictly followed?

In this study, we investigate the effect of selecting an alternative action rather than the optimal one suggested in the POMDP policy. We consider that the optimal POMDP policy is followed up to year 7, and at that point, the decision-maker is evaluating the potential economic implications

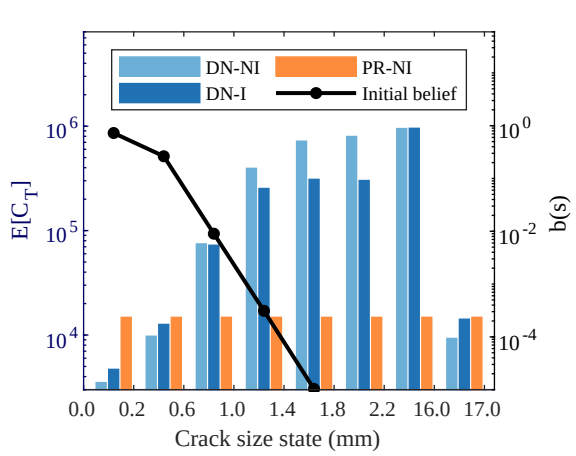

(a) Initial belief (black line) and state values of three α-vectors (colored bars).

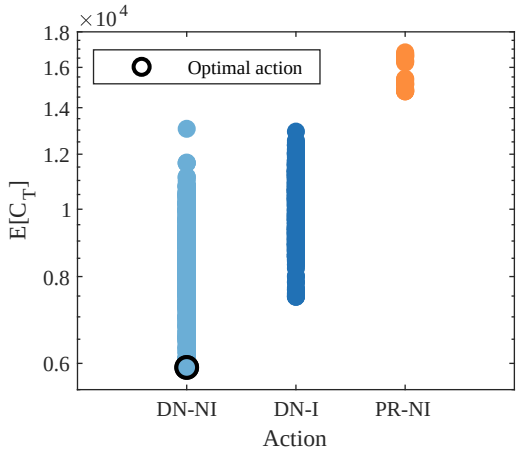

(b) Expected total cost associated with each α-vector (and corresponding action).

Fig. 3.1. Initial probability distribution over deterioration states (i.e., initial belief, $\mathbf{b}_0$) and expected total cost resulting from each α-vector at the initial decision step.

of avoiding a perfect repair maintenance intervention, which is the action suggested in the optimal POMDP policy, as showcased in Figure 3.2b. Previously, two crack detection inspection outcomes were reported at years 6 and 7, thus indicating a high structural failure risk, which could be effectively mitigated by conducting a repair action (Figure 3.3a). In that case, the structural condition will be restored, and the updated belief will transition to its initial deterioration condition, $\mathbf{b}_0$, as illustrated in Figure 3.2a with green markers. If the decision-maker opts, however, for an alternative action at year 8, the expected total cost and the regret, i.e. the extra cost associated with potentially suboptimal actions, can be straightforwardly computed through a Bellman backup operation, as:

$$V(\mathbf{b}_7) = \sum_{s \in S} b_7(s) R(s,a) + \gamma V(\mathbf{b}'_7), \tag{3.3}$$

where $\mathbf{b}$ and $\mathbf{b}'$ correspond to the current and updated beliefs, respectively, and $R(s,a)$ stands for the reward associated with the action taken. Specifically, the potential alternative actions at this decision point are:

- Do-nothing/No-inspection (DN-NI), in which the fatigue deterioration will naturally progress according to the defined transition model, as illustrated with grey markers in Figure 3.2a.
- Do-nothing/Inspection (DN-I), which can result in either a crack detection or no crack detection outcome. The corresponding updated beliefs are plotted in Figure 3.2a with red and blue markers, respectively. Since two inspection outcomes can be collected, in this

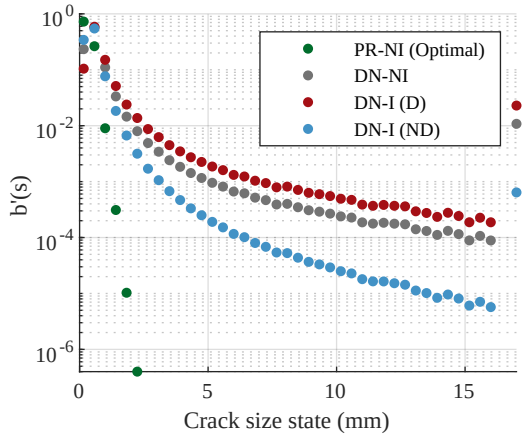

(a) Updated beliefs for all action-observation combinations.

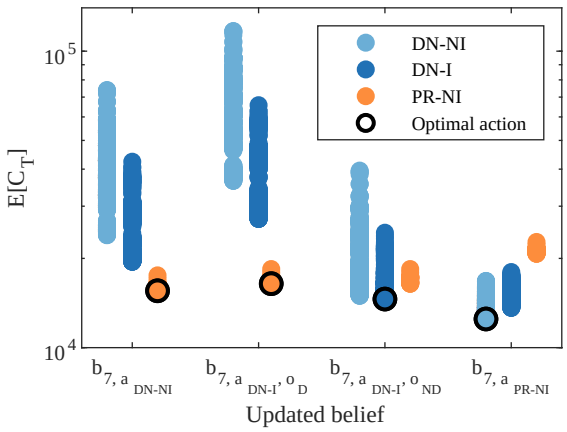

(b) Total expected cost for all action-observation combinations.

Fig. 3.2. Updated beliefs and corresponding total expected cost for all action-observation combinations at year 8.

case, the expected total cost estimated can be computed as:

$$V(\mathbf{b}_7) = \sum_{s \in S} b_7(s) R(s, a_{DN-I}) + \gamma \left[\sum_{o \in O} p(o \mid \mathbf{b'}_{7,a_{DN-I}}) \cdot V(\mathbf{b'}_{7,a_{DN-I},o}) \right], \tag{3.4}$$

where $p(o \mid \mathbf{b}')$ represents the probability associated with each inspection outcome.

Gathering all action-observation combinations, Figure 3.2b illustrates the expected cost associated with each action. The suggested action and expected total cost corresponding to each updated belief, $\mathbf{b}'$, can be directly computed through the α-vectors included in the original POMDP policy, as indicated in Equation (3.2). For instance, if the decision-maker follows the optimal policy and opts for a perfect repair action at year 8, the suggested subsequent optimal action is DN-NI (Figure 3.2b). Instead, if the decision-maker selects a DN-NI action at year 8, the logical suggested action is an immediate repair action the following year. Note that in the reported results, the total expected cost, i.e., $V(\mathbf{b}')$, is computed from the original POMDP policy. In order to exactly evaluate the economic implication of selecting suboptimal actions, the formulated POMDP can be solved again, considering $\mathbf{b}'$ as the initial belief in a reduced finite decision horizon, which corresponds, in this particular example, to twelve time steps. However, only minor differences in the estimation of the expected total cost between the two aforementioned approaches are observed in this study.

Further examining all alternative action-observation combinations available at year 8, Figure 3.3 showcases typical resulting policy realizations. As one could expect, a perfect repair is suggested after a DN-NI action is selected at year 8 (Figure 3.3b), then the structural component condition is restored, and no additional future interventions are usually needed. If a DN-I action is taken at year 8, the POMDP policy suggests a subsequent repair action after a crack detection inspection outcome is observed (Figure 3.3c), whereas if the inspection results in a no

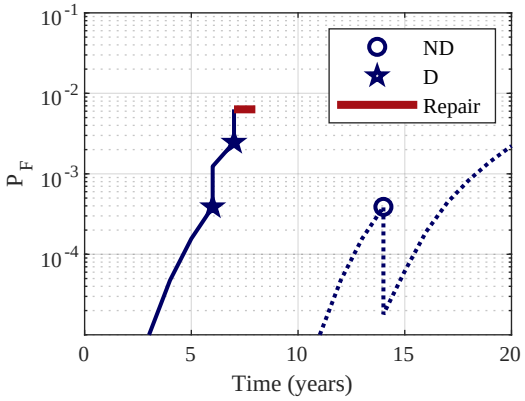

(a) Policy realization for a_{PR-NI} (following the original optimal POMDP policy).

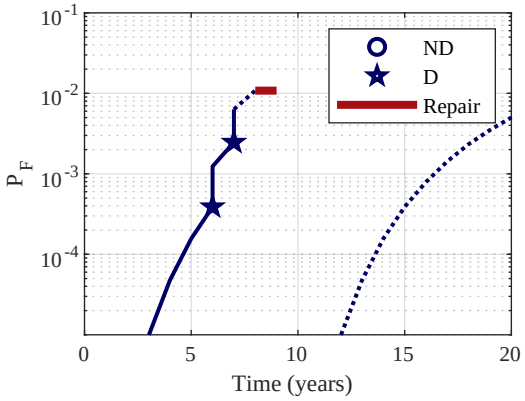

(b) Policy realization for a_{DN-NI}. A repair action is immediately suggested the next year.

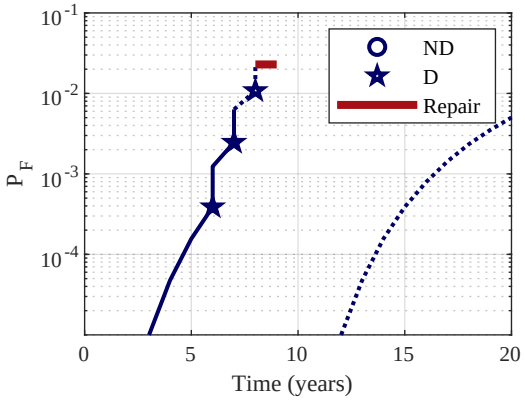

(c) Policy realization for a_{DN-I}, o_D. A repair action is immediately suggested the next year.

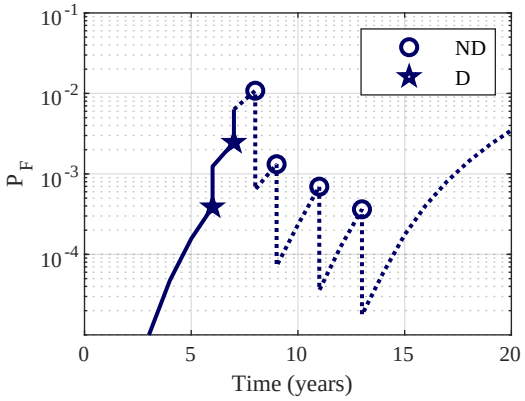

(d) Policy realization for a_{DN-I}, o_{ND}. A repair action is not subsequently suggested.

Fig. 3.3. Representation of typical policy realizations for all action-observation combinations available at year 8.

detection outcome, a repair is not planned, and instead, the policy realization shows a series of inspections for the remainder of the horizon (Figure 3.3d).

In summary, Table 3.1 lists the expected total cost associated with all available actions at year 8 and their corresponding regret $\mathbb{E}[C_P]$. It can be observed that, in this study, the most suboptimal choice is a DN-NI action, as it results in a higher failure risk after year 8, while a repair still needs to be allocated the following year. Interestingly, a DN-I action is less suboptimal, in this case, since subsequent no detection inspection outcomes can still be observed, thus slightly reducing the need of a perfect repair action.

3.4 Conclusion

In an offshore wind inspection and maintenance (I&M) planning context, we describe the fundamentals of partially observable Markov decision processes (POMDPs) -based policies and showcase their inherent adaptive and flexible properties. Through a typical offshore wind

Table 3.1. Regret incurred when selecting alternative actions other than the one suggested in the optimal original POMDP policy.

Action at year 8	$\mathbb{E}[C_T]$ (monetary units)	$\mathbb{E}[C_P](\%)$
PR-NI (Optimal)	12,493	-
DN-NI	15,552	24.5
DN-I $\begin{cases} p(o_{ND} \mid \mathbf{b'}_{7,a_{DN-I}}) = 0.5437 \\ p(o_D \mid \mathbf{b'}_{7,a_{DN-I}}) = 0.4563 \end{cases}$	15,428	23.5

I&M planning case study, we also demonstrate that decision-makers following POMDP-based strategies can efficiently and swiftly quantify the effect of selecting alternative actions rather than those suggested in the optimal POMDP policy. Based on the reported benefits offered by POMDP-based policies in terms of optimality [26, 27], adaptability [29, 43], and flexibility [133], along with the interpretability aspects introduced in this work, we encourage the adoption of POMDP-based I&M planning methods in the offshore wind industry.

Authorship contribution statement

Hlaing, N.: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing - Original draft, Writing - Review & editing, Visualization.
Morato, P. G.: Conceptualization, Validation, Formal analysis, Writing - Review & editing. **Papakonstantinou, K. G.:** Methodology, Validation, Formal analysis, Writing - Review & editing. **Andriotis, C. P.:** Methodology, Validation, Formal analysis, Writing - Review & editing. **Rigo P.:** Supervision, Project administration.

Acknowledgements

This research is funded by the Belgian Energy Transition Fund (FPS Economy) through PhairywinD and MaxWind projects.

Chapter 4

Probabilistic virtual load monitoring of offshore wind substructures: A supervised learning approach

Paper Hlaing, N., Morato, P. G. and Rigo, P. (2022). Probabilistic Virtual Load Monitoring of Offshore Wind Substructures: A Supervised Learning Approach. In *The Proceedings of the 32nd International Ocean and Polar Engineering Conference, ISOPE-2022*, volume 4, 3137–3144.

Abstract In this work, a virtual load monitoring framework is proposed for deriving a mapping from either high or low frequency (1-Hz /10-minute time averaged) SCADA data to load signals, while preserving the high frequency dynamic components of the latter. Specifically, the proposed virtual load monitoring scheme relies on a data-driven model that receives features retrieved from SCADA data and yields the probability distribution of the structural response. The constituent neural networks are trained via supervised learning based on the labelled data retrieved while strain sensors are still functional, since at that operational stage, both SCADA and structural response can be collected concurrently. Once the strain sensors are not functional, the trained deep neural network is deployed, providing structural response predictions from on-site SCADA data. The proposed virtual monitoring approach is tested on a monopile-supported offshore wind turbine and cross-validated in terms of the predicted stress range distribution of a structural connection located at the mudline. The results show good agreement between structural response predictions and measurements, thus demonstrating the efficacy and utility of the tested scheme.

4.1 Introduction

Manual or robotic inspections have conventionally been planned and conducted in order to collect useful information for the evaluation of the performance and structural degradation of offshore wind turbines (OWTs). In the context of increasing wind farm installations in deeper waters and farther from shore, monitoring solutions can provide timely information on wind turbine performance and/or structural health. Along with the advancement of sensor technologies, various monitoring instruments are currently installed on offshore wind turbines, e.g., SCADA system, accelerometers, and strain sensors. Monitoring fatigue deterioration has

become an important objective for assessing the health of OWT support structures since they are exposed to the interaction of environmental loads and structural dynamics induced by turbine loads. Strain gauges are often installed at fatigue critical locations to collect structural response data that can be effectively utilized to reduce uncertainties associated with fatigue deterioration.

An approach to reduce the epistemic uncertainty associated with the long-term load distribution relying on strain monitoring data has been introduced in [96], within the proposed decision-making framework for optimally managing offshore wind farms. In general, the long-term stress range exhibited by offshore wind structural components is typically described by a two-parameter Weibull distribution [117, 129]. Due to the physical and model uncertainties associated with load/ load effect estimations, a high coefficient of variation of the Weibull scale parameter is often assumed. The uncertainty might also further increase during the structure's operational life as a result of the time-varying behavior of wind turbines, influenced by variations in soil stiffness, scouring, etc. The stress range data retrieved from strain monitoring can be used to update the prior belief and to reduce the uncertainty of the Weibull scale parameter. In this context, the uncertainty of strain monitoring can be characterized as the probability distribution over the states of the 'observed' Weibull scale parameter given the 'true' Weibull scale parameter. The measurement quality of strain sensors combined with imperfect Weibull distribution fitting are, therefore, inherently considered in the uncertainty of strain monitoring. Bayesian inference is then conducted to update the scale parameter through Weibull distribution fitting of stress ranges for a certain period of strain monitoring (e.g., 1 year).

In several reported methods for the estimation of remaining useful lifetime, the fatigue damage computed from strain monitoring data serves as the basis to assess the remaining lifetime of the analyzed offshore wind turbines. Short-term 10-minute damages are, for instance, computed from strain measurements using Miner's rule and binned into different environmental conditions using SCADA and metocean data. Estimating the occurrence probability of several environmental conditions either from design documents or long-term monitoring of SCADA data, the computed short-term 10-min damages are extrapolated to lifetime damage. By comparing with the design values or target reliability threshold, the remaining useful fatigue life of OWTs can be reassessed [134, 135].

Strain monitoring data, as previously explained, can effectively reduce the epistemic uncertainty associated with offshore wind fatigue damage estimations, thus enabling more informed and rational inspection and maintenance decisions. However, under the influence of the harsh marine environment, the health of strain sensors also deteriorates over time, thereby affecting sensors' functionality, e.g., strain sensors' life might be limited to only 2 or 3 years of active operation. On the other hand, SCADA data is normally collected throughout the operational lifetime of wind turbines. This has directed offshore wind research towards data-driven load prediction methods, i.e., SCADA-based virtual load monitoring. The purpose of virtual load monitoring is to develop a model that relates SCADA data to the loads/load effects while the

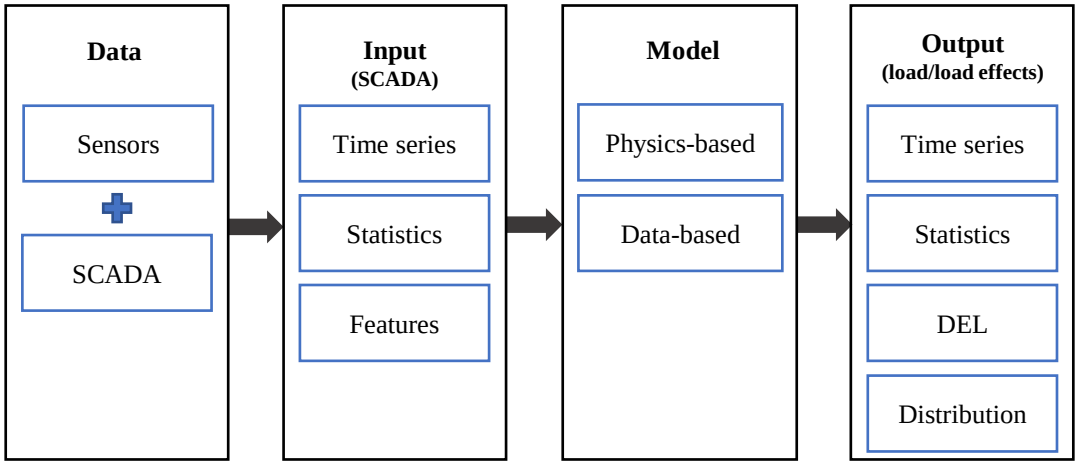

Fig. 4.1. Schematic representation of reported virtual load monitoring approaches [4].

strain sensors are able to collect data. When the strain sensors are not functional, the developed model becomes, therefore, a virtual sensor that provides information on loads/load effects from collected SCADA data.

The development process of a virtual load monitoring scheme is graphically described in Figure 4.1, showcasing various approaches at each step. The estimation models can be categorized as physics-based or data-based. Some examples of physics-based methods include finite element models, modal expansion, Kalman filters, or state estimation, whereas data-driven estimation models are supported by artificial neural networks or other statistical methods. Most existing methodologies have focused on learning the relationship between the features/ statistics of 10-minute averaged SCADA data to short-term damage equivalent load (DEL) [4, 134]. More recently, high frequency 1-Hz SCADA data has enabled the estimation of quasi-static load component in time series, for example, via deep learning [67]. The superposition of other sensor data, e.g., accelerations, is required so as to account for the dynamic component of load signals [62]. Nevertheless, the applicability of such methods is currently limited as long as high frequency SCADA data is expensive and not easily accessible.

In this work, a virtual load monitoring framework is proposed for predicting fatigue load from either high or low frequency (1-Hz /10-minute time averaged) SCADA data. The investigated data-driven prediction model relies on artificial neural networks that are trained via supervised learning based on the features retrieved from SCADA data and the probability distribution of concurrent structural response data. In contrast to other virtual load monitoring methods, the dynamic components of the load signals are preserved, even if the input SCADA data is encoded at low frequency. Implementation details needed for developing a supervised learning approach, able to map low frequency SCADA to strain data, are explained in the following section. Thereafter, the uncertainties associated with the introduced virtual load monitoring model are characterized. In order to showcase the practical applicability of the developed method, the fatigue damage of an offshore wind structural connection is calculated based on

the stress range prediction generated by an already trained neural network. The results show good agreement between predicted and labelled stress range distributions and fatigue damage estimates, thus verifying the efficacy of the proposed virtual load monitoring approach.

4.2 Methodology

4.2.1 Mapping SCADA to strain data

In typical offshore wind turbine monitoring systems, strain sensors usually sample data at high frequency, i.e., around 20-50 Hz or depending on the manufacturer and/or the monitoring objective. The recorded data is stored in one output file every 10 minutes. The sampling frequency of SCADA data is lower compared to strain monitoring. SCADA data is recorded at 1-Hz, however only 10-minute statistics are usually stored to reduce data storage. Nowadays, 1-Hz SCADA data can also be obtained with added expense.

A method to map either high or low frequency (1-Hz/ 10-minute averaged) SCADA data to the associated load effects is proposed herein. The framework presented in Figure 4.2 is based on 10-minute averaged SCADA data. The following procedures are performed for each 10-minute output file from a strain sensor. First, the strain/stress signals at the measured locations are extrapolated to the hotspot stress at the location where a fatigue assessment is to be performed [129]. The extrapolated stress signals then undergo rainflow counting into stress range bins. Normalizing the rainflow results by the total number of stress cycles, a probability distribution

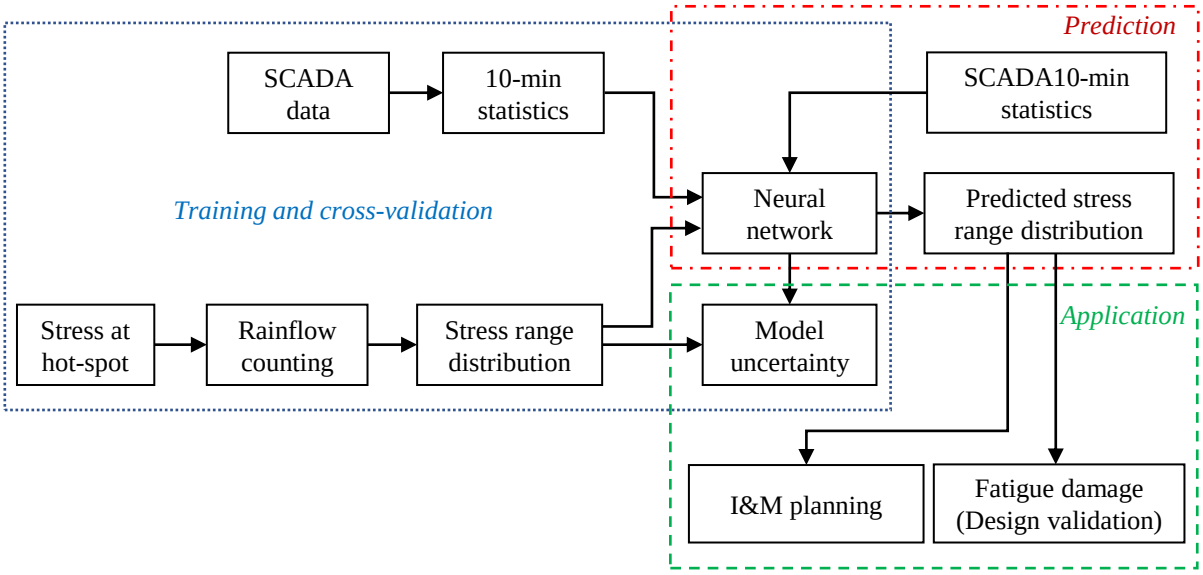

Fig. 4.2. Framework for the development and deployment of a virtual load monitoring model scheme that maps low frequency SCADA data to high frequency strain data.

of stress range occurrence during 10-minute period is obtained such that:

$$\sum_{i=1}^{S_n} P(s_i) = 1, \tag{4.1}$$

where S_n is the number of stress range bins and $P(s_i)$ stands for the occurrence probability of the i^{th} stress range bin.

Once the data has been processed, a deep neural network can be laid out and trained to estimate the mapping between 10-minute averaged SCADA data (inputs) and the associated stress range probability distribution (labels). Among various types of neural networks, fully connected feed-forward networks are relevant for this application due to the efficient gradient computation of the loss through back propagation. The choice of the network's architecture, i.e., the number of hidden layers and the number of neurons, is particularly dependent on the amount of training data as well as a trade-off between training time and accuracy.

The non-linear mapping between SCADA and strain data can be introduced through activation functions in each layer of the neural network. There exist several state-of-the-art differentiable activation functions, which are often used in machine learning and that enable the network to learn non-linear patterns contained in the dataset. Since the output layer represents a probability distribution of stress ranges, a *Softmax* function or normalized exponential function should be encoded in that layer, formulated as:

$$P(s_i) = \frac{e^{x_i}}{\sum_{j=1}^{S_n} e^{x_j}}, i = 1, 2, ...S_n, \tag{4.2}$$

where $\mathbf{x}$ is the input vector to the activation function. The *Softmax* function provides an output vector of probabilities $P(s_i) \in (0, 1)$ for $S_n > 1$, while the normalization term ensures the sum of all the outputs to be one.

Regarding the size of dataset, more than 50,000 (10-minute averaged) data files can be obtained to develop a virtual load monitoring model if a strain sensor continuously records for over one year. A major portion of the recorded data is used to train the network, whereas the rest is employed to evaluate the trained model. As mentioned previously, the proposed mapping is also applicable to high frequency SCADA data. In such case, one can follow one of the following two options:

1. The statistics/features are extracted from each 10-minute period of high frequency SCADA time series. Post-processing and rainflow counting can normally be performed on 10-minute output file of strain monitoring.

2. The strain/stress data is partitioned into 1-second time series and the rainflow counting is performed on the 1-second time series. It should be noted that there might be losses of some stress cycles due to splitting of the temporal load data. Another potential drawback

is the storage and loading of the training dataset which is 600 times larger than the first option.

4.2.2 Load prediction model uncertainty

The uncertainty characterization of a virtual load monitoring model is comprised of the following aspects [92]:

- Parametric variability due to different environmental/ operational conditions that the turbine might experience rather than those from the training dataset.
- Residual variability as a result of the elimination of random process's variations through averaging, e.g., averaging of SCADA data and imperfect Weibull distribution fitting.
- Observation error resulting from inaccurate measurements of SCADA devices and strain sensors.
- Model inadequacy revealed by the discrepancy between model predictions and the true mean value of the physical process.

Observation error and residual variability due to Weibull distribution fitting can be accounted as strain monitoring uncertainty [96]. Using the virtual load monitoring model additionally induces parametric variability and model inadequacy. The modulation of parametric variability, except for extreme events, can also be achieved through avoiding overfitting of the neural network and encouraging generalization during the training phase. Recalibration of the network's weights with new environmental/ operational conditions and sensor data can also reduce parametric variability.

Model inadequacy can be quantified by analyzing the neural network's predictions with respect to the sensor data, which, corrected from observation errors, is considered as the true mean value of the physical process. For a more detailed description of model uncertainty (inadequacy) quantification, the interested reader is referred to [2]. On a worth-noting remark, the wind turbine behaviour may change after the removal/damage of strain sensors and the neural network might then become less accurate in representing the real wind turbine behaviour. Consequently, the strain sensors might be re-installed after certain years so as to recalibrate and fine-tune the neural networks, thus revisiting and treating model inadequacy.

4.3 Virtual load monitoring of an offshore wind structural connection

4.3.1 Numerical simulations

Structural response signals from installed sensors and concurrently recorded SCADA data can be used to train a virtual load monitoring model. Post-processing of structural response signals is required to transform them into a probability distribution of load effects, i.e., stress ranges. Alternatively, structural response simulations can be performed via physics-based analysis of the wind turbine with the objective to produce a training dataset. In this paper, the simulated physics-based data of a monopile-supported wind turbine has been used for the training/testing of the neural networks. Note that, if this option is selected for practical applications, the physics-based model should be calibrated in order to ensure a consistent dynamic behavior with respect to on-site measurements, and the physics-based model uncertainty should also be adequately quantified.

Aero-hydro-servo-elastic simulations in both operational and parked conditions are performed for NREL 5MW wind turbine mounted atop a monopile with a rigid foundation in 20 m water depth. Specifications of the turbine and foundation are described in [136, 137]. The Ijmuiden site is chosen for the gathering of wind and wave climate data, which is located in the Dutch North Sea, near the already built Noordzeewind OWEZ project [138]. The water depth of the chosen site is also closely in line with that of the NREL 5MW monopile. The environmental and operational variables used in the simulations are listed in Table 4.1.

The input environmental parameters are obtained through ancestral sampling from the joint scatter diagrams of the chosen site. Firstly, the random samples of wind speed V_w are withdrawn from the Weibull distribution, and the significant wave height H_s, peak spectral period T_p, wind and wave directions $\theta_{wind}, \theta_{wave}$ are then sampled from two-dimensional scatter diagrams for each of the sampled wind speeds. The initial rotational speed is obtained as a function of the sampled wind speeds V_w based on the turbine specifications. The yaw error is assumed to be

Table 4.1. Environmental and operational parameters.

Variable	Description	Ref:
Wind speed (V_w)	$V_w \sim$ Weibull (scale = 10.49, shape = 2.08)	[138]
Wind direction (θ_{wind})	$P(\theta_{wind}, \theta_{wave} \mid V_w)$	[138]
Turbulence intensity (I)	IEC-3 ($I_{15m/s} = 0.14$)	[139]
Significant wave height (H_s)	$P(H_s, T_p \mid V_w)$	[138]
Peak period (T_p)	$P(H_s, T_p \mid V_w)$	[138]
Wave direction (θ_{wave})	$P(\theta_{wind}, \theta_{wave} \mid V_w)$	[138]
Rotational speed (ω)	$f(V_w)$	[136]
Yaw error (θ_{yaw})	$\theta_{yaw} \in [-10, 10]$	

independent of other parameters and considered as a uniform random distribution within 10 degrees clockwise and counterclockwise from the wind direction.

Coupled aero-hydro-servo-elastic simulations are run on OpenFAST, a multi-fidelity tool for wind turbine simulations. The simulation output consists of 10-minute time series of bending moments and forces at the mudline, sampled at a frequency of 50 Hz. The stress at a fatigue hotspot, which is located on the transverse Y-axis (0° nominal wind direction) of the substructure coordinate system, is computed from the simulation outputs according to:

$$\sigma_{hs} = SCF \cdot \left(-\frac{F_{Zss}}{A_{sub}} + \frac{M_{Xss} \cdot R_{sub}}{I_{sub}} \right), \tag{4.3}$$

where F_{Zss} and M_{Xss} stands for the reaction force and bending moment at the mudline obtained from wind turbine simulations. R_{sub} corresponds to the outer radius of the monopile. A_{sub} and I_{sub} are the cross-sectional area and the second moment of area respectively. The stress concentration factor SCF corresponding to the welded detail is estimated as recommended by industrial design standards [129]:

$$SCF = 1 + \frac{3(\alpha_m - \alpha_0)}{t} e^{-\alpha}, \text{ where } \alpha = \frac{0.91 L_w}{\sqrt{D_{sub} t_{sub}}}. \tag{4.4}$$

where α_m is the maximum misalignment and 5% of the monopile thickness is assumed. α_0 is the misalignment inherent in the SN data for butt welds and $\alpha_0 = 0$ is recommended to consider narrow fabrication tolerances [129]. L_w is the width of weld which is taken equal to the thickness of the monopile, and D_{sub} and t_{sub} are the diameter (6 m) and the thickness (60 mm) of the monopile, respectively.

Rainflow counting is then conducted on the hot-spot stress time series, computed from Equation (4.3), into a number of stress range bins. A discrete probability distribution (labels) is then obtained through the normalization by the total number of stress range cycles. On a worth-noting remark, the discretization scheme of the stress range bins is crucial since it will govern over the accuracy of the virtual monitoring scheme as well as on the fitting of a continuous probability distribution, envisaged for further applications. The error of probability distribution fitting approaches to zero when the number of stress range bins tends to infinite. However, it will also increase the number of zero-probability stress range bins, which, negatively affects the accuracy of virtual monitoring. In this paper, multiple trials are tested and the following discretization scheme with $n_{bin} = 50$ is finally selected:

$$\text{Edges} : \Delta\sigma_{hs} = \left[0, 0.5 : \frac{50 - 0.5}{n_{bin} - 2} : 50, \infty \right] \tag{4.5}$$

The input variables or stored pseudo-SCADA data consist of 10-minute mean values of wind speed V_w, wind direction θ_{wind}, significant wave height H_s, peak period T_p, wave direction θ_{wave}, nacelle position $\theta_{nac} = \theta_{wind} + \theta_{yaw}$, and rotational speed ω. The dataset includes SCADA (inputs) and stress range probability distribution (labels) of 50,000 numerical simulations (80%

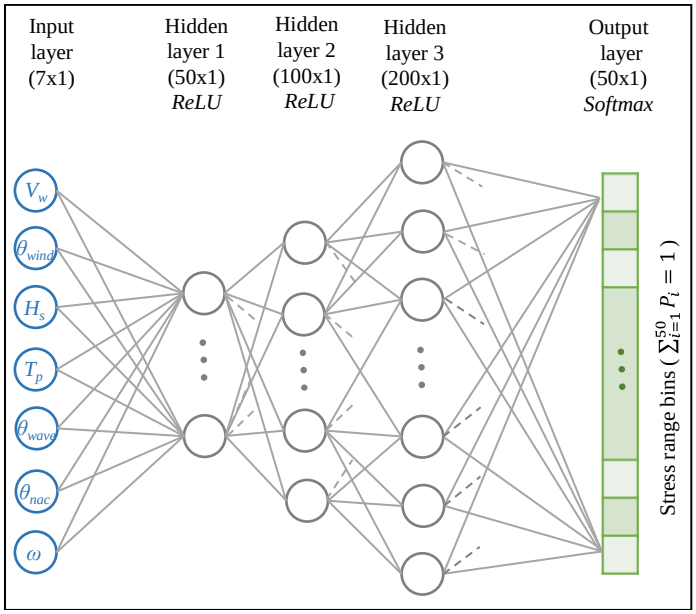

Fig. 4.3. Neural network architecture of the proposed virtual load monitoring scheme.

in the training set and 20% in the test set) under both operational and parked conditions. The link to the dataset compiled throughout this work is provided at the end of the chapter.

4.3.2 Training the neural network

The virtual load monitoring model is implemented using the deep learning library *Keras* [140], written in Python, and running on the machine learning platform *TensorFlow*. The input parameters with different ranges of values undergo data normalization into a standard normal distribution before they are fed to the neural network. After testing various model's architectures, a fully connected feed-forward neural network is chosen, consisting of three hidden layers with 50, 100, 200 neurons, respectively, as shown in Figure 4.3. The hidden layers include rectified linear unit ($ReLU$) activation functions. The output layer, featuring a *Softmax* function, yields a discrete probability distribution over 50 stress range bins.

The neural network is trained with *RMSprop* optimizer at the default learning rate of 0.001. The training of the network is evaluated in terms of the mean absolute error (MAE) over the stress range bins:

$$MAE = \frac{1}{N}\frac{1}{S_n}\sum_{j=1}^{N}\sum_{i=1}^{S_n} \mid P_{pred,j}(s_i) - P_{data,j}(s_i) \mid, \tag{4.6}$$

where P_{pred} is the predicted probability from the neural network and P_{data} is the probability from the data. S_n is the total number of stress range bins and N is the number of data samples. During the training, the MAE continuously decreases as the number of epochs increases since in fact, the network starts to overfit the training data. To avoid this, the network is evaluated on the validation dataset after every training epoch. The training set is, therefore, randomly split into 80% for training and 20% for validation during the training process. The MAE

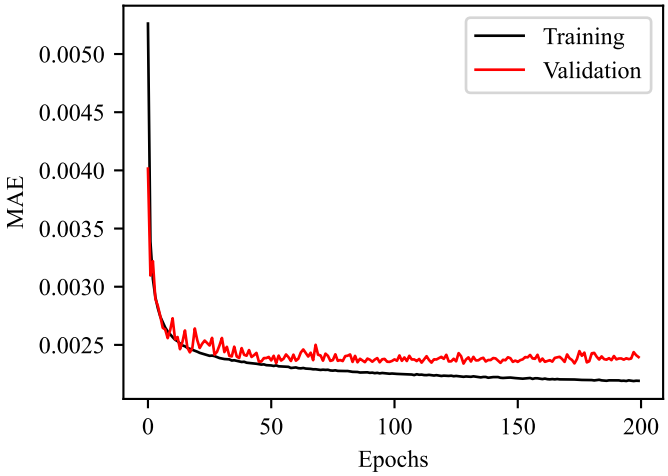

Fig. 4.4. Training and validation of the neural network, tracking the evolution of the mean absolute error (MAE) over epochs.

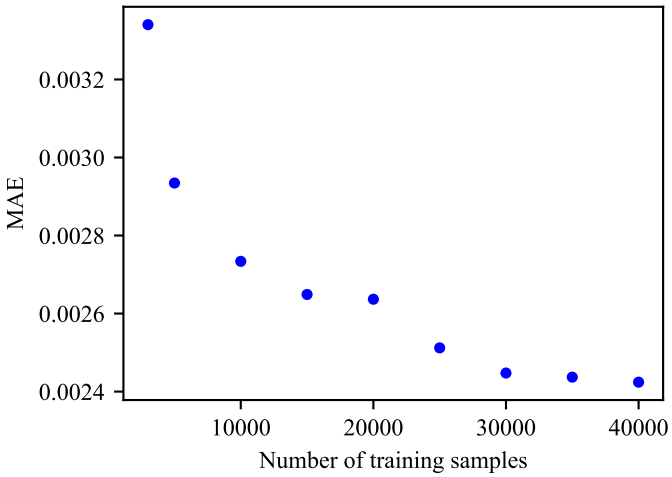

Fig. 4.5. Reduction of the test MAE for increasing size of training data sets.

on the training and validation datasets tracked over 200 epochs is shown in Figure 4.4. The performance on the validation dataset plateaus from 30^{th} epoch, whereas the training MAE continues decreasing due to overfitting. An early stopping callback is introduced, i.e., the training task is stopped if there is no improvement of the validation MAE, thus preventing the risk of overfitting.

4.3.3 Training and cross-validation results

First, the amount of training data on the performance of the neural network has been analyzed. Several neural networks of the same architecture are trained with varying amounts of training data from 3,000 to 40,000 samples. The whole test set is fed to the trained networks to evaluate their performance. Figure 4.5 plots the mean absolute errors on the test set against the number of training samples. As the size of training dataset is increased, the reduction in the error can

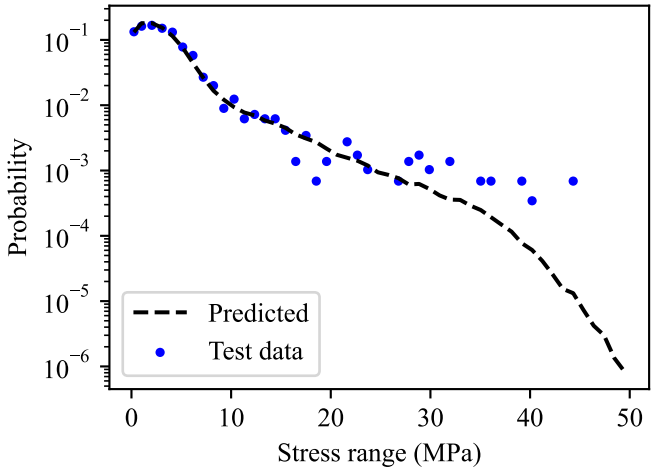

Fig. 4.6. Comparison between predicted and simulated stress range distribution for a 10-minute random test sample.

be observed. So far as the overfitting is controlled, the neural network benefits from larger size of training datasets. However, the decrease in MAE is less significant after the number of training samples is more than 30,000, and therefore, 1 year of SCADA and strain data is, in general, sufficient to train a virtual load monitoring neural network.

Thereafter, comparisons between the probability distributions from simulated data and the neural network's prediction (with the training set size of 40,000 samples) are showcased. Figure 4.6 compares the stress range distribution of a 10-minute random test sample. In the test data, some stress range bins belong to a zero probability from the rainflow counting results. However, the output of the neural networks is always non-zero due to the nature of Softmax function, see Equation (4.2). The function instead outputs very small probabilities. This might result in an underestimation of the probability over other stress range bins, e.g., stress range bins 30-40 MPa, shown in Figure 4.6. This effect can be adverse when higher number of stress range bins are used in the output layer, whereas a finer discretization of stress range is favorable for Weibull distribution fitting.

Figure 4.7 compares the accumulated stress range distribution of the whole test set ($\sim$ 10000 samples). Despite the aforementioned underestimations in low probability regions ($< 10^{-3}$), accurate neural network's predictions can be observed in smaller stress ranges. In addition, the smaller stress cycles are more critical than high stress ranges in terms of fatigue as they contribute more to the damage.

4.3.4 Prospective applications

Fatigue damage estimation

Some prospective applications of the proposed virtual load monitoring are also presented in Figure 4.2. When the strain sensors installed are not functional, one can rely on SCADA data,

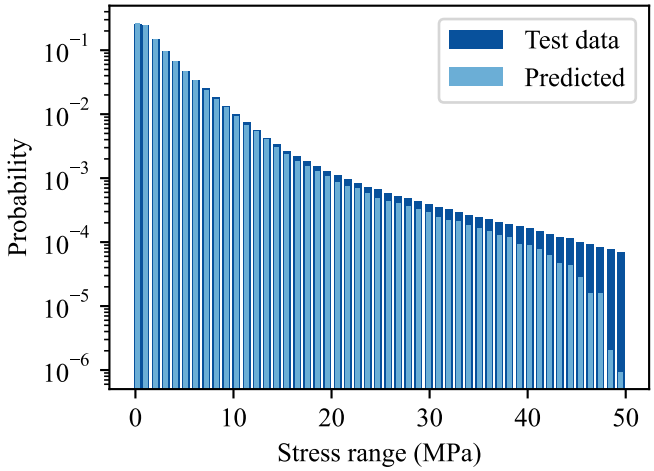

Fig. 4.7. Comparison between predicted and simulated stress range distribution, normalized over the test set.

in conjunction with virtual load monitoring, in order to predict the load effect distribution. The predicted stress range distribution can then be used to compute the fatigue damage, e.g., for design validation. The short-term damage can be easily computed through a forward run of the trained network using the newly recorded SCADA, followed by the application of Miner's rule, considering a deterministic number of stress cycles n within 10 minutes. For a linear SN curve, the fatigue damage can be calculated as:

$$D = n \sum_{i=1}^{S_n} \frac{P(s_i)}{N_{f_i}}, \text{ with } \log_{10} N_{f_i} = \log_{10} K - m \log_{10} s_i. \tag{4.7}$$

where m and K are parameters of the SN curve. $P(s_i)$ is the probability in the i^{th} stress range bin and s_i is the reference stress range value of the bin.

In this hypothetical example, $n = 1200$ cycles/10-minute are considered. The SN curve parameters to calculate the damage at the mudline (free corrosion) are specified as $\log_{10} K = 11.687$, $m = 3$. The network is fed by SCADA data (from the test set) to predict the stress range distribution and short-term fatigue damages are estimated for each 10-minute interval, following Equation (4.7). The comparison of 10-minute fatigue damage from the simulated data and the neural network's predictions is shown in Figure 4.8. The color bar describes the percentage of the test dataset. As the dotted black line represents 100% accuracy, adequate performance of the neural network can be observed on around 70-80% of the test data, yet some spread in the predicted damage can be seen on the remaining 20% of the test data.

Similarly, the long-term fatigue can also be efficiently estimated by damage accumulation while continuously recording and feeding the collected SCADA data, benefiting from the cheap computational cost of retrieving neural network's predictions.

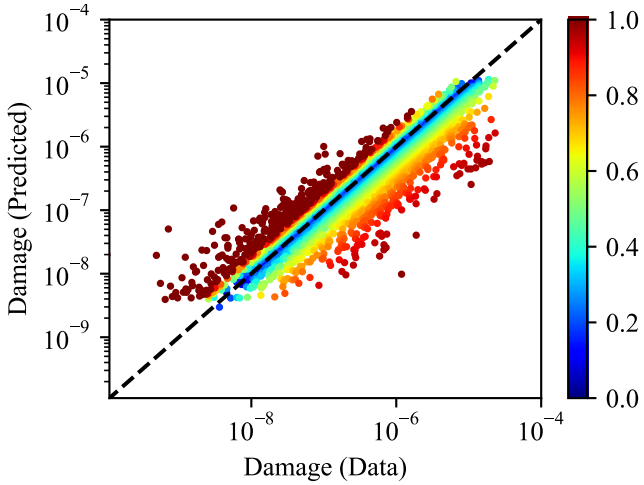

Fig. 4.8. Comparison of short-term fatigue damage between the simulated data and neural network's predictions.

Bayesian inference of load effects

The proposed virtual load monitoring can also be applied to probabilistic fatigue assessments, e.g., decision-making for inspection and maintenance planning, lifetime extension. In this application, the predicted stress range information from the neural network model is used to update the prior stress range distribution. The accumulated probability distribution of the stress bins over a certain period (e.g., 1 year) is fitted to a Weibull distribution, thus constituting an observation of the Weibull's scale parameter. Quantifying the aforementioned uncertainties of virtual load monitoring, the neural network's prediction enables the updating of the scale parameter through Bayesian inference.

Assuming that the stress range distribution follows a Weibull, Miner's rule, formulated in Equation (4.7), becomes:

$$D(t) = nt\left[\frac{q^m}{K}\Gamma\left(1+\frac{m}{h}\right)\right], t = 0, 1, 2, ...T_d, \tag{4.8}$$

where the SN parameter K is now a random variable owing to the uncertainty of SN curves, n is the number of stress cycles per year and often assumed deterministic, Γ is the gamma function, whereas q and h are the parameters of the long-term stress range distribution.

A common practice to obtain the expected value of the scale parameter in inspection and maintenance planning tasks is to calibrate the damage to the design's limit, i.e., $D(T_d) = 1/FDF$ where T_d is the design lifetime and FDF is the fatigue design factor [115]. Table 4.2 lists the values of SN variables used in this example. Owing to the lack of knowledge on potential changes in turbine dynamics, it is assumed that the uncertainty of Weibull scale parameter slightly increases every year according to:

$$q_{t+1} = q_t + \varepsilon, \tag{4.9}$$

where ε is a Gaussian distribution with zero mean and a standard deviation of 0.1.

Table 4.2. Fatigue assessment representative parameters.

Variable	Distribution	Mean	Std (CoV)
FDF	Deterministic	3	
m	Deterministic	3	
$log_{10}K$	Normal	11.687+2Std	0.2
h	Deterministic	0.8	
q	Normal	*1.9776	(0.2)
n	Deterministic	$6.32 \cdot 10^7$	
U_O	Normal	0	0.25
U_{PR}	Normal	0	0.25
U_{MI}	Normal	0	0.5

* calibrated for $T_d = 20$ years and $FDF = 3$.

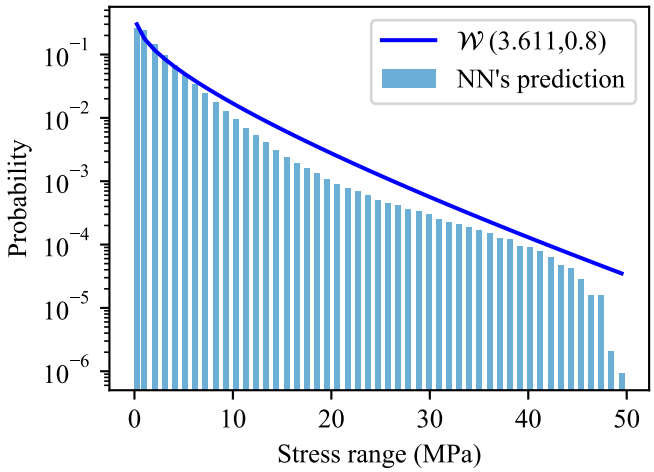

Fig. 4.9. Goodness-of-fit between a continuous Weibull distribution and the discrete stress range distribution obtained from the neural network.

The training set and test set are merged in order that the stress range distribution spans over a long-enough period, and the merged dataset is fed to the neural network. Accumulating the obtained results, the final stress range distribution is fitted to a Weibull distribution with a fixed shape parameter of 0.8. Figure 4.9 illustrates the goodness-of-fit between the fitted Weibull distribution of $q_{pred} = 3.611$ and the neural network's prediction. U_O, U_{PR} and U_{MI} represent the uncertainties of the virtual monitoring associated with observation errors, combined residual and parametric variability, and model inadequacy, respectively. The observed scale parameter q_{obs} can be then quantified as:

$$q_{obs} = q_{pred} + U_O + U_{PR} + U_{MI}, \tag{4.10}$$

where q_{pred} stands for the scale parameter of the fitted Weibull distribution.

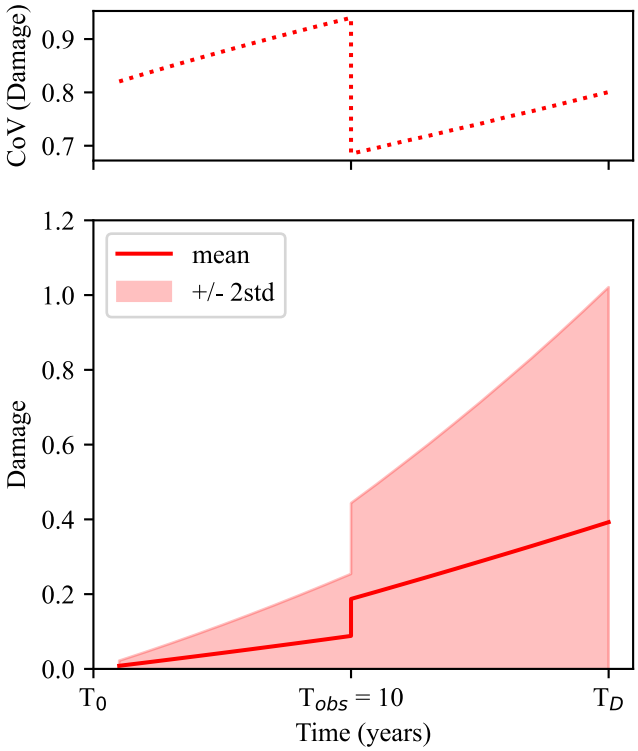

Fig. 4.10. Evolution of estimated fatigue damage over time.

Assuming the SCADA is collected at half of the design lifetime, the scale parameter q is updated using virtual monitoring data, according to Bayes' rule:

$$P(q' \mid q_{obs}) \propto P(q_{obs} \mid q)P(q). \tag{4.11}$$

The fatigue damage in Equation (4.8) can subsequently be updated with the posterior information of q. Figure 4.10 illustrates the evolution of the fatigue damage and its coefficient of variation (CoV) over time. Due to the uncertainty propagation of the prior scale parameter q, the CoV of the fatigue damage increases. Updated by the posterior distribution of q, the mean and the standard deviation of the fatigue damage increase since the observed scale parameter is higher than the prior expected value. However, the uncertainty of the estimated fatigue damage (CoV) is significantly reduced at that year.

Due to limited dataset size, the stress range distribution and fatigue damage are updated only once in this example. In practical scenarios, the SCADA data is collected throughout the operational life of the offshore wind turbines. Bayesian inference can then be conducted every one or two years, thus continuously reducing the statistical uncertainty. It is worth-noting that, the Weibull scale parameter can be considered as a time-invariant random variable in certain decision-making tasks in order to attenuate the computational demand associated to modeling the stochastic deterioration process [29, 133]. In that case, the uncertainty of the Weibull scale parameter, once updated through strain sensor data, will not increase, and the applicability of the virtual monitoring will be limited.

4.4 Conclusions and outlook

The nonlinear relationship between operational SCADA data and other monitoring signals, e.g., strain measurements, can be learned while the information is concurrently collected. Once some sensors are not functional due to, for instance, deterioration experienced in the marine environment, and if operational data is still available, the trained data-driven model can predict other measurements of interest. In most current monitoring setups, operational information is stored at low frequency (10-minute averaged), whereas structural response data is normally collected at a much higher frequency, thus hindering the implementation of informative data-driven predictive models. As demonstrated in this paper, by specifying the output signal as a probability distribution, a virtual load monitoring scheme, able to treat both low and high frequency input signals, can be efficiently implemented via supervised learning. The conducted numerical experiments show that the proposed virtual load monitoring approach can be applied to multiple practical scenarios, from fatigue damage estimations of offshore structures to Bayesian updating of load effects, among others.

Besides the mentioned practical applications, this work unfolds further research opportunities. Robust feature selection methods are, for instance, worth exploring in order to select the most influencing variables and statistical values among the wide spectrum of signals that can be monitored on-site, while avoiding redundant variables. Further research efforts can also be directed to the development of probabilistic deep learning methods, thus enabling the intrinsic quantification of physical and measurement uncertainties by modeling the neural network parameters, i.e., weights and biases, as random variables. The inherent uncertainty associated with data-driven models can be, therefore, reduced as more data is collected and fed to the neural network.

Data availability statement

The simulated and post-processed datasets generated throughout the development of this work are freely available online at the following link: https://doi.org/10.5281/zenodo.5957394.

Authorship contribution statement

Hlaing, N.: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing - Original draft, Writing - Review & editing, Visualization. **Morato, P. G.:** Conceptualization, Validation, Formal analysis, Writing - Review & editing, Supervision. **Rigo, P.:** Supervision, Project administration, Funding acquisition.

Acknowledgements

This research is funded by the Belgian Energy Transition Fund (FPS Economy) through PhairywinD and MaxWind projects. We further acknowledge the “Consortium des Équipements de Calcul Intensif”, at the University of Liege (Belgium), for granting us access to high performance computing facilities.

Chapter 5

Farm-wide virtual load monitoring for offshore wind structures via Bayesian neural networks

Paper Hlaing, N., Morato, P. G., de Nolasco Santos, F., Weijtjens, W., Devriendt, C. and Rigo, P. (2023). Farm-wide virtual load monitoring for offshore wind structures via Bayesian neural networks. *Structural Health Monitoring*. DOI:10.1177/14759217231186048

Abstract Offshore wind structures are exposed to a harsh marine environment and are subject to deterioration mechanisms throughout their operational lifetime. Even if the deterioration evolution of structural elements can be estimated through physics-based deterioration models, the uncertainties involved in the process hurdle the selection of life-cycle management decisions, e.g., lifetime extension. In this scenario, the collection of relevant information through an efficient monitoring system enables the reduction of uncertainties, ultimately driving more optimal life-cycle decisions. However, a full monitoring instrumentation implemented on all wind turbines in a farm might become unfeasible due to practical and economical constraints. Besides, certain load monitoring systems often become defective after a few years of marine environment exposure. Addressing the aforementioned concerns, a farm-wide virtual load monitoring scheme directed by a fleet-leader wind turbine offers an attractive solution. Fetched with data retrieved from a fully instrumented wind turbine, a model can be firstly trained and then deployed, thus yielding load predictions of non-fully monitored wind turbines, from which only standard data remains available, e.g., SCADA. During its deployment stage, the pretrained virtual monitoring model might, however, receive previously unseen monitoring data, thus often producing inaccurate load predictions. In this paper, we propose a virtual load monitoring framework formulated via Bayesian neural networks (BNNs) and we provide relevant implementation details needed for the construction, training, and deployment of BNN data-based virtual monitoring models. As opposed to their deterministic counterparts, BNNs intrinsically announce the uncertainties associated with generated load predictions and allow to detect inaccurate load estimations generated for non-fully monitored wind turbines. The proposed virtual load monitoring is thoroughly tested through an experimental campaign in an

operational offshore wind farm and the results demonstrate the effectiveness of BNN models for 'fleet-leader'-based farm-wide virtual monitoring.

5.1 Introduction

Offshore wind turbines are continuously exposed to a combined wind and wave load excitation, thus inducing fatigue deterioration and other mechanical stressors throughout their service life. The evolution of fatigue damage can be estimated through physics-based engineering models, yet the resulting deterioration predictions contain significant uncertainties. Combined with engineering models, manual and/or robotic inspections can be conducted in order to reduce the uncertainties associated with deterioration estimations, hence supporting more rational and informed maintenance decisions [133, 141]. With the advent of modern sensor technologies, monitoring systems are increasingly being deployed with the objective of continuously monitoring the deterioration experienced by offshore wind structures, thus also enabling decision-makers to make timely and informed decisions [142, 143]. For example, fatigue load monitoring through strain gauges provides valuable information that can be used to estimate the remaining useful fatigue lifetime [135, 144, 145] and/or to update probabilistically modeled time-varying deterioration mechanisms [96, 146, 147].

However, strain gauges, and other monitoring systems, are also prone to deterioration in a harsh marine environment and their operational lifespan is normally shorter than the service life considered for an offshore wind turbine. In this context, virtual load monitoring, either physics- or data-based, offers an adequate solution, providing load information once strain sensors are no longer functional [81]. Each approach features its own advantages and disadvantages, and the choice should be mainly based on one's knowledge in physics and on data availability [148]. Physics-based approaches demand the development of a numerical model in order to simulate the loads under actual environmental and operational conditions (EOCs), which are collected via a supervisory control and data acquisition (SCADA) system. Based on collected on-site structural response information, physics-based models can be periodically updated [66, 83]. Whereas physics-based approaches require less data than fully data-based schemes, the development and implementation of complex dynamical models demand specific knowledge, e.g., thrust coefficient, soil parameters, that is usually hard to retrieve in real case scenarios.

Alternatively, one can benefit from easily-accessible data, e.g., SCADA data, accelerometers, among others, to train a deterministic or probabilistic data-based model [61, 147]. The majority of previous research studies investigate the application of deterministic models for virtually monitoring offshore wind turbines, mostly focusing on learning relationships between SCADA data and fatigue loads [4, 67]. In a few cases, high frequency acceleration measurements are combined with SCADA data in order to adequately capture dynamic load components [62]. Generally, deterministic models do not explicitly indicate the uncertainty associated with the generated predictions, except for model verification analysis with respect to sensor

data [92, 147]. On the other hand, probabilistic approaches can intrinsically provide an indication of the uncertainty contained in the produced predictions. From the limited number of reported probabilistic virtual monitoring methods, [82] and [149] investigated the applicability of heteroscedastic Gaussian processes for probabilistically modeling turbine loads, and similarly, a Gaussian process-based method was also proposed with the objective of extrapolating monitored fatigue damage from an instrumented tower section to any other level, based on a covariance matrix defined via acceleration signals [71].

An additional challenge faced by virtual load monitoring schemes is their applicability at a farm-wide level. In this perspective, physics-based approaches can easily become burdensome since a complex numerical model of the entire wind farm is needed, where each turbine structural design is tailored to a specific water depth, soil conditions, etc. More recently, certain data-based virtual monitoring schemes have adopted a fleet-leader approach, in which load measurements collected from monitored turbines are extrapolated to non-instrumented ones. For instance, a recent study extrapolates short-term damage measurements from an instrumented fleet-leader to the wind farm relying on binned SCADA data and turbine conditions (i.e., operational or parked) [150]. However, other relevant information for the estimation of fatigue damage evolution, e.g., structural dynamics variations among wind turbines, might not be appropriately captured in SCADA data [151]. In this regard, data-based models that combine both SCADA data and response information collected from accelerometers can provide farm-wide fatigue predictions more effectively [72].

As previously explained during the literature survey, recent investigations often rely on deterministic data-based approaches for the formulation of virtual load monitoring models, yet their inability to detect potential conflicts during the deployment stage, e.g., inaccurate load predictions when the network is fed with previously unseen input data, limits their applicability to farm-wide virtual monitoring implementations. Unless additional sensors are installed with the objective of retrieving 'ground truth' load measurements, the model uncertainty associated with predictions generated for non-fully monitored wind turbines cannot be quantified. In this paper, we cast a virtual load monitoring framework that relies on Bayesian neural networks and probabilistic deep learning methods in order to provide farm-wide load predictions and enable the intrinsic quantification of aleatory uncertainty (emerging due to the random nature of the physical system) and epistemic uncertainty (arising due to lack of knowledge of the system). The proposed virtual monitoring framework is tested on a dataset collected from three wind turbines that are currently operating in a Belgian offshore wind farm. In particular, a Bayesian neural network is trained based on SCADA and accelerometer data (inputs) along with the corresponding damage equivalent moments (labels) collected from a specific offshore wind turbine. The reduction of model uncertainty with increasing dataset size is thoroughly quantified and the resulting Bayesian model is cross-validated for the same turbine as well as for other turbines located in the same wind farm. Interestingly, the results showcase that Bayesian models are able to intrinsically report higher model uncertainties when tested on a wind turbine

characterized with a dynamic behavior different from the one employed during the training, thus demonstrating the ability of the proposed virtual load monitoring scheme to automatically inform if the provided predictions are inaccurate and whether further information collection actions are needed.

5.2 Bayesian neural networks

Most recent applications of Bayesian neural networks for offshore wind energy settings have mainly focused on wind speed and power forecasting [152, 153, 154]. Up to the knowledge of the authors, Bayesian neural networks-based structural health monitoring methods have not yet been formally proposed for offshore wind applications. In this section, we briefly introduce Bayesian neural networks from a general theoretical perspective, since this will facilitate the application and understanding of the proposed virtual monitoring framework. In essence, a Bayesian neural network (BNN) is a stochastic artificial neural network trained via Bayesian inference, featuring the combined strength of deep learning and Bayesian theory in order to provide a rich probabilistic interpretation of the generated predictions.

The key defining characteristic of Bayesian neural networks with respect to conventional artificial neural networks (ANNs) is their stochastic neural network components, i.e., stochastic activations and/or weights [155], as illustrated in Figure 5.1, enabling this way multiple model parametrizations $\boldsymbol{\theta}$, each of them associated with a probability distribution $P(\boldsymbol{\theta})$. In most applications, the output prediction of a BNN is also formulated as a probability distribution, thereby quantifying uncertainties associated with the underlying process, i.e., aleatory uncertainties that naturally arise from random physical phenomena or inherent noise in the training data, and model (epistemic) uncertainty due to the limited information used for training the networks.

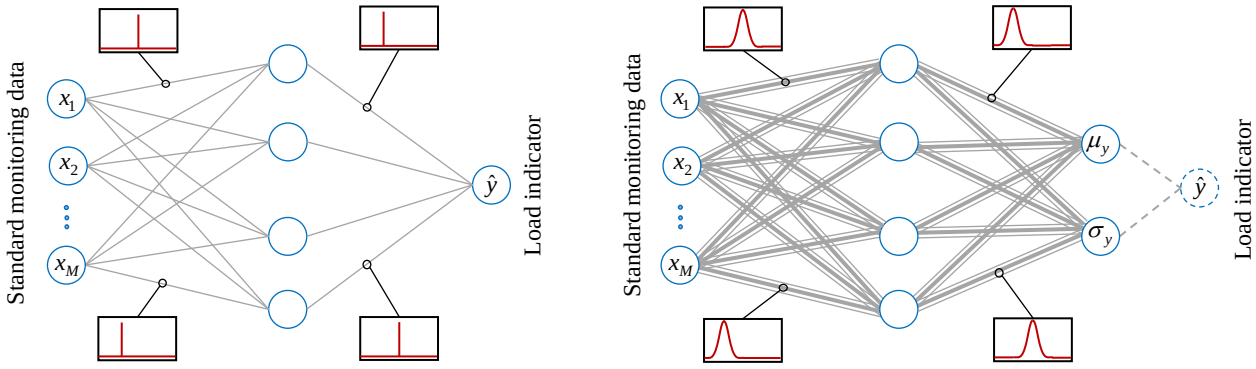

(a) Deterministic neural network with weights and biases specified as point estimates $\boldsymbol{\theta}$ along with a deterministic output.

(b) Bayesian neural network with weights and biases specified as probability distributions $P(\boldsymbol{\theta})$ along with a stochastic output.

Fig. 5.1. Schematic diagrams comparing the topology and constituents of a standard deterministic neural network (DNN) and a Bayesian neural network (BNN), both mapping standard input monitoring data $\boldsymbol{x}$ to a load indicator y.

Whereas aleatory uncertainty is irreducible, model uncertainty can be reduced as the networks learn from additionally considered training data. In contrast to their deterministic counterparts, BNNs report high epistemic uncertainty in regions where only few (or none) training points are available.

When modeling BNNs, and similarly to ANNs, the selection of the network architecture plays a key role. Various ANNs' topologies commonly used in machine learning applications, e.g., feed-forward, convolutional, and recurrent neural networks, are also applicable to BNNs. Additionally, a stochastic model should also be defined, i.e., a prior distribution of model parameters $P(\boldsymbol{\theta})$. While the choice of prior distributions is arbitrary, a Gaussian prior with zero mean and diagonal covariance $\mathcal{N}(\mathbf{0}, \sigma \boldsymbol{I})$ is commonly adopted. Gaussian priors are often preferred due to their advantageous mathematical properties, e.g., its logarithmic formulation is the cornerstone of most learning algorithms.

5.2.1 Inference methods for BNNs

Conditioned to the training dataset D, the posterior probability distribution of neural network weights $P(\boldsymbol{\theta} \mid D)$ can be computed via Bayes' theorem:

$$P(\boldsymbol{\theta} \mid D) = \frac{P(D \mid \boldsymbol{\theta}) P(\boldsymbol{\theta})}{\int P(D \mid \boldsymbol{\theta}) P(\boldsymbol{\theta}) d\boldsymbol{\theta}}. \tag{5.1}$$

The calculation of the posterior distribution is usually intractable for continuous probabilistic settings. Therefore, various methods have been developed in order to estimate the Bayesian posterior, e.g., Laplace approximation [156], variational inference [157], Markov chain Monte Carlo sampling [158].

Among various Markov chain Monte Carlo (MCMC) algorithms, Metropolis-Hastings [159, 160] has been widely used in Bayesian statistics [161], benefiting from the fact that only a proportional distribution to the posterior is needed. Despite MCMC algorithms enable the estimation of posterior distributions through sampling processes, their applicability is still limited to small datasets and medium complex models, e.g., 10 to 100 variables. Alternatively, variational inference (VI) has been widely used for settings featuring large datasets and highly complex models with thousands to millions of parameters and can be applied to most neural network architectures [162]. The interested reader is directed to Jospin et al. [155] for a detailed overview of VI and other inference methods applicable to Bayesian neural networks.

The objective of variational inference (VI) is to approximate the potentially complex posterior distribution of weights by a simpler one, denoted as variational distribution. Gaussian distributions are often used to estimate the posteriors, whose parameters $\lambda = (\mu_{\boldsymbol{\theta}}, \sigma_{\boldsymbol{\theta}})$ are, therefore, commonly known as variational parameters. VI methods adjust λ so that the variational distribution $q_\lambda(\boldsymbol{\theta})$ closely resembles the posterior $P(\boldsymbol{\theta} \mid D)$. The similarity or divergence between the two distributions is formally described by the Kullback-Leibler (KL) divergence [163], a non-symmetric and information-theoretic measure of the statistical difference between

two probability distributions. An optimal solution for the variational distribution $q_\lambda(\boldsymbol{\theta})$ is then obtained by minimizing KL divergence between $q_\lambda(\boldsymbol{\theta})$ and the posterior $P(\boldsymbol{\theta}|D)$. Mathematically, KL corresponds to the expected value of the difference between the logarithmic probabilities associated with the two distributions:

$$\begin{aligned}\mathbb{KL}\left(q_\lambda(\boldsymbol{\theta}) \,||\, P(\boldsymbol{\theta} \mid D)\right) &= \int q_\lambda(\boldsymbol{\theta}) \log \frac{q_\lambda(\boldsymbol{\theta})}{P(\boldsymbol{\theta} \mid D)} d\boldsymbol{\theta} \\ &= \int q_\lambda(\boldsymbol{\theta}) \log \frac{q_\lambda(\boldsymbol{\theta})}{\frac{P(\boldsymbol{\theta}, D)}{P(D)}} d\boldsymbol{\theta} \\ &= \int q_\lambda(\boldsymbol{\theta}) \log P(D) d\boldsymbol{\theta} - \int q_\lambda(\boldsymbol{\theta}) \log \frac{P(\boldsymbol{\theta}, D)}{q_\lambda(\boldsymbol{\theta})} d\boldsymbol{\theta} \\ &= \log P(D) - \int q_\lambda(\boldsymbol{\theta}) \log \frac{P(D \mid \boldsymbol{\theta}) P(\boldsymbol{\theta})}{q_\lambda(\boldsymbol{\theta})} d\boldsymbol{\theta}.\end{aligned} \tag{5.2}$$

Since the first term is independent with respect to the variational parameters, minimizing $\mathbb{KL}\left(q_\lambda(\boldsymbol{\theta}) \,||\, P(\boldsymbol{\theta} \mid D)\right)$ is equivalent to maximizing the second term, often denoted as the evidence lower bound objective (ELBO):

$$\begin{aligned}\text{ELBO} &= \int q_\lambda(\boldsymbol{\theta}) \log \frac{P(D \mid \boldsymbol{\theta}) P(\boldsymbol{\theta})}{q_\lambda(\boldsymbol{\theta})} d\boldsymbol{\theta} \\ &= -\int q_\lambda(\boldsymbol{\theta}) \log \frac{q_\lambda(\boldsymbol{\theta})}{P(\boldsymbol{\theta})} d\boldsymbol{\theta} + \int q_\lambda(\boldsymbol{\theta}) \log P(D \mid \boldsymbol{\theta}) d\boldsymbol{\theta}.\end{aligned} \tag{5.3}$$

In particular, the loss function to be minimized corresponds to the negative ELBO and the optimal variational parameter λ^* can be, therefore, formulated as:

$$\lambda^* = \text{argmin}\{\mathbb{KL}\left(q_\lambda(\boldsymbol{\theta}) \,||\, P(\boldsymbol{\theta})\right) - \mathbb{E}_{\boldsymbol{\theta} \sim q_\lambda}\left[\log\left(P(D \mid \boldsymbol{\theta})\right)\right]\}. \tag{5.4}$$

The first term represents the KL divergence between the variational distribution $q_\lambda(\boldsymbol{\theta})$ and the known prior $P(\boldsymbol{\theta})$ and it makes sure the variational distribution is close to the prior distribution. When priors are selected with zero mean, minimizing $\mathbb{KL}\left(q_\lambda(\boldsymbol{\theta}) \,||\, P(\boldsymbol{\theta})\right)$ resembles the concept of regularization, i.e., driving weight estimates toward zero [164]. The second term of Equation (5.4) computes the expected negative log-likelihood of the training data given the weight $\boldsymbol{\theta}$ distributed according to $q_\lambda(\boldsymbol{\theta})$. Minimizing this term controls that BNN's produced predictions match training target data.

When a BNN is being trained, the loss function cannot be back-propagated through $\boldsymbol{\theta}$ since it follows a probability distribution. In this scenario, the derivative of the loss with respect to the variational parameters cannot be obtained. However, the following reparameterization trick enables the formulation of $\boldsymbol{\theta}$ as a deterministic function of the variational parameters:

$$\boldsymbol{\theta} = \mu_{\boldsymbol{\theta}} + \sigma_{\boldsymbol{\theta}} \cdot \varepsilon, \text{ where } \varepsilon \sim \mathcal{N}(0, 1). \tag{5.5}$$

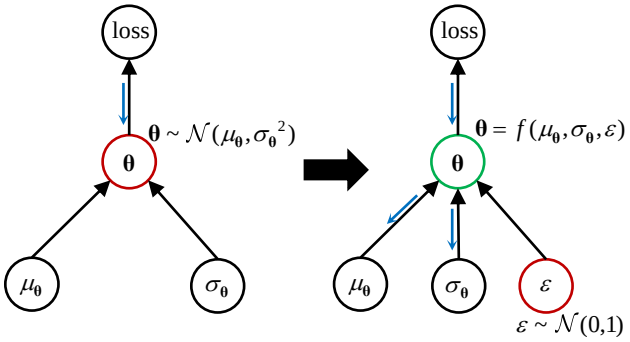

Fig. 5.2. Graphical representation of the reparametrization trick, where by reformulating stochastic network parameters $\boldsymbol{\theta}$ as a function of statistical distribution parameters and additional stochastic inputs, the back-propagation of the loss with respect to variational parameters can be effectively computed.

Through this formulation, one can compute the derivative of the loss with respect to the variational parameters, as shown in Figure 5.2. During forward prediction runs, $\boldsymbol{\theta}$ is obtained through sampling from a standard normal distribution ε, instead of sampling directly from the variational distribution $q_\lambda(\boldsymbol{\theta})$ so as to facilitate the implementation of the aforementioned reparametrization formulation. There are also other possible solutions for computing the gradient when random variables are included in the neural network, e.g., score function estimator [165], VarGrad [166], straight-through estimator [167], among others. The reparametrization approach described before is widely adopted in practice owing to its capability for generating unbiased gradient estimates.

It is also worth-noting that, to reduce computation efforts, the negative log-likelihood can be evaluated, in some cases, for only one $\boldsymbol{\theta}$ sample, instead of computing the expectation over several realizations, as described in Equation (5.4). The resulting gradient descent is noisy, yet it can still find its path toward the minimum loss. At the expense of drastically increasing the computational demand, a more accurate gradient can be computed by sampling multiple realizations.

5.3 Farm-wide virtual load monitoring through Bayesian neural networks

In general, the full instrumentation of an offshore wind farm with strain sensors becomes economically impractical due to the elevated installation and maintenance costs associated with the process. In this context, virtual load monitoring offers an efficient alternative by providing load information - denoted hereafter as 'load indicator' - based on readily available monitoring data - denoted hereafter as 'standard monitoring data', e.g., SCADA and accelerometer data. Thus, virtual load monitoring constitutes a natural solution for the implementation of a farm-wide monitoring strategy, i.e., one or a set of representative turbines, commonly designated as

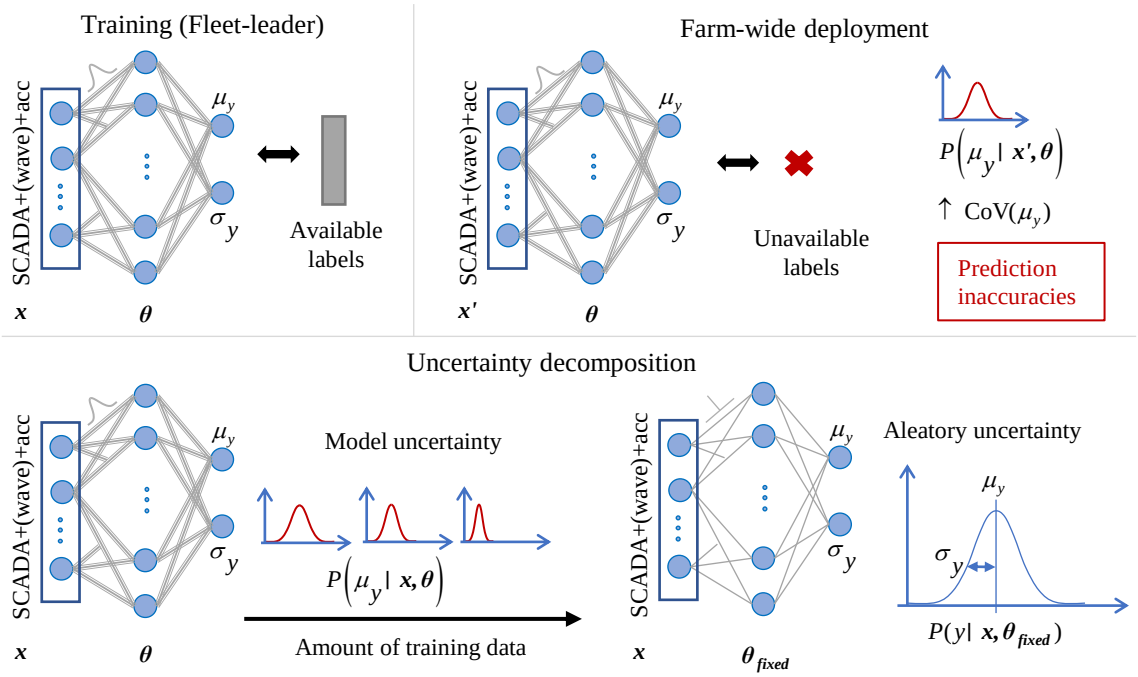

Fig. 5.3. Rationale of the proposed farm-wide virtual load monitoring framework featuring Bayesian neural networks as data-based virtual sensors. (Top left) A fleet leader BNN is trained based on available load measurement labels. (Top right) At the deployment stage (measurement labels are no longer available), the pre-trained BNN indicates whether the generated predictions might be inaccurate by reporting a high model uncertainty. (Bottom) Uncertainty decomposition is enabled by the proposed BNN approach, yielding information on: (i) the need to collect more data for improving the model's performance, (ii) the intrinsic variability of the analyzed phenomena.

a fleet-leader, is fully instrumented enabling the training of a data-based model, which then provides load indicator predictions to the other non-fully instrumented wind turbines. As opposed to conventional deterministic virtual load monitoring schemes, this paper proposes a probabilistic virtual monitoring method, which by indicating the uncertainty associated with the 'load indicator', intrinsically informs the quality of the generated predictions. The overarching rationale of the proposed virtual monitoring method is illustrated in Figure 5.3, highlighting BNNs' capabilities for automatically detecting potential prediction inaccuracies when the virtual model is deployed to wind turbines where load measurements are not available. Moreover, it is also showcased in the figure how the overall involved uncertainty can be decomposed into model and aleatory components, providing information on (i) whether more data is needed to improve model's performance, and (ii) capturing the intrinsic variability associated with the analyzed phenomena, respectively. A more specific description of the general procedure for implementing a BNN-based virtual load monitoring model is summarized in Figure 5.4, and will be further explained in the following subsections.

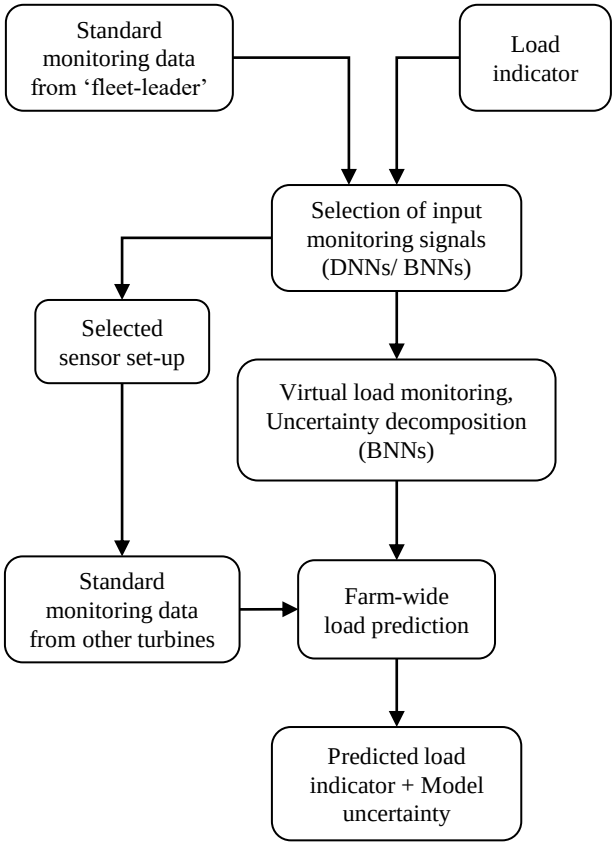

Fig. 5.4. Flowchart diagram illustrating the steps needed for the implementation of the proposed farm-wide virtual load monitoring framework.

5.3.1 Selection of the input monitoring signals

In principle, various environmental, operational, and/or motion signals, e.g., SCADA, wave data, accelerations, can be monitored and fed as inputs to the virtual load monitoring model. However, a reduced selection of meaningful monitoring signals as inputs to the model will ease the instrumentation setup of non-fully monitored wind turbines, and overall alleviate practical constraints during the deployment of the farm-wide virtual monitoring strategy. In order to select the reduced set of monitoring signals, a data-based model can be tested for various potential configurations, and by observing the generalization error associated with each considered setting, the setup that results in the desirable trade-off between accuracy and monitoring equipment complexity can be then chosen.

The generalization error can be estimated through either deterministic or probabilistic approaches, yet the metrics employed to assess the error vary between them. Whereas deterministic neural networks (DNNs) are constructed assuming point estimates for the constituent weights and biases, Bayesian neural networks (BNN) consider weights and biases as probability distributions, as shown in Figure 5.1. Besides, even if both DNNs and BNNs similarly receive input monitoring signals $\boldsymbol{x} = \{x_1, x_2, ...x_M\}$, DNNs provide a deterministic load indicator as output, while BNNs' output becomes a probabilistic load indicator. For the case of deterministic

neural networks, the standard metrics, e.g., mean absolute error (MAE) or root mean squared error (RMSE), can be used as reference to compare the performance of the analyzed virtual load monitoring models, as:

$$MAE = \frac{1}{N}\sum_{i=1}^{N}|y_i - \hat{y}_i|, \tag{5.6}$$

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}, \tag{5.7}$$

where N stands for the number of test samples, and $\hat{y}$ represents the model prediction, compared against the ground truth (label), y.

On the other hand, a probabilistic output is provided by BNNs, as mentioned previously, from which random samples can be drawn. More specifically, the output layer features the statistical parameters of a specified probability distribution, e.g., a Gaussian $y \sim \mathcal{N}(\mu_y, \sigma_y)$, which are then fed to an additional distribution layer to be able to draw random samples of the load indicator $\hat{y}$, as well as to compute likelihood of the label, i.e., $P(y \mid \boldsymbol{x}, \boldsymbol{\theta})$. To assess the performance of a BNN, one cannot rely on MAE or RMSE, since the model outputs are random realizations. Instead, metrics that provide a probabilistic interpretation should be observed, e.g., the expected log-likelihood of the label given the prediction model, defined as:

$$\mathbb{E}[\mathcal{L}(y)] = \frac{1}{N}\frac{1}{N_f}\sum_{i=1}^{N}\sum_{N_f}\log\left(P(y_i \mid \mu_{y,i}, \sigma_{y,i})\right), \tag{5.8}$$

where μ_y and σ_y stand for the output statistical parameters predicted by the model. Note that the output statistical parameters are, for the case of BNNs, also stochastic, resulting from the random realizations drawn from the network's stochastic weights and biases, and thus statistical properties of the likelihood can be retrieved via numerical simulations, i.e., sampling N_f forward predictions.

5.3.2 Bayesian neural networks uncertainty decomposition

As explained before, the statistics of BNN's predicted results can be computed at the deployment stage through numerical simulations, e.g., Monte Carlo sampling. For instance, one can estimate the expected value and the predictive uncertainty of the load indicator $\hat{y}$ given newly acquired standard monitoring data $\boldsymbol{x}$:

$$\mathbb{E}[\hat{y} \mid \boldsymbol{x}] = \frac{1}{N_f}\sum_{N_f} f(\hat{y} \mid \boldsymbol{x}, \boldsymbol{\theta}), \tag{5.9}$$

$$\mathbb{V}(\hat{y} \mid \boldsymbol{x}) = \frac{1}{N_f}\sum_{N_f}\left(f(\hat{y} \mid \boldsymbol{x}, \boldsymbol{\theta})\right)^2 - \left(\mathbb{E}[\hat{y} \mid \boldsymbol{x}]\right)^2, \tag{5.10}$$

where the network parameters, $\boldsymbol{\theta}$ are randomly drawn from the posterior weights and biases associated distributions, and f symbolizes the Bayesian network model itself. Note that, in this case, the retrieved predictive uncertainty estimate $\mathbb{V}(\hat{y} \mid \boldsymbol{x})$ encompasses both aleatory and epistemic contributions.

On the one hand, aleatory uncertainties arise from the inherent randomness of physical phenomena and/or the presence of noise in sensing devices. While the physical uncertainty is irreducible [168], measurement uncertainty can be reduced by modulating the noise level of sensors, albeit it cannot be controlled by adjusting the model. On the other hand, epistemic uncertainties are induced by the quality of the model and can be reduced by improving the model. For instance, at the deployment stage, a trained BNN might indicate high epistemic uncertainty if data outside of the training domain is fed to the network, and after the model is retrained from representative data in the reported high uncertainty region, the BNN's epistemic uncertainty can be further reduced. Theoretically, the epistemic uncertainty will be totally dissipated in the limit of infinite available training data. In practice, however, there exists no perfect model for predicting the response of complex engineering systems, i.e., the model might not consider all representative features, and the additional uncertainty associated with the missing or unavailable latent variables is sometimes also denoted as aleatory uncertainty [169].

A decomposition of the overall uncertainty retrieved by the BNNs, into its aleatory and epistemic contributions, becomes highly informative when deploying the trained network to the farm-wide level. An indication of high global predictive uncertainty does not inherently report the need for retraining the model, since the variability might correspond to noise present in the observations (labels). However, further data collection and model retraining actions can be planned as a result of observed high model uncertainty. Not isolated to virtual load monitoring applications, uncertainty decomposition is an active topic within the probabilistic machine learning community [168, 169, 170]. The overall uncertainty can be decomposed, according to the law of total variance, as follows:

$$\mathbb{V}(\hat{y} \mid \boldsymbol{x}) = \mathbb{E}[\sigma_y^2 \mid \boldsymbol{x}] + \mathbb{V}(\mu_y \mid \boldsymbol{x}). \tag{5.11}$$

In general, complex engineering systems are exposed to aleatoric physical phenomena, i.e., the system response for a given set of input parameter values does not correspond to a single output value. Since conventional numerical simulators are very often deterministic, the intrinsic variability is normally accounted for by running multiple simulations with different random seeds for each input combination, e.g., wind and wave conditions specified following offshore wind design practices and recommendations [90, 91]. In BNNs, the aleatory uncertainty is captured by learning the variance parameter σ_y^2, thus inherently yielding a probabilistic output. The first term $\mathbb{E}[\sigma_y^2 \mid \boldsymbol{x}]$ in Equation (5.11) can be interpreted as the aleatory component, computed as:

$$\mathbb{E}[\sigma_y^2 \mid \boldsymbol{x}] = \frac{1}{N_f} \sum_{N_f} (f(\sigma_y \mid \boldsymbol{x}, \boldsymbol{\theta}))^2, \tag{5.12}$$

The epistemic uncertainty is encapsulated in the probability distribution of the network's weights and biases. The variance of the BNN's predicted means $\mathbb{V}(\mu_y \mid \boldsymbol{x})$, computed as in the following equation, therefore explains the epistemic uncertainty,

$$\mathbb{V}(\mu_y \mid \boldsymbol{x}) = \frac{1}{N_f} \sum_{N_f} \left(f(\mu_y \mid \boldsymbol{x}, \boldsymbol{\theta})\right)^2 - \left(\mathbb{E}[\mu_y \mid \boldsymbol{x}]\right)^2 . \tag{5.13}$$

where,

$$\mathbb{E}[\mu_y \mid \boldsymbol{x}] = \frac{1}{N_f} \sum_{N_f} f(\mu_y \mid \boldsymbol{x}, \boldsymbol{\theta}), \tag{5.14}$$

It should be noted that there also exists the uncertainty of BNN's predicted variance $\mathbb{V}(\sigma_y^2 \mid \boldsymbol{x})$, yet its contribution to the overall uncertainty is insignificant enough to be neglected.

5.3.3 Farm-wide load prediction

Once the reduced set of input monitoring signals, i.e., standard monitoring data, has been identified by quantifying the generalization error through either deterministic or Bayesian neural networks, and the BNN model corresponding to the fleet-leader has also been trained, the following steps are to be implemented for farm-wide load prediction:

- Collection and treatment of the required standard monitoring data $\boldsymbol{x}$ from non-fully monitored turbines.
- Deployment of the trained BNN on non-fully monitored turbines, from which multiple forward simulations are run, thus randomly drawing load indicator realizations, $\hat{y}$.
- Computation of the load indicator expected value $\mathbb{E}[\hat{y} \mid \boldsymbol{x}]$.
- Estimation of the model epistemic uncertainty $\mathbb{V}(\mu_y \mid \boldsymbol{x})$, and the performance metric $\mathbb{E}[\mathcal{L}(y)]$ if the target labels are available, and further information collection actions might be decided depending on the observed model uncertainty metric.

5.3.4 Epistemic Bayesian neural network

In certain cases, one might only be interested in a virtual sensor model that provides a mapping between the inputs and deterministic predicted output(s). In that case, an epistemic variant of the proposed BNN can be implemented by assigning zero to the output variance node(s) (i.e., $\sigma_y \to 0$), hence disregarding the potential aleatory uncertainty associated with the output response and only seeking the prediction of the mean response output node (i.e., μ_y). Theoretically, the response of typical engineering systems normally contains physical aleatory uncertainty, yet more data is required to train a probabilistic output response compared to a point estimate. In any case, an epistemic BNN can yield model uncertainty information because the weights and biases are still described by probability distributions. Since no aleatory

uncertainty is now incorporated into the model, i.e., the first term in Equation (5.11) becomes zero and the total predictive uncertainty directly corresponds to model uncertainty. Whether aleatory uncertainty is included or not, farm-wide deployment can be executed in both cases since the BNN model will inform valuable model uncertainty metrics about its confidence in the generated predictions.

5.4 Experimental campaign: Probabilistic virtual load monitoring in an offshore wind farm

The proposed framework for farm-wide virtual load monitoring is hereafter implemented and tested for the specific case of an existing offshore wind farm. In this study, we do not rely only on environmental and operational data, but also incorporate acceleration signals within the standard monitoring data that will be used to predict the load indicator, as motion information provides an indication of wind turbines' structural dynamics and significantly influences fatigue load estimations [72].

5.4.1 Monitoring setup and dataset description

This investigation relies on the data collected during the course of a 2-year monitoring campaign, from early 2018 to the end of 2019, conducted on three wind monopile-supported turbines located in a Belgian offshore wind farm. The overall monitoring setup installed above the mean sea level is graphically illustrated in Figure 5.5. The turbines are equipped with a standard supervisory, control and data acquisition (SCADA) system, continuously collecting both environmental and operational data. More specifically, the collected SCADA data contains 10-minute statistics of wind speed and direction, turbine rotational speed, yaw and pitch attitude, and instantaneous power. The sensor setup does not feature wave measurements, yet data from Meetnet Vlaamse Banken [171]'s Westhinder wave buoy was additionally collected, thus providing wave height, wave period and wave direction information. As aforementioned, the monitored turbines are also equipped with three accelerometers installed on the transition piece and tower at 17 m, 38 m, and 77 m above lowest astronomical tide (LAT), respectively. The collected accelerometer data reports acceleration statistics, i.e., minimum, maximum, and root-mean-square (rms), from 10-minute time series of fore-aft (FA) and side-to-side (SS) accelerations.

Along with the standard monitoring data mentioned previously, strain gauges are installed at 17 m above the lowest astronomical tide (LAT) level of the wind turbines, so as to collect load signals, i.e., bi-axial bending moments, M_{tl} in SS direction and M_{tn} in FA direction. Following common offshore wind industrial and scientific practices [90], the time series of the monitored loads are post-processed into representative damage equivalent loads (DELs), i.e., equivalent load range such that the damage caused by a pre-defined number of the equivalent load amplitude cycles N_{eq} equals the damage D_{SN} caused by the original load time series, as computed using

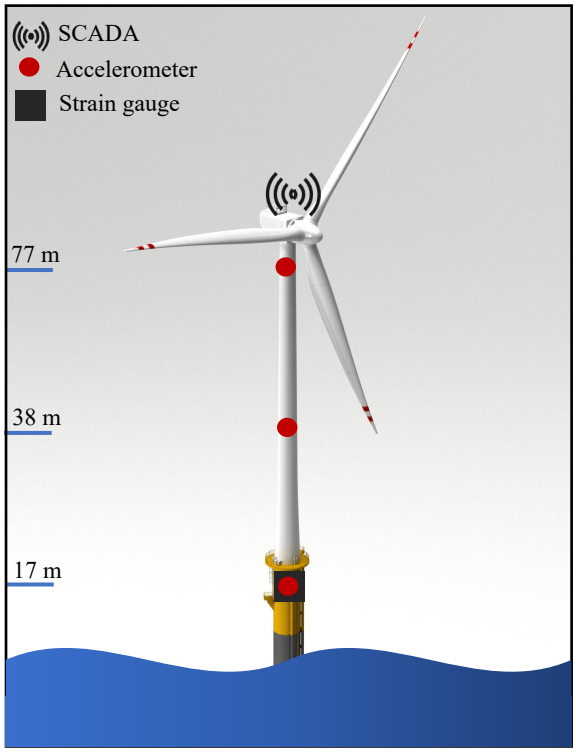

Fig. 5.5. Illustration depicting the monitoring setup installed on an operational offshore wind turbine, from which data was continuously collected during the course of the experimental campaign. The monitoring setup includes a standard SCADA system, accelerometers at three different levels, and strain gauges installed at the lowest level.

Miner-Palmgren's rule:

$$D_{SN} = \sum_{i=1}^{N_s} \frac{n_i}{k \cdot S_{r,i}^{-m}} = \frac{N_{eq}}{k \cdot DEL^{-m}}, \tag{5.15}$$

where k and m correspond to linear S-N curve parameters, N_s stands for the number of stress range bins in the load spectrum, whereas $S_{r,i}$ and n_i represent the reference value for the i^{th} stress range bin and the number of cycles inside that bin, respectively. Note that the equivalent stress cycles N_{eq} are commonly specified as 10^7. In this work, the damage equivalent moment (DEM) is computed for each retrieved 10-minute time series of bending moment measurements, as follows:

$$DEM = \left(\frac{\sum_{i=1}^{N_m} n_i \cdot M_{r,i}^m}{N_{eq}} \right)^{1/m}, \tag{5.16}$$

where N_m stands for the number of bins in the load spectrum, $M_{r,i}$ and n_i represent the reference moment value for the i^{th} bin and the number of cycles in that bin. One can straightforwardly compute the fatigue damage from the calculated DEM by also considering the geometrical properties (i.e., thickness, second moment of area), as well as the specified SN curve parameters k and m. In this application, the estimated damage equivalent moments DEM_{tl} and DEM_{tn} constitute the output load indicators that are provided as labels to the virtual load monitoring model.

Table 5.1. Description of the dataset.

	Sensor	Monitoring signal	Symbol	Units
Input	SCADA	Rotational speed (mean)	$\mu[RPM]$	rpm
		Yaw angle (mean)	$\mu[Yaw]$	deg
		Pitch angle (mean)	$\mu[Pitch]$	deg
		Power (mean)	$\mu[Power]$	kW
		Wind speed (mean)	$\mu[WSpd]$	m/s
		Wind speed (std)	$\sigma[WSpd]$	m/s
		Wind direction (mean)	$\mu[WDir]$	deg
	Wave buoy	Wave height	H_s	cm
		Average wave period	T_p	s
		Wave direction	θ_w	deg
	Accelerometers	FA acceleration (max)	$max[acc_{FA}]$	g
	- LAT-017	FA acceleration (min)	$min[acc_{FA}]$	g
	- LAT-038	FA acceleration (rms)	$rms[acc_{FA}]$	g
	- LAT-077	SS acceleration (max)	$max[acc_{SS}]$	g
		SS acceleration (min)	$min[acc_{SS}]$	g
		SS acceleration (rms)	$rms[acc_{SS}]$	g
Output	Strain gauges	DEM (side-to-side)	DEM_{tl}	MNm
	- LAT-017	DEM (fore-aft)	DEM_{tn}	MNm

In summary, Table 5.1 lists and describes each dataset considered in this work, overall containing 28 monitoring input signals: 7 from the SCADA system, 3 from the wave buoy, and 18 from the accelerometers; and 2 monitoring output signals retrieved from the strain gauges. As aforementioned, the data is recorded for three wind turbines within the same offshore wind farm, and during the investigation, the data retrieved from the fleet-leader wind turbine is employed for the development and training of the virtual load monitoring model, where the data is randomly split into 75% for training and 25% for testing. The data collected from the other two turbines is fully reserved for farm-wide load prediction purposes.

5.4.2 Selection of the input monitoring signals

Following the procedures provided in the framework, a reduced set of representative and informative input monitoring data is carefully selected through both deterministic and Bayesian neural network approaches. In order to decide the reduced set of standard monitoring data, the generalization error is computed for the following model configurations of monitoring input signals:

1. SCADA + wave
2. SCADA + wave + accelerometer (LAT-017)

3. SCADA + wave + accelerometer (LAT-038)
4. SCADA + wave + accelerometer (LAT-077)
5. SCADA + wave + accelerometers (LAT-017, 038)
6. SCADA + wave + accelerometers (LAT-017, 038, 077)
7. SCADA
8. SCADA + accelerometer (LAT-017)
9. SCADA + accelerometer (LAT-038)
10. SCADA + accelerometer (LAT-077)
11. SCADA + accelerometers (LAT-017, 038)
12. SCADA + accelerometers (LAT-017, 038, 077)

Note that a distinction has been made in the selection process between model configurations which include, or do not include, wave data. It is worth exploring alternatives that purely use the turbine's monitoring data without relying on the secondary wave data, thus naturally simplifying the later implementation of a farm-wide load monitoring.

5.4.3 Deterministic neural networks

The deterministic neural networks (DNNs) implemented in this investigation rely on *Keras* [140] API, which forms part of the machine learning and artificial intelligence library TensorFlow. The topology of the DNNs features a fully connected feed-forward neural network consisting of three hidden layers with 64, 128, 64 neurons, respectively. Naturally, the width of the input layer varies according to the tested input monitoring signal combination, and the output layer features two neurons corresponding to the predictions of damage equivalent moments in the side-to-side (DEM_{tl}) and fore-aft (DEM_{tn}) directions. Note that all network layers are equipped with rectified linear unit ($ReLU$) activation functions.

During the course of the training task, the neural networks are trained via *Adamax* optimizer at the default learning rate of 0.001, minimizing the loss function corresponding to the mean absolute error (MAE):

$$MAE = \frac{1}{2N} \sum_{i=1}^{N} \left(\left| DEM_{tl,i} - \hat{DEM}_{tl,i} \right| + \left| DEM_{tn,i} - \hat{DEM}_{tn,i} \right| \right), \tag{5.17}$$

where N stands for the total number of training samples, DEM and $\hat{DEM}$ represent the ground truth (measurements) and predicted damage equivalent moments, respectively. For each tested combination, the training task conducted 200 epochs, running stochastic gradient descent based

on randomly collected 32-sample batches. Moreover, potential overfitting conflicts are avoided by implementing an 'early stopping callback', i.e., the training task is stopped if there is no improvement in the validation MAE [147]. In this regard, the training dataset is randomly split into 80% for training and 20% for validation purposes.

As described before, we have overall tested 12 input monitoring signal setups. Figure 5.6a illustrates the performance of each DNN model with respect to the considered input monitoring setup, indicating in the generalization error computed as:

$$Error(\%) = \frac{100}{N}\sum_{i=1}^{N}\left|\frac{DEM_i - D\hat{E}M_i}{DEM_i}\right|, \tag{5.18}$$

where N stands for the total number of test samples. Interestingly, including wave data as an input to the model is beneficial in configurations in which information on accelerations is not available, where one can observe a reduction of approximately 6.3% (DEM_{tn}) and 10.1% (DEM_{tl}) in the computed generalization error. However, the benefit of feeding the network with wave data becomes negligible once acceleration information becomes also available, resulting in around 1% reduction of the generalization error in all accelerometer-integrated setups. As a result, only primary wind turbine monitoring signals, i.e., combination of SCADA system and accelerometers, can be deemed, in this case, as a satisfactory input monitoring setup, thus avoiding the need of relying on secondary wave data during the deployment stage of the virtual load monitoring model.

Furthermore, it can also be observed that the generalization error of the model including SCADA + accelerometer (LAT-077) as inputs remains very similar to the error reported for the input setup that features SCADA + accelerometers (LAT-017, 038, 077). Specifically, the generalization error for the case in which all accelerometers information is considered results in 3.0% (DEM_{tl}) and 5.3% (DEM_{tn}), whereas the error corresponding to the setup where only the top accelerometer is included as input to the model results in 3.1% (DEM_{tl}) and 5.8% (DEM_{tn}), respectively. Installing two additional accelerometers at the lower levels reduces only 0.1-0.5% in the errors. Since fatigue is primarily driven by the first structural mode to which the top accelerometer is more sensitive compared to the other two lower levels, it provides the most informative data to predict the loads. Therefore, the input monitoring signal setup SCADA + top accelerometer interestingly outperforms the monitoring input combination of SCADA + accelerometers (LAT-017, 038), even though the load prediction is conducted for a location situated at the lowest level (LAT-017).

5.4.4 Bayesian neural networks

The Bayesian neural networks introduced in this investigation are implemented with the support of the probabilistic deep learning library TensorFlow Probability. As for the case of the DNNs, the width of the Bayesian neural network (BNN) input layer is defined based on the specified input monitoring data, along with three hidden layers equipped with $ReLU$

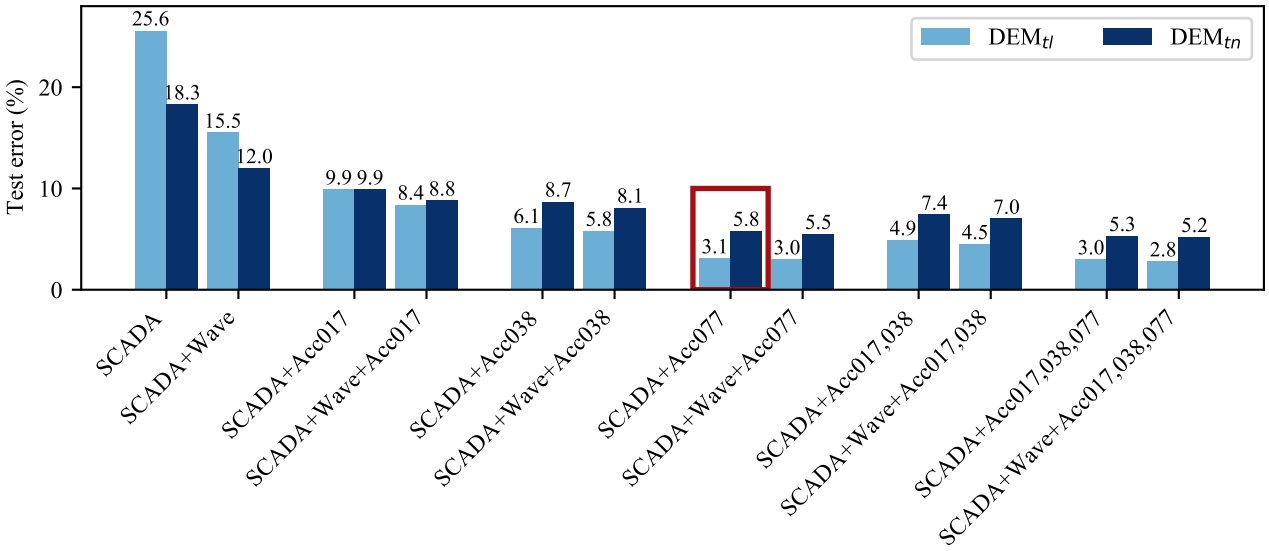

(a) Errors reported during the testing stage of the implemented deterministic neural networks. The lower error indicates a better performance of the model.

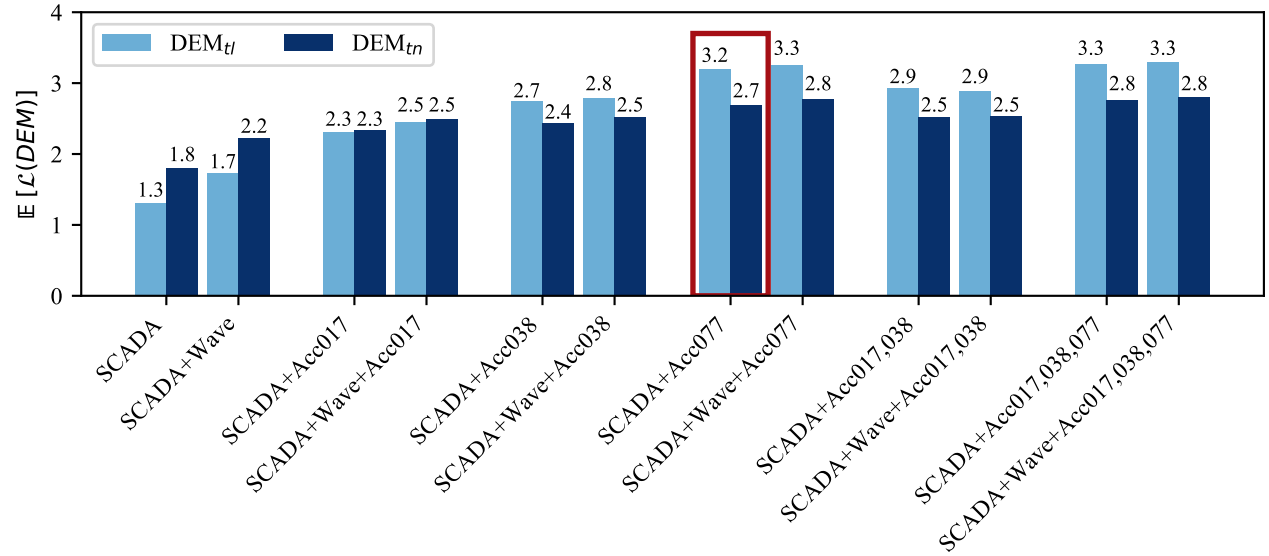

(b) Expected log-likelihood reported during the testing stage of the implemented Bayesian neural networks. The higher expected log-likelihood indicates a better performance of the model.

Fig. 5.6. Illustration showcasing the performance of load prediction data-based models. Each bar corresponds to a model specified with a specific set of input monitoring signals, i.e., SCADA, wave, and/or acceleration data. The red box indicates the selected reduced set of input monitoring signals.

activation functions, and an output layer with 4 neurons, from which the output statistical parameters $\mu_{DEM_{tl}}, \sigma_{DEM_{tl}}, \mu_{DEM_{tn}}, \sigma_{DEM_{tn}}$ are estimated. Note that, for the case in which the standard deviation of the output load indicator distribution is known and fixed, the BNN can be alternatively laid out with a reduced output layer containing only the mean statistical parameter that drives the resulting load distribution. More specifically, both hidden and output layers of the BNNs are built through *DenseFlipout*, implementing Bayesian variational inference via a Flipout estimator. Since the training samples in each batch share the stochastic weights ε, there is potential correlation in the resulting gradients which can lead to inefficient training. Flipout is an efficient method to improve variance reduction by implicitly sampling pseudo-independent stochastic weights for each training data and therefore decorrelating the gradients within a batch. For more detailed description of the Flipout estimator, the reader is referred to Wen et al. [172].

A thorough description of the topology and training environment considered for both deterministic and Bayesian neural networks is showcased in Appendix A, providing details for each tested input monitoring setup. In general, the number of neurons included in the hidden layers of BNNs is lower than the neurons specified for the deterministic counterparts. If the same architecture would be specified for both DNNs and BNNs, the training for the latter will naturally take longer as Bayesian inference normally demands more computational resources than classical backpropagation. However by considering the network parameters as probability distributions, BNNs are intrinsically more informative, despite having less neurons, and additionally provide beneficial properties to avoid overfitting with respect to standard DNNs. As shown in Figure 5.7, a DNN starts to overfit the training data after some training

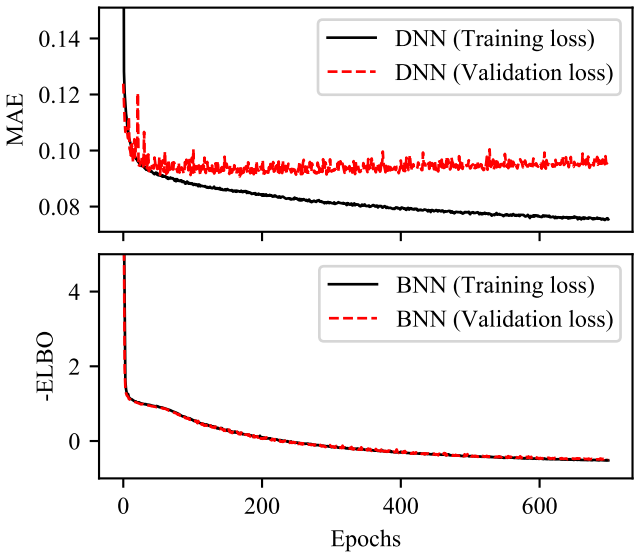

Fig. 5.7. Graphical representation and comparison between the usual training behavior of (top) standard deterministic neural networks (DNNs) and (bottom) Bayesian neural networks (BNNs). Training and testing losses are plotted for both models over epochs. The automatic overfitting control featured by BNNs can be observed in the illustration.

episodes, i.e., the validation loss plateaus while the training loss keeps decreasing. One therefore needs to add regularizers and/or test on a validation set after each training epoch. In this study, an early stopping callback based on validation loss metrics is implemented in order to prevent potential overfitting. On contrary, a BNN does not overfit as it contains a regularizer by default as explained in the previous section, and a separate validation dataset is no longer needed, and the early stopping callback in BNN monitors the training loss.

The prior weights' distributions of the Bayesian neural networks are assigned to follow a multivariate standard normal distribution. Since the BNN needs to minimize the negative ELBO, i.e., the sum of KL divergence and negative log-likelihood as shown in Equation (5.4), the built *DenseFlipout* layers add the KL divergence between the posteriors and their respective priors to the specified loss function. The Bayesian model is then trained on 1024-sample data batches applying the *Adam* optimizer. On a worth-noting remark, we used a very small learning rate to train BNNs in this work since otherwise, the network tends to converge to a local minimum owing to the stochastic nature of the model combined with the complexity of the dataset.

The training is established for 2000 epochs but the BNN early stopping criteria demands to stop if there is no significant improvement in the training loss. The evolution of the weights can also be tracked while the Bayesian neural network is being trained. Figure 5.8 illustrates the evolution of the mean, the standard deviation and the coefficient of variation of a model bias in the neural network. Evidently, rapid reduction of the uncertainty can be observed as the network learns from the dataset.

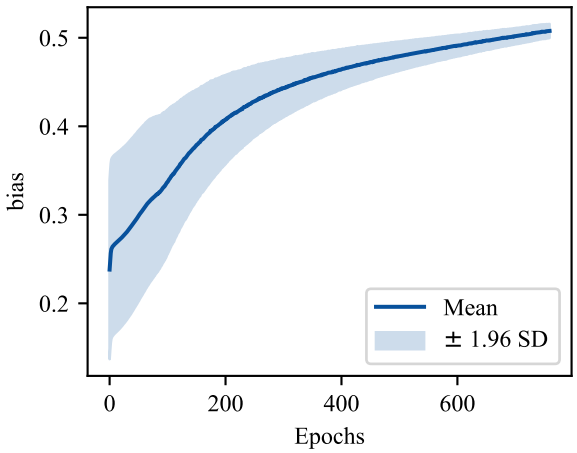

(a) Illustration representing the bias' mean with a continuous dark blue line and confidence intervals (± 1.96 SD) with a light blue color shade.

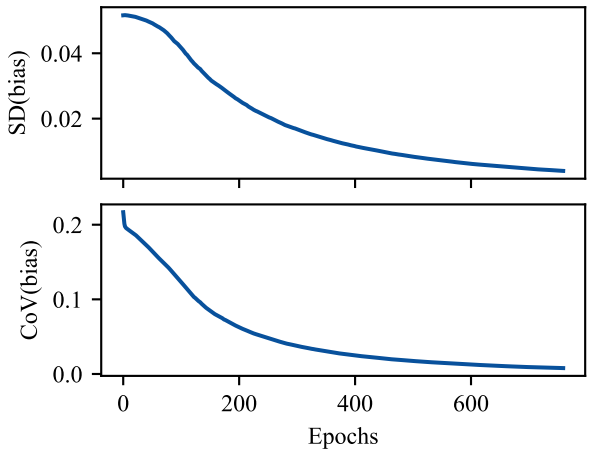

(b) [Top] Bias' standard deviation, SD(bias), represented over training epochs. [Bottom] Bias' coefficient of variation, CoV(bias), plotted during the course of the training stage.

Fig. 5.8. Evolution of a specific bias from the neural network over training epochs. The reduction of model uncertainty can be appreciated by observing the plunge of the bias' coefficient of variation (CoV) over the course of the training task.

The comparison of BNN models is shown in Figure 5.6b. The expected log-likelihood on the test dataset is computed from 10,000 forward runs according to Equation (5.8). Similar to the deterministic case, the model which maximizes the expected log-likelihood is the one with all available information, however, the model fed with reduced information - SCADA + accelerometer (LAT-077) - also provides satisfactory results. Furthermore, since both deterministic and probabilistic approaches provide consistent results, one can perform the sensor sensitivity analysis only in a deterministic manner, if preferred, without particular interest on model uncertainty.

Whereas the optimization of sensor placement can be application specific [173], here we aimed to maximize the model performance, yet keeping the required instrumentation as minimal as possible to facilitate farm-wide deployment. Therefore, based on the presented results, we selected the SCADA + accelerometer (LAT-077) combination for subsequent steps.

5.4.5 Fleet-leader's virtual monitoring model

Following the process undertaken for the selection of a reduced set of input monitoring signals, we further investigate the uncertainty associated with the virtual load monitoring model with respect to specific data collection periods. For each conducted assessment, BNNs featuring the same architecture are trained and fed with data gathered throughout a certain period, i.e., from 3 to 24 months. Note that a test set randomly sampled from the longest available period, i.e., 25% percentage of 'fleet-leader' full dataset, is used as reference to fairly evaluate all tested models. More specifically, Figure 5.9 represents the mean of the model's performance over the test dataset and as a function of 3-month data collection periods. By examining the separation between tick values along the x-axis, one can notice the scarcity of monitoring data during certain periods, which might be potentially associated with the inactivity of certain sensors.

In order to quantify the performance of the analyzed BNN models, two metrics are reported: the expected log-likelihood of the DEM, $\mathbb{E}[\mathcal{L}(DEM)]$, and the BNN's model uncertainty. The representation of the latter can be seen in Figure 5.9b, which is computed over 10,000 forward model runs. For better interpretability, the model uncertainty, usually indicated by the standard deviation of the predicted means, $\mathrm{SD}(\mu_{DEM})$, is normalized with respect to the expectation, $\mathbb{E}[\mu_{DEM}]$, as:

$$\mathrm{CoV}(\mu_{DEM}) = \frac{\mathrm{SD}(\mu_{DEM})}{\mathbb{E}[\mu_{DEM}]}, \tag{5.19}$$

where $\mathrm{CoV}(\mu_{DEM})$ corresponds to (normalized) model uncertainty, and $\mathbb{E}[\mu_{DEM}]$ and $\mathrm{SD}(\mu_{DEM})$ are computed according to Equations (5.14) and (5.13), respectively. Note that the standard deviation, SD(.), is equal to the square root of the variance, $\mathbb{V}(.)$. Even if the model uncertainty, i.e., $\mathrm{CoV}(\mu_{DEM})$, is quantified from BNN predicted results, it does not directly assess the accuracy of the generated load predictions with respect to DEM_{tl} and DEM_{tn} (labels) measurements. The considered BNN models are, therefore, additionally evaluated with respect to the expected log-likelihood of load measurements, $\mathbb{E}[\mathcal{L}(DEM)]$, computed according to Equation (5.8) over

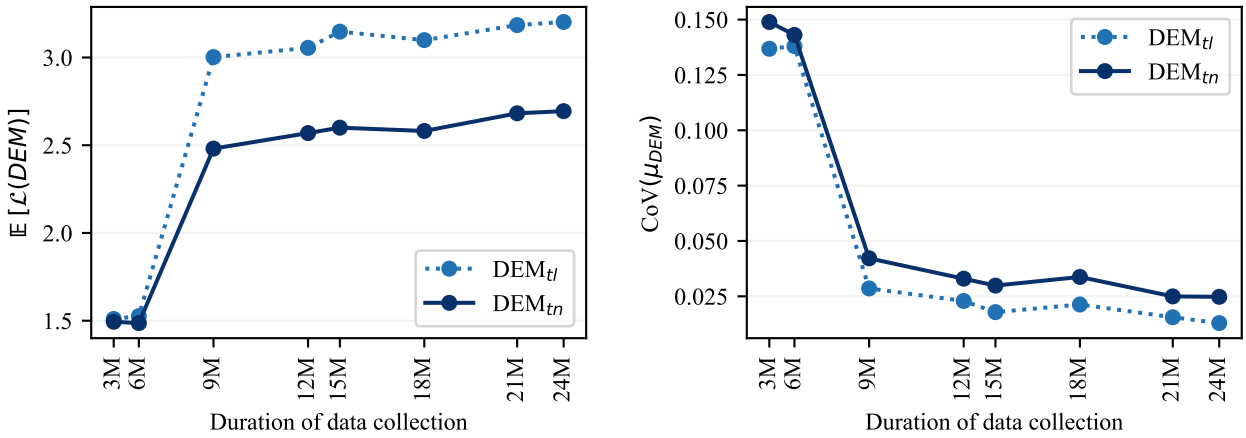

(a) Expected log-likelihood of the DEM, $\mathbb{E}[\mathcal{L}(DEM)]$, represented as a function of the data collection period.

(b) Normalized model uncertainty $\text{CoV}(\mu_{DEM})$ represented as a function of the data collection period.

Fig. 5.9. Illustration of the virtual monitoring model performance over specific data collection periods. It can be noticed that the amount of data collected differs for each period, e.g., the information retrieved over the course of the second trimester is scarce. Note that the plots indicate the mean of the model's performance over the test dataset.

also 10,000 forward model runs, and the results are plotted in Figure 5.9a. A similar trend in the performance reported by both metrics can be observed in the figure.

In general, BNNs (and other deep learning models) will benefit from additional training data, especially if the information is collected for regions where previously available training data was limited. The amount of data required to establish a robust BNN is case-dependent as it is influenced by the number of neurons considered and the complexity of the inherent physical process, among others. For the specific case of offshore wind turbines, representative environmental and operational data can be collected within a short term ($\sim$ 1-2 years) [72, 147]. As shown in Figure 5.9, the reduction of model uncertainty reported by the BNN steadily decreases over 12 months, where the model uncertainty reaches a stagnation point. One can thus conclude that, for this application, a robust BNN can be trained from 1-2 years of training data. Even if enough representative data has been collected, wind turbine dynamics might change at some point in the operational life, in which case, the BNN will automatically indicate to the user (e.g., operator) that additional data might be required. More specifically, an increased BNN's model uncertainty reported by the fleet-leader might suggest that the initially trained model is no longer adequate and strain gauges should be re-installed and the BNN model should re-trained. Based on the reported findings, the BNN model is trained over the full training dataset, i.e., data collection spanning over 24 months, and the resulting BNN virtual model is then deployed to other wind turbines in the subsequent farm-wide monitoring study.

5.4.6 Farm-wide deployment of virtual load monitoring model

Once the BNN load model has been trained based on the monitoring output signals (labels) collected from the fleet-leader wind turbine, the resulting BNN model is deployed in order to predict the loads of the other two wind turbines in the same wind farm, denoted in the text as 'MP01' and 'MP02'. Following the same monitoring setup as the fleet-leader, the input standard monitoring signals fed to the network during the deployment stage correspond to the combination of SCADA and acceleration data collected at the level LAT-077 (top accelerometer), and the virtual load monitoring model predicts the damage equivalent moments at the level LAT-017. Throughout the study, the damage equivalent moments predicted for every 10-minute collected input data are randomly sampled, thus enabling an uncertainty quantification of the retrieved load predictions. Note that even if the predictions generated by the model are studied and deployed to only two wind turbines in this investigation, the proposed framework can be easily deployed to the whole wind farm.

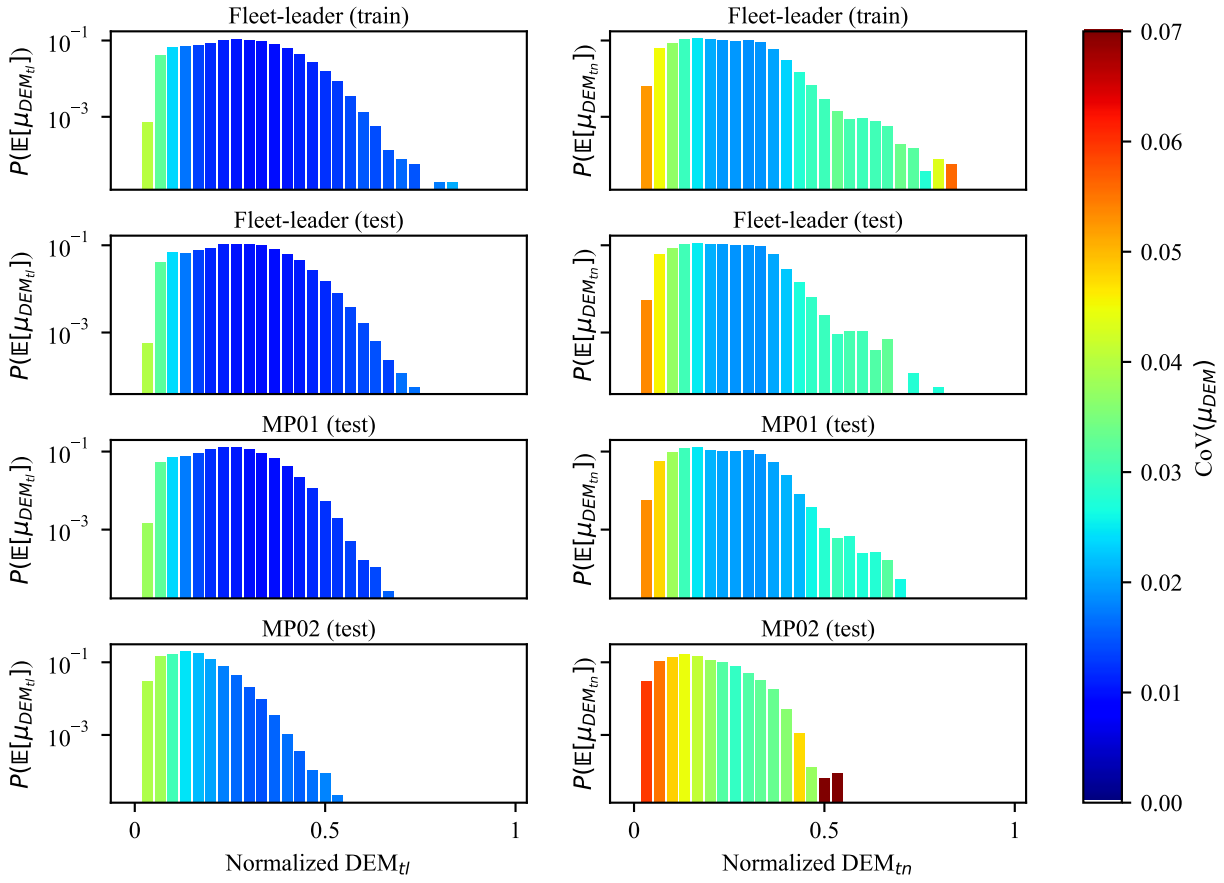

Fig. 5.10. Load predictions generated by the Bayesian neural networks at the deployment stage for all analyzed offshore wind turbines. The retrieved expected damage equivalent moments, $\mathbb{E}[\mu_{DEM}]$, are classified into discrete bins colored according to their associated model uncertainty $\mathrm{CoV}(\mu_{DEM})$. The height of each bar represents its probability and the color intensity indicates its associated model uncertainty.

To observe the load predictions generated by the BNNs, Figure 5.10 illustrates the expected load indicator $\mathbb{E}[\mu_{DEM}]$ and its respective (normalized) model uncertainty $\text{CoV}(\mu_{DEM})$ for all investigated turbines. Specifically, the expected load indicators $\mathbb{E}[\mu_{DEM}]$ are allocated into DEM bins and normalized by the total number of samples, obtaining, therefore, a probability distribution over the damage equivalent moment. Note that $\mathbb{E}[\mu_{DEM}]$ is also equal to $\mathbb{E}[D\hat{E}M]$. The model uncertainty associated with each bin is then computed according to Equation (5.19). As one can observe in the figure, the damage distribution is similar between the fleet-leader and MP01, resulting in higher probabilities in the low damage region (i.e., normalized $DEM < 0.5$) and smaller probabilities for medium to severe damage regions (i.e., normalized $DEM > 0.5$). However, MP02 rarely experiences medium to severe damage and many test samples are classified in the low damage region. Correspondingly, the model performance metric $\text{CoV}(\mu_{DEM})$ also indicates a good agreement between the fleet-leader's training and testing, as well as MP01's load predictions. The model in general predicts more accurately DEM_{tl} loads than DEM_{tn}, achieving the lowest uncertainty in the medium DEM_{tl} region of the fleet-leader and MP01 turbines. The model uncertainties reported for the case of MP02 are visibly higher than other turbines, announced in the illustration by a darker red color.

The showcased farm-wide load indicators and model uncertainty results are only based on the outputs retrieved from the deployed BNN, yet the accuracy of the predicted values with respect to the labels is not explicitly considered. Since MP01 and MP02 turbines have also been equipped with strain gauges during the monitoring campaign, the measurements (labels) are available to further analyze the obtained results. Figure 5.11 summarizes the model uncertainty for all considered offshore wind turbines, compared against the expected log-likelihood of the measured load indicator given the BNN outputs. Both performance metrics are represented by box plots that span over the interquartile range (IQR), i.e., between 25th and 75th percentiles, along with whiskers that extend up to 2.5th and 97.5th percentiles, respectively. Unsurprisingly, the expected log-likelihood outcomes are in agreement with the BNN-provided model uncertainties. MP02 wind turbine's results are characterized with higher model uncertainty as well as lower expected log-likelihood than for the other turbines, thus indicating potential conflicts with the generated predictions, e.g., the input monitoring data used during the training stage of the BNN might significantly differ from the input monitoring data collected for wind turbine MP02.

To further clarify the potential differences between the input monitoring data used during the training of the fleet-leader and the input monitoring data available for the other wind turbines at the deployment stage, the minimum Euclidean distance, r_{min}, of the corresponding input test data, $\boldsymbol{x}_{test}$, with respect to the training dataset, $\mathbf{X}$, is also quantified:

$$r_{min}(\boldsymbol{x}_{test}, \mathbf{X}) = \min_{\boldsymbol{x} \in \mathbf{X}} ||(\boldsymbol{x}_{test} - \boldsymbol{x})||_2, \tag{5.20}$$

where $\boldsymbol{x}, \boldsymbol{x}_{test} \in \mathbb{R}^M, \mathbf{X} \in \mathbb{R}^{M \times N}$, M and N stand for the dimension of input variables and the total number of training samples. $\mathbf{X}$ represents the matrix of input variables for all training data, $\boldsymbol{x}$ and $\boldsymbol{x}_{test}$ indicate the vectors of input variables for each training and testing sample,

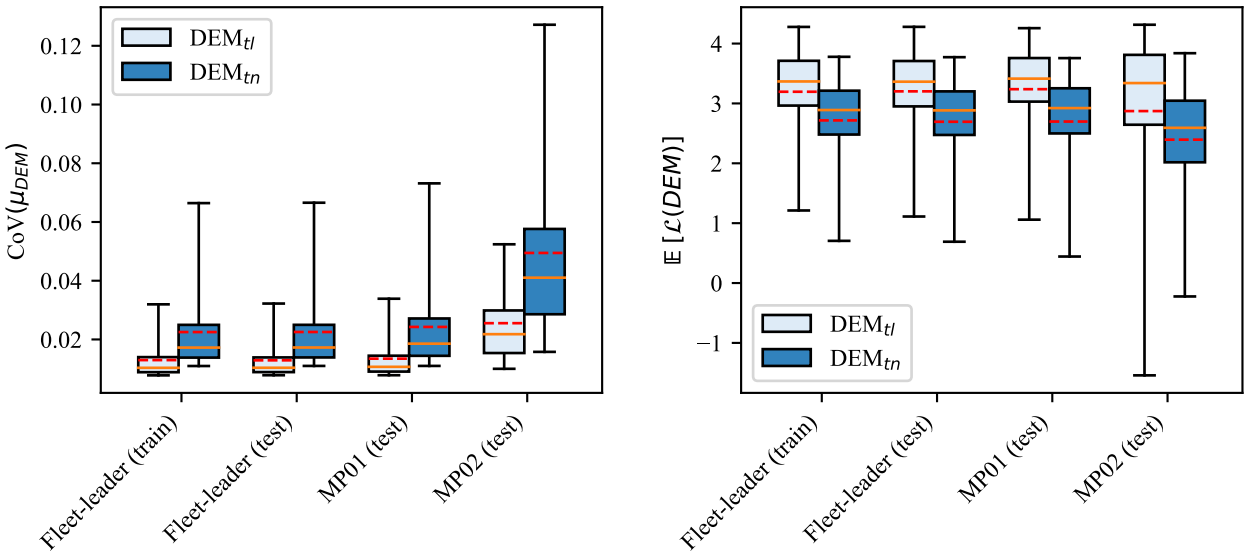

(a) Model performance metric: Normalized model uncertainty announced by the implemented BNN, CoV(μ_{DEM}).

(b) Model performance metric: Expected log-likelihood of the measured damage equivalent load, $\mathbb{E}[\mathcal{L}(DEM)]$.

Fig. 5.11. Representation of BNN's model performance for farm-wide load prediction. In particular, the model uncertainty is reported for the fleet-leader (both train and test datasets), MP01, and MP02 offshore wind turbines. In the figure, the orange line and the red-dotted line represent, respectively, the median and mean values of CoV(μ_{DEM}) over the corresponding dataset, and the boxes span between 25th and 75th percentiles, whereas the whiskers extend up to 2.5th and 97.5th percentiles.

and r_{min} denotes the minimum Euclidean distance from each test sample to its nearest training sample. A high r_{min} value implies that the corresponding test sample is far from the training region and, consequently, the predicted results rendered by the model might be inaccurate or highly uncertain. As one can observe in Figure 5.12, MP01 turbine's input monitoring test dataset is in good agreement with the fleet-leader's train dataset, also characterized with similar Euclidean distances when compared with the fleet-leader test dataset. On the other hand, the observed high Euclidean distances with a wider spread over the test set demonstrate that MP02's input test dataset substantially differs from the training dataset. Intrinsically, a BNN-based virtual monitoring scheme adjusts and reports higher model uncertainty, thus detecting potential conflicts that might emerge when input monitoring data at the deployment stage corresponds to unexplored data during the training stage.

In particular, the lack of correspondence between training and MP02 testing datasets can be explained by their divergent structural dynamics behavior. To better visualize this, the marginal probability distributions corresponding to each input variable are compared for all tested turbines, and the results are presented in Appendix B. Whereas SCADA input signals are fairly consistent, acceleration data clearly differ in MP02 wind turbine, as shown in Figure B1.

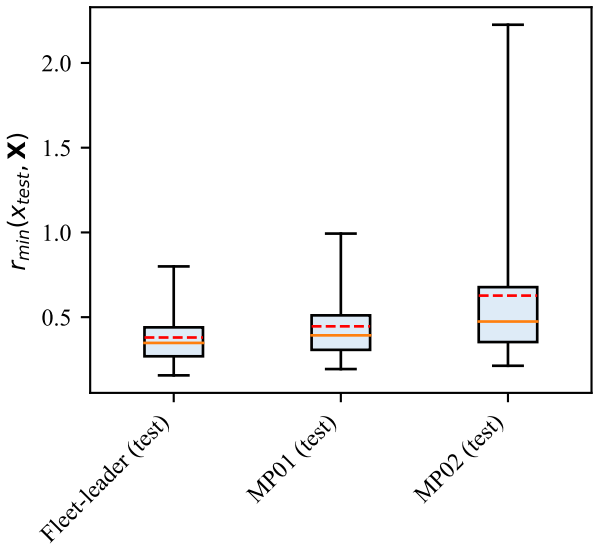

Fig. 5.12. Representation of the minimum Euclidean distance from each wind turbine's input test dataset to the fleet-leader's input training dataset. The minimum Euclidean distances are plotted for the fleet-leader (test dataset), MP01, and MP02 offshore wind turbines. In the figure, the orange line and the red-dotted line represent, respectively, the median and mean values of $r_{min}(x_{test}, \boldsymbol{X})$ distances over their corresponding dataset, and the boxes span between 25th and 75th percentiles, whereas the whiskers extend up to 2.5th and 97.5th percentiles.

5.4.7 Comparative study between DNNs and epistemic BNNs

We have hitherto focused on the BNN model in which both aleatory and model uncertainties are encapsulated. This fully probabilistic BNN model is not directly comparable to the DNN since the former yields a probabilistic load estimate while the latter only generates a point estimate. In order to objectively compare BNNs and DNNs in terms of error point estimates, an epistemic BNN (i.e., the variance of the prediction output is intentionally set up to 0) is trained based on the fleet-leader dataset and deployed to other wind turbines. In particular, this BNN modality quantifies the model uncertainty of the predicted load mean, yet the aleatory uncertainty is disregarded, as previously explained in the theoretical section. The outcome of the comparative analysis is represented in Figure 5.13, comparing the errors associated with DNN's deterministic outputs, $D\hat{E}M$, and BNN's predictions, $\mathbb{E}[D\hat{E}M]$, quantified following Equation (5.18). Note that the networks are trained only with the fleet-leader data and directly deployed to MP01 and MP02 without fine-tuning. The figure reveals that the epistemic BNN yields slightly more accurate point load estimates than its DNNs counterparts, except for MP02 turbine. Interestingly, the epistemic BNN also reports higher model uncertainty for MP02, illustrated in Figure 5.13 with blue-colored bars, agreeing with the computed test accuracy (MAE).

One should also keep in mind that the modeled epistemic BNN features fewer neurons than the considered DNN, potentially reducing its generalization capabilities. A thorough description of the implemented neural network architectures and training parameters is presented in

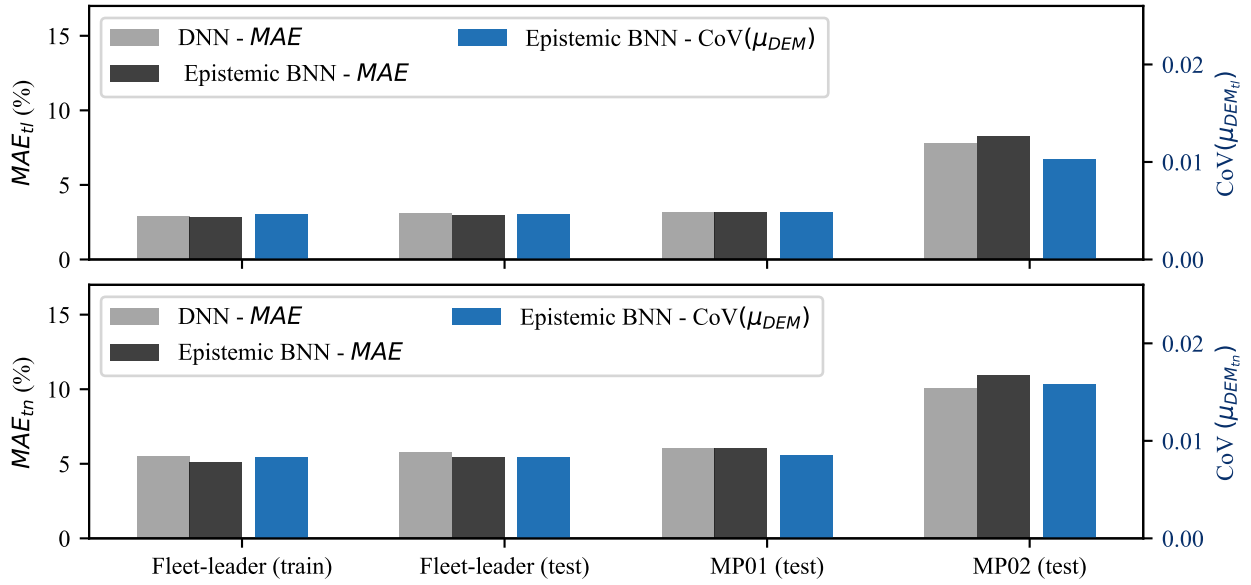

Fig. 5.13. Prediction error associated with DNN and epistemic BNN predictions. The mean absolute error (MAE) corresponding to DNNs and epistemic BNNs is represented with light and dark grey bars, respectively. Additionally, model uncertainty metrics, CoV(μ_{DEM}), reported by the epistemic BNN (without the need for ground truth labels) are represented with blue bars.

Appendix A. Whereas the generalization capabilities of different neural network architectures and hyperparameters are case-dependent, the proposed BNN-based virtual monitoring method constitutes a general framework for detecting potential load inaccuracies without the explicit need for a target (load measurement), which is especially relevant when dealing with farm-wide monitoring applications. In this specific case study, measurements are available for all analyzed turbines, yet this will most likely not be the case in practical scenarios as it is economically unfeasible to fully instrument all turbines. While the DNNs do not explicitly report model uncertainty estimates, BNNs are able to yield consistent model uncertainty information without the need for ground truth labels. This is confirmed by the reported results, where potential high prediction inaccuracy is automatically announced for the wind turbine MP02.

5.4.8 BNNs model uncertainty

BNNs are able to automatically announce the model uncertainty associated with the generated predictions independently of whether the output predictions are modeled as a probability distribution or as a point estimate. As mentioned in the theoretical section, a BNN that can capture both epistemic and aleatory uncertainty information would naturally require more training data. To analyze this, a further investigation has been conducted, comparing the model uncertainty reported by (i) a BNN that captures both aleatory and epistemic uncertainty components and (ii) an epistemic BNN (i.e., the variance of the prediction output is intentionally set up to 0). The architecture and training hyperparameters are the same for both tested BNNs. Further details can be found in Appendix A.

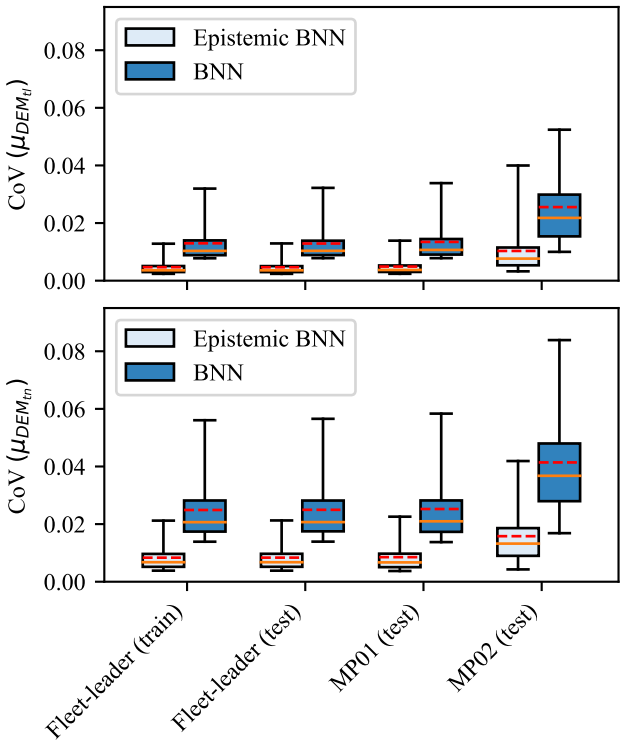

Fig. 5.14. Model uncertainty associated with the load predictions generated by the investigated Bayesian neural networks. BNNs capturing only epistemic uncertainties are colored in light blue. Spreading over each wind turbine dataset, box plots represent the corresponding model uncertainty, CoV(μ_{DEM}), within the interquartile range, with whiskers that span from 2.5th to 97.5th percentiles. Additionally, the median and mean are indicated with orange and red-dotted lines, respectively.

The results are represented in Figure 5.14, indicating with box plots the spread in model uncertainty, i.e., CoV(μ_{DEM}), over the training and testing datasets corresponding to all tested wind turbines. A similar trend can be observed for the analyzed BNNs, reporting a higher model uncertainty associated with the predictions generated for wind turbine MP02, thus verifying farm-wide applicability of both variants. A further inspection of Figure 5.14 reveals that the aleatoric BNN yields higher model uncertainty compared to its epistemic counterpart. This is justified by the fact that the training of a model that can predict a probability distribution given a certain input combination requires more training data, as mentioned previously, and especially considering the input is high-dimensional in this case, i.e., 13 input variables. Note that the prediction accuracy cannot be directly compared here because one of the tested BNN models generates an output probability distribution, whereas the epistemic BNN only provides the prediction of a point estimate.

5.5 Conclusions

In this paper, we propose and examine the effectiveness of a data-based virtual load monitoring framework formulated with Bayesian neural networks for 'fleet-leader' farm-wide load monitoring, i.e., a data-based model is trained with data collected from a fully monitored wind turbine, and once the training task is completed, the resulting model is deployed, thus yielding load predictions for other non-fully monitored wind turbines. Within the investigation, we carefully assess relevant advantageous properties offered by Bayesian neural networks (BNNs), e.g., uncertainty quantification, automatic overfitting regulation, with respect to standard deterministic neural networks (DNNs), and we test the proposed virtual load monitoring framework through an experimental monitoring campaign conducted in an existing offshore wind farm.

The results observed throughout the conducted experimental campaign reveal that, by reporting an epistemic model uncertainty indicator, and as opposed to their standard DNNs counterparts, BNN-based virtual monitoring approaches are able to intrinsically identify potential conflicts with the generated load predictions, providing therefore an informative proxy for controlling the accuracy of the deployed farm-wide virtual monitoring model. For instance, a BNN-based virtual monitoring model, whether aleatory uncertainty is inclusively modeled or not, will automatically report high model uncertainty during its deployment if the input monitoring data features outliers, i.e., unexplored data with respect to the data fed to the model during the training stage. Besides their useful uncertainty management capabilities, BNNs overall training task, even if more computationally demanding than for DNNs, is automatically regulated by Bayesian inference principles, thus avoiding the risk of overfitting and eluding the need of a separate cross-validation dataset.

From all available standard monitoring signals, a reduced set of informative input monitoring signals has been selected in this work by quantifying the generalization error resulting from each considered monitoring setup. The selection process relies, in this case, on already collected input data from the installed load monitoring system. Benefiting from BNNs' internal properties in terms of uncertainty quantification, we motivate further research efforts toward farm-wide sensor placement studies capable of allocating monitoring installation and maintenance actions by following optimal adaptive management policies, e.g., asset management strategies identified via Markov decision processes and/or deep reinforcement learning methods [27, 29, 38, 43].

In a more specific note, we also encourage the exploration of sophisticated principled metrics that can be computed during the testing stage of BNN-based virtual monitoring models, and we suggest, for instance, the investigation of entropy-based metrics [169]. Moreover, a detailed survey of data-based probabilistic virtual monitoring models is also recommended, e.g., comparing kernel- and neural network -based methods, which could potentially be complemented with a thorough uncertainty decomposition assessment.

Authorship contribution statement

Hlaing, N.: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing - Original draft, Writing - Review & editing, Visualization. **Morato, P. G.:** Conceptualization, Validation, Formal analysis, Writing - Review & editing. **de Nolasco Santos, F.:** Validation, Formal analysis, Data curation, Writing - Review & editing, Visualization. **Weijtjens, W.:** Validation, Formal analysis, Data curation, Writing - Review & editing. **Devriendt, C.:** Data curation, Supervision. **Rigo P.:** Supervision, Project administration, Funding acquisition.

Data availability

Due to its proprietary nature and confidentiality concerns, supporting data of this research cannot be made publicly available.

Acknowledgements

This research is funded by the Belgian Energy Transition Fund (FPS Economy) through PhairywinD (https://www.phairywind.be) and MaxWind projects. We further acknowledge OWI-lab (https://www.owi-lab.be) for supporting with the wind farm dataset to develop this paper.

Chapter 6

Quantifying the value of virtual monitoring in life-cycle management of offshore wind support structures

Paper Hlaing, N., Morato, P. G. and Rigo, P. Quantifying the Value of Virtual Monitoring in Life-cycle Management of Offshore Wind Support structures. *Engineering Structures*, under internal review.

Abstract This paper focuses on life-cycle management planning for offshore wind support structures by leveraging data-based virtual monitoring. Bayesian neural networks are proposed as virtual sensors, providing useful information and predictive uncertainty when monitoring systems are no longer available. The probabilistic virtual monitoring model is integrated with the decision-making framework in order to improve the life-cycle management policies of deteriorating structures. Featuring high-dimensional state, action and observation spaces, the decision-making problem is formulated as a decentralized partially observable Markov decision process (POMDP) and solved via multi-agent reinforcement learning (MARL) algorithms, advising decisions on when and where to inspect, monitor, or maintain. Life-cycle management planning is performed both with and without virtual monitoring, quantifying the value of virtual information. The outcomes of this research not only verify the added benefit of virtual sensors in decision-making under uncertainty, but also demonstrate the effectiveness of MARL in identifying life-cycle management strategies for complex engineering systems.

6.1 Introduction

Structures and infrastructures, including bridges and offshore wind turbines, are subject to various gradual deterioration mechanisms such as fatigue and corrosion, which can significantly degrade the performance or even trigger failure. To effectively control the risk of structural failure and mitigate its consequences, it becomes crucial to strategically allocate maintenance actions throughout the lifetime. Given the stochastic nature of the deterioration processes mentioned earlier, planning maintenance for these infrastructures requires a comprehensive approach involving probabilistic analysis and a decision-making framework under uncertainty.

Structural health information can be collected through inspection and monitoring, reducing the uncertainties in deterioration estimations and improving the inspection and maintenance decisions. Whereas non-destructive inspection methods previously stood as primary source of information, the concept of structural health monitoring (SHM) has recently gained prominence thanks to the rapid advancement of sensor technologies [93]. SHM which involves the use of advanced sensing technologies to continuously monitor the condition of structures offers an array of benefits, ranging from early detection of structural defects to enabling decision-makers to make timely and informed decisions [96, 143].

However, collecting additional information through either inspections or SHM not only incurs a significant cost but also concerns with measurement uncertainty. One therefore needs a value of information (VoI) analysis to quantitatively assess the potential benefit prior to execution. Straub has presented a formal VoI framework from the structural reliability perspective quantifying the value of perfect as well as imperfect information [174]. The frameworks for quantifying the value of structural health monitoring have also been reported [36, 47, 175]. Generally, uncertainties in the deterioration estimations are reduced upon gathering of additional data, thus having an effect on maintenance strategies. One can quantify the net value of information as the difference between: the achieved reduction of expected costs due to additional information, and the cost of information collection. In some I&M planning methods, the expense of collecting information via inspections or structural health monitoring is already inherently considered within the cost modeling of the decision-making framework [36], and the net VoI is directly obtained as the difference in the expected costs.

Leveraging the advancement of artificial intelligence and machine learning tools, virtual sensing has increasingly used to enhance the amount of information collected. Virtual sensing which refers to the ability to gather information about a specific system without directly using physical sensors, is particularly useful in situations where physical sensors are expensive or difficult to deploy [59, 60]. Even if not more accurate than physical sensors, virtual sensing tools can provide insights on a system's performance at no cost. Virtual sensing models have been widely implemented in offshore wind structures [68, 70, 79, 80] and other engineering applications [65, 84]. Quantifying the value of virtual sensing information, however, imposes some challenges such as to obtain the uncertainty of the virtual sensors and consequently the quality of virtual observations. Even in the posterior analysis, the uncertainty of the already observed virtual information may be unknown in some cases, if the target labels cannot be collected.

In this paper, we adopt a probabilistic virtual sensing model [176] and integrate it in life-cycle management framework, with the objective of strategically allocating inspection, monitoring and maintenance actions. Virtual monitoring based decision-making methods can be found in the literature [97, 98, 99], assuming the existence of digital twin as virtual sensors since the beginning of decision horizon. Unlike the existing methodologies, our proposed method relies on the pre-posterior decision analysis to optimally allocate virtual monitoring options over the lifetime

of structures. In particular, the life-cycle management framework is formulated considering the possibility of developing a virtual sensor after data collection, and continuous information retrieval afterwards. Beliefs about deterioration states are updated with not only inspection outcomes but also physical or virtual monitoring observations. The uncertainty of virtual monitoring is stochastically propagated over the decision horizon, reflecting the variations in the structural dynamics. Optimization of the decision-making problem adheres to the principles of partially observable Markov decision processes (POMDPs) while the computational complexity of identifying optimal policies, is alleviated by implementing a multi-agent reinforcement learning approach. We apply the proposed methodology to the case study of managing an offshore wind turbine support structure consisting of multiple components. We thoroughly analyze how the optimal policies are altered by including virtual monitoring, and the benefit of virtual monitoring is quantified relying on the value of information principles.

6.2 Virtual sensing

Sensors are often installed on engineering systems to monitor certain physical quantities, performance indices or operational information. In many practical applications, it might, however, be difficult or even impossible to obtain direct measurements of some quantities due to challenges in terms of accessibility and/or sensor installation and maintenance expenses, among others. Virtual sensors are therefore developed and used to estimate or infer unavailable information [60]. Virtual sensing has been adopted in several engineering applications, including wireless networks [59], intelligent buildings [58], and structural health monitoring [70, 79, 176].

A virtual sensor consists of a software or mathematical model, whether physics-driven (white-box), data-driven (black-box) or combined (grey-box), which relies on measurable data to indirectly estimate unmeasured information. Physics-driven virtual sensors, often developed as numerical models, involve the use of governing equations and physics laws and occasional calibration of model parameters with respect to in-situ observations. Although physics-based approaches require limited observation data, the development and implementation in complex engineering systems demand specific knowledge of model parameters that are usually hard to retrieve in real case scenarios.

Data-driven virtual sensors, often including regression models or neural networks, learn to obtain empirical relationships between input variables and target responses. Data-based models are flexible and can efficiently identify meaningful patterns in the dataset, especially being useful to model complex behavior of engineering structures. However, they require substantial amount of data for effective learning, and their 'black-box' nature might lack interpretability, making it challenging to explain the rationale behind predictions or estimations. Additionally, ensuring generalization beyond the training data and adapting to different operational conditions or system dynamics can be challenging.

Physics-informed grey-box models often bridge the gap between traditional physics-based models and the flexibility of machine learning models, offering a powerful tool for estimating or inferring unmeasured information within complex engineering systems [76, 87]. These models leverage domain-specific knowledge by incorporating the underlying physics laws or governing equations into the learning process, ensuring that the models are consistent with the known behavior of the system, offering an advantage in situations with limited available data [148].

The most critical aspect of leveraging virtual sensing information for decision-making lies in effectively quantifying the uncertainty inherent in virtual models, as emphasized before. Virtual sensing methods in this context can be categorized into deterministic models and probabilistic models. Deterministic models which operate with fixed model parameters generally lack explicit indications of the uncertainty associated with their predictions, relying primarily on physical sensor data ('ground truths') for verification and validation purposes [92]. However, when it is difficult to acquire target labels for evaluating virtual sensing information's quality, conservative uncertainty assumptions may be employed, often based on engineering judgment.

On the other hand, probabilistic approaches whose parameters are represented by random variables, can intrinsically provide an indication of the uncertainty contained in the output predictions. Such models often include stochastic simulators, e.g., stochastic finite elements methods [177], and probabilistic machine learning models, e.g., Gaussian processes, Bayesian neural networks. Gaussian process models are frequently used as virtual monitoring tools in small to medium dimensional cases [71, 82, 149]. Probabilistic neural networks such as BNNs, deep ensembles, etc., are capable of handling high dimensional and complex problems especially when large amount of data is available, and are frequently used in engineering applications [154, 176, 178].

6.3 Value of information in decision-making

This section briefly describes a theoretical framework for quantifying the value of information in decision-making. A more detailed and comprehensive explanation can be found in [174]. The uncertainties in the underlying process can be characterized by a set of random variables $\mathbf{X}$, each of which follows a specific probability distribution. In decision analysis under uncertainty, the objective is to identify actions that minimize the expected cost $\mathbb{E}(C)$. When no information is available, the decision-making is purely based on the prior knowledge of uncertain variables $\mathbf{X}$. This optimization problem is referred to as a prior decision analysis:

$$a_{opt} = \text{argmin}_a \mathbb{E}[c(a, \mathbf{X})] = \text{argmin}_a \int_{\mathbf{X}} c(a, \mathbf{x}) f(\mathbf{x}) d\mathbf{x}. \tag{6.1}$$

The random variables are discretized into a finite number of discrete states $s \in \mathcal{S}$, then the optimization problem is:

$$a_{opt} = \text{argmin}_a \sum_{s \in \mathcal{S}} c_s(a) p(s), \tag{6.2}$$

where $p(s)$ is the probability of being in the state s and $c_s(a)$ is the cost of action a if taken in the state s. The expected prior cost is also associated with optimal decision:

$$\mathbb{E}[C_{prior}] = \sum_{s \in \mathcal{S}} c_s(a_{opt}) p(s). \tag{6.3}$$

Optionally, additional information can be gathered before making the action decisions in order to reduce the uncertainties in random variables $\mathbf{X}$. If it is assumed that the information is already collected, the prior knowledge on $\mathbf{X}$ can be updated using the observation o via Bayesian inference:

$$p(s \mid o) = \frac{P(o \mid s) p(s)}{\sum_{s \in \mathcal{S}} P(o \mid s) p(s)}, \tag{6.4}$$

where the likelihood $P(o \mid s)$ is the probability of observing o given the state s. The decision optimization, which is now based on the updated information, is denoted as the posterior analysis, obtaining the optimal decision as:

$$a_{opt|o} = \text{argmin}_a \sum_{s \in \mathcal{S}} c_s(a) b(s \mid o). \tag{6.5}$$

The conditional value of information for a particular observation o is computed as the difference in the costs of prior and posterior optimal actions in the conditional belief $b(s \mid o)$, as follows:

$$CVoI_o = \sum_{s \in \mathcal{S}} c_s(a_{opt}) b(s \mid o) - \sum_{s \in \mathcal{S}} c_s(a_{opt|o}) b(s \mid o). \tag{6.6}$$

If the optimal decision remains the same before and after information collection, the CVoI will be equal to zero.

When additional information is collected, there are several possible outcomes, e.g., inspection outcomes can be binary: detection and no-detection, or multivariate: crack size measurement. One needs to consider all those possibilities in the VoI analysis ahead of information collection decision. The VoI is then the expected value of the CVoI with respect to all possible observations $VoI = \mathbb{E}[CVoI]$. Assuming a finite number of possible observations, the VoI is computed as:

$$\begin{aligned} VoI &= \sum_{o \in \mathcal{O}} CVoI_o P(o) \\ &= \sum_{o \in \mathcal{O}} P(o) \left[\sum_{s \in \mathcal{S}} c_s(a_{opt}) b(s \mid o) - \sum_{s \in \mathcal{S}} c_s(a_{opt|o}) b(s \mid o) \right] \\ &= \mathbb{E}[C_{prior}] - \sum_{o \in \mathcal{O}} P(o) \sum_{s \in \mathcal{S}} c_s(a_{opt|o}) b(s \mid o). \end{aligned} \tag{6.7}$$

The second term of the equation is denoted as pre-posterior decision analysis, in which the decisions are planned by already taking into account all potential outcomes.

In addition, it is also useful to compute the net value of information (NVoI) which also includes the cost of additional information C_i,

$$NVoI = VoI - C_i. \tag{6.8}$$

If the NVoI is positive, the collected information provides an added value whereas the negative NVoI indicates that the information is more expensive than its added benefit.

Quantifying the value of information in life-cycle management planning involves sequential prior and pre-posterior decision analyses, i.e., decisions are made at multiple time steps. This imposes significant challenges since the computational cost is exponentially increasing with the number of decision steps. In traditional risk-based inspection and maintenance planning, this has lead to the identification of heuristic decision rules such as inspections which are planned at equidistant time intervals, or before exceeding a failure probability threshold, and maintenance actions being planned according to inspection outcomes. However, since a sub-space of the vast policy space has been evaluated, heuristics-based optimization can be sub-optimal, resulting in the inaccurate estimation of the VoI.

To address this limitation, more advanced decision-making frameworks have been explored. Such approaches involve the use of Markov decision processes (MDPs) [18, 26, 27, 179] or reinforcement learning techniques [38, 175]. In the former approach, the life-cycle management planning problem is formulated as a partially observable Markov decision process (POMDP). The system state is characterized by the deterioration condition of the asset and the decision-maker's belief about this condition. At each decision point, the decision-maker selects an action (e.g., inspection, maintenance) based on the current belief, i.e., a probability distribution over the deterioration states. Since the belief is updated after taking an action or collecting an observation, the current belief is a sufficient statistic of all the past actions and observations. The objective is to derive a policy that dictates the optimal action at each decision point to maximize the expected reward (or minimize the expected cost) over the entire life-cycle. Dynamic programming algorithms, e.g., value iteration or policy iteration, or more sophisticated point-based solvers are then employed to identify this optimal policy. POMDP-generated policies are adaptive and assure optimality if exactly solved [29].

Reinforcement learning algorithms are adopted to approximate POMDP solution for life-cycle management planning of civil infrastructures such as bridges and wind turbines [44, 180, 181]. In this paradigm, each agent representing a specific component interacts with the environment (the deterioration of the asset in this context) and learns to make decisions that maximize the cumulative reward. The agent receives feedback in the form of rewards or penalties based on the outcomes of its actions. Through learning from these experiences, the agent refines its decision-making strategy over time. While reinforcement learning approach can handle sequential decision-making in high-dimensional state and action spaces, optimality is not guaranteed as the training may converge to a local optimum. However, I&M strategies obtained from state-of-the-art reinforcement learning algorithms such as DDMAC, QMIX, FACMAC, etc., are

demonstrated to outperform calendar-based or condition-based planning with substantial cost savings in the management of large infrastructure systems [40, 43, 49]. Furthermore, when life-cycle management planning consists of multiple information collection possibilities, e.g., inspection, sensor installation, and multiple maintenance actions, e.g., minor repair, major repair, replacement, etc., it becomes difficult to define heuristic patterns, especially for multi-component systems.

6.4 Virtual monitoring integrated life-cycle management framework

6.4.1 Decentralized POMDP description

Life-cycle management planning aims to minimize the expected sum of discounted costs, including inspections, sensor installations, repairs, and failure risk. In this section, we formulate the life-cycle planning of a multi-component system as a decentralized partially observable Markov decision process (Dec-POMDP), defined by a tuple $[\mathcal{S}, \mathcal{A}, \mathcal{O}, T, Z, R, n, \gamma]$. A detailed description of each tuple element is provided in the following sections. In essence, at every time step t of the decision horizon, n agents simultaneously choose an action while being in the state $s \in \mathcal{S}$ where $\mathcal{S}$ is the set of states. In an uncertain environment, the state s is not observable, and hence the agent only perceives the belief b(s), i.e., the probability distribution over the states. Each agent $l \in 1, ...n$ takes an action $a_l \in \mathcal{A}_l$ from the joint action space $\mathcal{A} = \mathcal{A}_1, ...\mathcal{A}_n$, thereby transitioning the belief state $b(s)$ according to the transition model $T(s', s, a) = P(s' \mid s, a)$. The agent then receives an imperfect observation $o \in \mathcal{O}$ with the probability $Z(o, s', a) = P(o \mid s', a)$, and the reward $R(s', s, a)$ discounted by a factor γ is collected by all agents. The optimal POMDP policy involves a sequence of actions that maximize the accumulated reward over the decision horizon.

States

Each agent observes its respective component's damage state and the current time step. However, since the damage state cannot be perfectly known, the agent instead perceives a probability distribution. The probability distribution of each component's damage condition transitions according to maintenance actions, and can also be updated based on the information collected from inspection, physical or virtual sensors.

When virtual monitoring option is integrated in life-cycle management, additional information related to virtual sensors need to be incorporated into the state space. In this work, we additionally consider two additional variables - the sensor status and the model uncertainty of a virtual sensor, which are both fully observable. The sensor status consists of three states- 'physical sensor', 'virtual sensor' and 'no sensor', in which the actual state is indicated by a

one-hot vector. The sensor status therefore informs whether there is a physical or virtual sensor from which the information can be retrieved.

It should be noted that the measurement uncertainty of physical sensors can be pre-collected from sensor manufacturers. The parameter describing the uncertainty of physical sensors ε_p is therefore deterministic, and not included in the state space. On the other hand, the model uncertainty is informed by a (probabilistic) virtual sensor during its deployment, i.e., only in the posterior case. In the pre-posterior analysis, a virtual sensor's uncertainty ε_v is considered as a random variable, and encoded in the state space $\mathcal{S}$.

Actions

The action space $\mathcal{A}$ includes not only maintenance actions such as 'perfect repair a_{PR}', or 'do nothing a_{DN}', but also decisions related to information collection, such as 'inspection a_I', 'install sensor a_S', etc. Moreover, maintenance- and observation-related decisions can also be combined as additional action possibilities [26]. For example, it is possible to install a sensor immediately after a maintenance action is performed, allowing data collection in the following year.

Observations

The observation space $\mathcal{O}$ is defined according to the inspection and monitoring techniques available. Non-destructive inspections usually provide only two observations - detection and no-detection of cracks. In some inspection techniques, direct measurement of cracks can also be obtained. Whereas monitoring systems output continuous data stream often in time series, a systematic post-processing is required to convert into discrete observations. For instance, rainflow counting can be used to convert load signals from a strain gauge into a load spectrum [147] or damage equivalent load [176].

Transition probabilities

A transition matrix for each maintenance action $a \in \mathcal{A}$ is defined as the probability of the component damage transitioning from the state $s_t \in \mathcal{S}$ to the state $s_{t+1} \in \mathcal{S}$. For example, the transition matrix of the action 'do nothing' $T(s_{t+1}, a_{DN}, s_t)$ follows the stochastic deterioration process of interest. The transition model for a 'perfect repair' action $T(s_{t+1}, a_{PR}, s_t)$ can be constructed such that the component with any belief $\mathbf{b}$ returns to its initial condition $\mathbf{b_0}$ [29]:

$$P(s_{t+1} \mid s_t, a_{PR}) = \begin{pmatrix} b_0(s_0) & b_0(s_1) & \cdots & b_0(s_k) \\ b_0(s_0) & b_0(s_1) & \cdots & b_0(s_k) \\ \vdots & \vdots & \ddots & \vdots \\ b_0(s_0) & b_0(s_1) & \cdots & b_0(s_k) \end{pmatrix}. \tag{6.9}$$

Despite using only two maintenance actions in this paper, other repair transition matrices can also be defined, e.g., the component damage after an 'imperfect repair' action will go back to a healthier state but not the initial one.

The transition matrices for the sensor status h are also separately defined for each action. Specifically, the sensor status indicates the presence of a 'physical sensor' once a sensor is installed. In the following years until the next sensor installation, the sensor status indicates the presence of a 'virtual sensor' unless a maintenance action is executed. In the latter case, it is assumed that the virtual sensing model no longer represents the structural component after repair, thus transitioning back to the 'no sensor' state. To preserve this dynamics, the transition matrices are modeled for 'install sensor', 'do nothing', 'inspection', and 'perfect repair' actions, respectively.

$$P(h_{t+1} \mid h_t, a_{DN}) = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}; \qquad P(h_{t+1} \mid h_t, a_I) = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix};$$
$$P(h_{t+1} \mid h_t, a_S) = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}; \qquad P(h_{t+1} \mid h_t, a_{PR}) = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix}. \tag{6.10}$$

On a remark, the lifetime of a physical sensor is modeled as only one year in the above setting. One can also consider a longer lifetime by increasing the dimension of the sensor status state and its transition matrices. Moreover, stochastic lifetime of physical sensors can also be modeled by changing the transition probabilities.

Observation probabilities

An observation matrix defines the probability of collecting an observation $o \in \mathcal{O}$ for the component damage being in state $s' \in \mathcal{S}$ after taking an action $a \in \mathcal{A}$. Observation outcomes are frequently in the form of binary indication or continuous indication. Non-destructive inspection techniques often provide detection or no-detection outcomes, of which the likelihood is quantified by means of probability of detection (PoD) curves [182]. The PoD is the probability of collecting a detection observation o_D as a function of the current state $s' \in \mathcal{S}$, that is $P(o_D \mid s') = PoD(s)$.

On the other hand, physical or virtual monitoring systems, and sometimes certain inspection techniques, provide continuous indication, e.g., measuring the loads and/or crack size. In this case, the likelihood of observations is modeled by a continuous distribution:

$$P(\mathbf{o} \mid s') \sim \mathcal{N}(s', \varepsilon), \tag{6.11}$$

where ε is the uncertainty of observation techniques. It should be noted that the observation space for continuous indications is usually discretized into a finite number of observed states for

efficient Bayesian inference [96]. In this work, the observations are also conditional on the sensor status, i.e., a monitoring-related observation is collected only if the sensor status indicates a physical or virtual sensor.

Rewards

In decentralized POMDP, each agent being responsible for a component takes actions $a \in \mathcal{A}$ over the decision horizon t_N. An immediate reward corresponding to the taken actions is collected by all agents at each time step, however the ultimate objective is to maximize the accumulated rewards, i.e., to minimize the expected life-cycle cost of the system:

$$C_{tot} = \sum_{t=0}^{t_N} \left[\gamma^t \left(R_{F_t} + C_{camp_t} + \sum_{l=1}^{N_c} \left\{ C_{ins_t}^{(l)} + C_{sens_t}^{(l)} + C_{rep_t}^{(l)} \right\} \right) \right]. \tag{6.12}$$

The system failure risk is defined as $R_F = C_F \cdot (P_{f_{sys},t+1} - P_{f_{sys},t})$, where $P_{f_{sys}}$ is the system failure probability and C_F is the associated consequences of a failure event. On the other hand, inspection, sensor installation and repair actions also incur costs, $C_{ins}, C_{sens}, C_{rep}$ respectively. Optionally, a global campaign cost C_{camp} can also be considered if at least one component in the system is inspected, monitored or repaired, plus a surplus, $C_{ins}, C_{sens}, C_{rep}$ per inspected/monitored/repaired component. γ is the discount factor which assigns higher importance to immediate decisions over long-term decisions. All I&M associated costs and failure risks are formally included in the Dec-POMDP framework as negative rewards.

6.4.2 POMDP dynamics

The algorithmic POMDP dynamics is presented hereafter, implementing the belief evolution via dynamic Bayesian networks [11]. Although the state space $\mathcal{S}$ is usually augmented to jointly model all the involved variables in classical POMDP configurations, we decouple and consider individual belief states for deterioration $\mathbf{b}_{d_t,q_t}$, deterioration rate $\mathbf{b}_{\tau_t}$, sensor status $\mathbf{b}_{h_t}$ and virtual sensor's uncertainty $\mathbf{b}_{\varepsilon_{vs,t}}$. It should be noted that physical or virtual sensors do not measure the damage d_t directly, but observe the underlying parameter q_t instead. Therefore, the deterioration belief state is in fact a joint distribution of d_t and q_t, also illustrated in Figure 6.1. The POMDP dynamics is also specified in conditional formations to alleviate the dimensionality and computational complexity. The deterioration rate $b(\tau_t)$ of each component transitions according to $P(\tau_{t+1}|\tau_t, a_t)$:

$$b(\tau_{t+1} \mid a_t) = \sum_{\tau_t \in \mathcal{S}_\tau} P(\tau_{t+1}|\tau_t, a_t) b(\tau_t). \tag{6.13}$$

The transition of the component deterioration $\mathbf{b}_{d_t,q_t}$ is also defined as:

$$b(d_{t+1}, q_{t+1} \mid a_t) = \sum_{d_t \in \mathcal{S}_d} \sum_{q_t \in \mathcal{S}_q} \sum_{\tau_{t+1} \in \mathcal{S}_\tau} P(d_{t+1}, q_{t+1}|d_t, q_t, \tau_{t+1}, a_t) b(d_t, q_t) b(\tau_{t+1}). \tag{6.14}$$

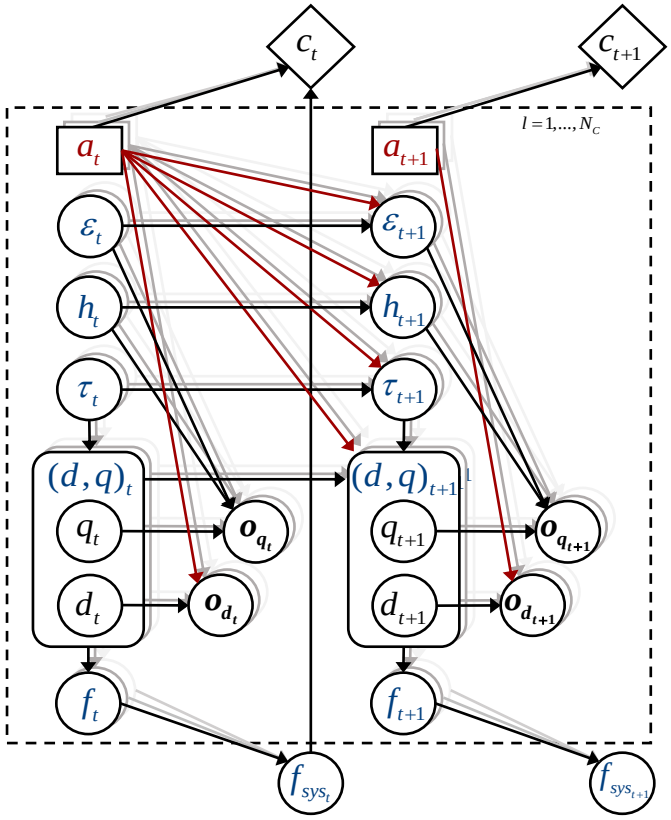

Fig. 6.1. Graphical representation of a POMDP for life-cycle management including virtual monitoring. The evolution of a component's deterioration mechanism, parametrized by the damage d_t, an underlying random variable q_t, the deterioration rate τ_t, is controlled by the actions a_t and informed by the imperfect observations $\boldsymbol{o_{d_t}}$ and $\boldsymbol{o_{q_t}}$. The observation o_{q_t} is also influenced by sensor status h_t and the sensor's measurement uncertainty ε_t. The system cost depends on the taken actions a_t and the system failure state f_{sys_t}, defined conditional on the components' failure states f_{c_t}.

On the other hand, the sensor status $\mathbf{b}_{h_t}$ also evolves, according to the transition model $P(h_{t+1}|h_t, a_t)$, as:

$$b(h_{t+1} \mid a_t) = \sum_{h_t \in \mathcal{S}_h} P(h_{t+1}|h_t, a_t)\, b(h_t). \tag{6.15}$$

If the sensor status $\mathbf{b}_{h_t}$ indicates the presence of a physical or virtual sensor, the agent will collect an observation of the underlying parameter o_q. If an inspection is also performed, the agent observes both o_q and o_d. Upon collection of the evidence (observation), the updated belief of each component's deterioration can be estimated through the normalization of:

$$P(d_{t+1}, q_{t+1} \mid o_{d_{t+1}}, o_{q_{t+1}}) \propto P(o_{d_{t+1}}, o_{q_{t+1}} \mid d_{t+1}, q_{t+1}) P(d_{t+1}, q_{t+1}). \tag{6.16}$$

The measurement uncertainty of physical sensors ε_p can be obtained prior to the installation, for instance, directly from the manufacturers or laboratory tests. However, the uncertainty of

virtual sensors is unknown prior to physical sensor installation, data collection, training and validation. In this context, the uncertainty of virtual sensors can be considered as a probability distribution $\varepsilon_v \sim \mathcal{N}(\mu_{\varepsilon_v}, \sigma_{\varepsilon_v})$, in which the distribution parameters should be defined according to statistical data, or previously reported results of the selected modeling approach. Alternatively, a virtual monitoring model can be trained during the computation of the pre-posterior analysis.

The uncertainty of virtual sensors is nevertheless fully observable, i.e., a probabilistic virtual sensor will also provide uncertainty information along with the predictions [176]. In this context, an uncertainty realization $\hat{\varepsilon}_v$ is drawn from the initial distribution ε_{v_t}. In the next time step, the virtual sensor's uncertainty is then updated by replacing the Gaussian mean by the current realization:

$$\varepsilon_{v_{t+1}} \sim \mathcal{N}(\hat{\varepsilon}_v, \sigma_{\varepsilon_v}). \tag{6.17}$$

In this formulation, the uncertainty of virtual sensors can change over time, being rational in real-world scenarios where virtual sensors are likely to receive both in-training and out-of-training operational conditions as inputs. In certain cases, even the dynamic behaviour of the structure may sometimes change, for example, offshore wind turbines often experience scouring which can cause alternations in their structural dynamics [93]. When significant variations are expected in the lifetime, a higher value of σ_{ε_v} can be considered.

The failure probability of each component can then be estimated by specifying a node which represents the component's failure state f_t conditional on the deterioration state of the analyzed component $p(f_{t+1}|d_{t+1}, q_{t+1})$. The component failure probability can be computed as:

$$b(f_{t+1}) = \sum_{d_{t+1} \in \mathcal{S}_d} \sum_{q_{t+1} \in \mathcal{S}_q} P(f_{t+1}|d_{t+1}, q_{t+1})\, b(d_{t+1}, q_{t+1}). \tag{6.18}$$

The probability of a system failure event $P_{f_{sys}}$ is subsequently defined, conditional on the components' failure probabilities, specified as:

$$b(f_{sys_{t+1}}) = \sum_{\mathbf{f_{t+1}} \in S_F} P(f_{sys_{t+1}}|\mathbf{f_{t+1}})\mathbf{b}(\mathbf{f_{t+1}}), \tag{6.19}$$

where $\mathbf{f_{t+1}}$ is the joint probability distribution of failure/survival states among all components. In some cases, one could assume full observability of the system failure state. In that case, the component failure states and component damage states can also be inferred from the observed system failure state $\hat{f}_{sys_{t+1}}$. However, we assume in this paper that the system failure is not fully observable.

All algorithmic steps regarding the transitioning, observation probabilities and belief updates in the virtual sensing integrated environment are summarized in Algorithm 1.

Algorithm 1 POMDP dynamics in virtual sensing integrated environment

for $t = 1, t_N$ **do**
  for $l = 1, N_c$ **do**
    Select action $a_t^{(l)}$.
    $b\left(\tau_{t+1}^{(l)}\right) \leftarrow P\left(\tau_{t+1}^{(l)} | \tau_t^{(l)}, a_t^{(l)}\right)\, b\left(\tau_t^{(l)}\right)$ ▷ propagation step
    $\tilde{b}\left(d_{t+1}^{(l)}, q_{t+1}^{(l)}\right) \leftarrow P\left(d_{t+1}^{(l)}, q_{t+1}^{(l)} | d_t^{(l)}, q_t^{(l)}, \tau_{t+1}^{(l)}, a_t^{(l)}\right)\, b\left(d_t^{(l)}, q_t^{(l)}\right)\, b\left(\tau_{t+1}^{(l)}\right)$
    $b\left(h_{t+1}^{(l)}\right) \leftarrow P\left(h_{t+1}^{(l)} | h_t^{(l)}, a_t^{(l)}\right)\, b\left(h_t^{(l)}\right)$
    $b\left(f_{t+1}^{(l)}\right) \leftarrow p\left(f_{t+1}^{(l)} | d_{t+1}^{(l)}, q_{t+1}^{(l)}\right)\, \tilde{b}\left(d_{t+1}^{(l)}, q_{t+1}^{(l)}\right)$ ▷ component failure probability
    $\hat{h}_t^{(l)} \sim b(h_t^{(l)})$ ▷ sensor status state
    if $\hat{h}_t^{(l)}$ = 'physical sensor' **then** ▷ sensor's measurement uncertainty
      $\varepsilon^{(l)} \leftarrow \varepsilon_p$
    else if $\hat{h}_t^{(l)}$ = 'virtual sensor' **then**
      $\hat{\varepsilon}_v^{(l)} \sim \varepsilon_{v,t}^{(l)}$
      $\varepsilon^{(l)} \leftarrow \varepsilon_p + \hat{\varepsilon}_v^{(l)}$
      $\varepsilon_{v,t+1}^{(l)} \leftarrow \mathcal{N}(\hat{\varepsilon}_v^{(l)}, \sigma_{\varepsilon_v})$
    else if $\hat{h}_t^{(l)}$ = 'no sensor' **then**
      skip
    end if
    $P(\mathbf{o}_{d_{t+1}}^{(l)} \mid \tilde{\mathbf{b}}_{d_{t+1}}^{(l)}) \leftarrow P(o_d \mid d)\, \tilde{\mathbf{b}}_{d_{t+1}}^{(l)}$ ▷ observation probabilities
    if $\hat{h}_t^{(l)}$ = 'physical sensor' **or** $\hat{h}_t^{(l)}$ = 'virtual sensor' **then**
      $P(\mathbf{o}_{q_{t+1}}^{(l)} \mid \tilde{\mathbf{b}}_{q_{t+1}}^{(l)}, \varepsilon^{(l)}) \leftarrow P(o_q \mid q, \varepsilon^{(l)})\, \tilde{\mathbf{b}}_{q_{t+1}}^{(l)}$
    else
      $P(o_{q_{t+1,(1)}}^{(l)} \mid \tilde{\mathbf{b}}_{q_{t+1}}^{(l)}) \leftarrow 1$
    end if
    $b\left(d_{t+1}^{(l)}, q_{t+1}^{(l)}\right) \leftarrow p\left(o_{d_{t+1}}^{(l)}, o_{q_{t+1}}^{(l)} | d_{t+1}^{(l)}, q_{t+1}^{(l)}, h_t^{(l)}, a_t^{(l)}\right)\, \tilde{b}\left(d_{t+1}^{(l)}, q_{t+1}^{(l)}\right)\, b\left(h_t^{(l)}\right)$ ▷ inference
    $b\left(d_{t+1}^{(l)}, q_{t+1}^{(l)}\right) \leftarrow b\left(d_{t+1}^{(l)}, q_{t+1}^{(l)}\right) / \left\{p\left(o_{d_{t+1}}^{(l)}, o_{q_{t+1}}^{(l)} | \mathbf{b}_{d_{t+1}, q_{t+1}}^{(l)}, \mathbf{b}_{h_{t+1}}^{(l)}, a_t^{(l)}\right)\right\}$
  end for
  $b\left(f_{sys_{t+1}}\right) \leftarrow p\left(f_{sys_{t+1}} | \mathbf{f_{t+1}^{(l)}}\right)\, \mathbf{b}\left(\mathbf{f_{t+1}^{(l)}}\right)$ ▷ system failure probability
end for

6.4.3 Policy optimization via multi-agent reinforcement learning

Life-cycle management problems, formulated as (decentralized) POMDPs can be solved either by dynamic programming, e.g., exact value iterations in small state spaces or point-based solvers from medium to high dimensional state spaces [22, 24]. On the other hand, for the life-cycle planning for very large infrastructures, reinforcement learning algorithms offer an attractive solution, parameterizing value function or policies with artificial neural networks [38, 43]. A benchmark on state-of-the-art multi-agent reinforcement (MARL) algorithms can be found in [49], with the specific application in infrastructure management planning. In this work, we adopt QMIX algorithm [54] to approximate the solution of the formulated Dec-POMDP problem. A brief description of the algorithm is provided below, and the interested reader is directed to [54] for a more elaborated explanation.

For each component l, there is one agent network that represents its individual action-value function $Q^{(l)}(\mathbf{b}, a)$, denoted $Q_{(l)}$ for short. Each agent network takes the component's current belief states $\mathbf{b}_t^{(l)}$ and the action from the previous time step $a_{t-1}^{(l)}$, and outputs the estimation of individual Q function:

$$Q_{(l)} = Q^{\pi,(l)}(\mathbf{b}, a) = \mathbb{E}_{\mathbf{b}_{t+1:t_N}, a_{t+1:t_N}^{(l)}} \left[R_t \mid \mathbf{b}_t, a_t^{(l)}, \pi \right], \tag{6.20}$$

where R_t is the sum of discounted rewards as follows:

$$R_t = \sum_{i=0}^{t_N - t} \gamma^i \; r_{t+i}. \tag{6.21}$$

The outputs of the agent networks are then fed to the mixing network which maps to the joint Q function, denoted Q_{tot}:

$$Q_{tot} = Q^{\pi}(\mathbf{b}, \mathbf{a}) = \mathbb{E}_{\mathbf{b}_{t+1:t_N}, \mathbf{a}_{t+1:t_N}} \left[R_t \mid \mathbf{b}_t, \mathbf{a}_t, \pi \right]. \tag{6.22}$$

The key property of the mixing network is the monotonicity constraint:

$$\frac{\partial Q_{tot}}{\partial Q_{(l)}} > 0, \tag{6.23}$$

so that as the individual $Q_{(l)}$ values increase, the joint value Q_{tot} also increases, reflecting a cooperative relationship among agents. To ensure the monotonicity constraint, the weights of the mixing network are restricted to be non-negative. The mixing network also receives the global information in the form of concatenated components' beliefs. However, Q_{tot} is monotonic only with respect to each $Q_{(l)}$ value, but not necessarily to the global belief. Therefore, a separate hypernetwork is used to estimate the non-negative weights of the mixing network using the global belief as inputs. The mixing network helps the agents learn coordinated policies in a centralized training, but once the policies are learned, the mixing network is not used during decentralized execution, and instead, each agent independently selects its action from the agent networks.

During the training, QMIX uses a loss function that is similar to the temporal difference loss used in Q-learning, i.e., the difference between the predicted and target Q_{tot} values:

$$\begin{aligned} \mathcal{L}(\theta) &= \sum_u \left(Q_{tot}^{pred}(\mathbf{b}, \mathbf{a}) - Q_{tot}^{target}(\mathbf{b}, \mathbf{a}) \right)^2 \\ &= \sum_u \left(Q_{tot}(\mathbf{b}, \mathbf{a}; \theta) - \left(r(\mathbf{b}, \mathbf{a}) + \gamma \max_{\mathbf{a}'} Q_{tot}(\mathbf{b}', \mathbf{a}'; \theta^-) \right) \right)^2 \end{aligned} \tag{6.24}$$

where r is the immediate reward and u is the batch size. Random batches of experiences are sampled from the replay buffer, where the agent stores past experiences (tuples of belief state, action, reward, next belief state). The main networks' parameters are described by θ and

those of the target networks are described by θ^-. Target networks are used to stabilize the learning process by providing more consistent target values, and are periodically updated with the parameters of the main networks.

6.4.4 Value of virtual health monitoring

We aim to quantify the value of virtual health monitoring in the integrated life-cycle management framework. In this context, the life-cycle management planning is performed with and without including virtual monitoring. When there is no virtual monitoring option, the sensor status h directly transitions from 'physical sensor' to 'no sensor' state as follows:

$$P(h_{t+1} \mid h_t, a_{DN}) = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix}; \qquad P(h_{t+1} \mid h_t, a_I) = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix};$$
$$P(h_{t+1} \mid h_t, a_S) = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}; \qquad P(h_{t+1} \mid h_t, a_{PR}) = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix}. \tag{6.25}$$

The VoI is calculated as the difference between the expected cost in which virtual monitoring is not included $\mathbb{E}(C_0)$ and the expected cost in which monitoring is included $\mathbb{E}(C_1)$:

$$VoI = \mathbb{E}(C_0) - \mathbb{E}(C_1). \tag{6.26}$$

Since the virtual monitoring models are cost-free, i.e., the computational cost and software/hardware costs are neglected, the VoI is also equal to the NVoI. Theoretically, NVoI of unbiased monitoring information is expected to be non-negative if the policies are optimal. However, it can be zero if the information does not improve the I&M strategies which means the virtual monitoring is not useful in the decision-making context.

6.5 Numerical experiments

The presented life-cycle management framework with virtual monitoring is implemented and tested for life-cycle management of an offshore wind support structure under fatigue deterioration. The objective is to plan inspections, maintenance and structural health monitoring of an OWT support structure over a lifetime of 20 years in order to minimize the expected life-cycle cost.

6.5.1 Fatigue deterioration modeling

The design of offshore wind foundations is often based on fatigue damage modeling along with SN curves and Miner's cumulative damage rule. A component is assigned as failed when the

cumulative fatigue damage exceeds the failure limit Δ corrected by the fatigue design factor (FDF), formulating the probabilistic fatigue limit state g_{SN} as follows:

$$g_{SN}(t) = \Delta - D_{SN}(t), \tag{6.27}$$

whereas temporal fatigue accumulation $D_{SN}(t)$ is governed by the load distribution and SN curves. The long-term stress range in offshore wind substructures is often described by a Weibull distribution with scale parameter q and shape parameter λ, and the fatigue resistance is empirically parameterized by bi-linear SN curves with slopes m_1 and m_2 changing at the reference point of S_1, along with the corresponding SN curve intercepts $C_{1,SN}$ and $C_{2,SN}$. The temporal accumulated fatigue damage is then computed as follows:

$$D_{SN}(t) = \nu t\left[\frac{q^{m_1}}{C_{1,SN}}\gamma_1\left\{1+\frac{m_1}{\lambda};\left(\frac{S_1}{q}\right)^{\lambda}\right\} + \frac{q^{m_2}}{C_{2,SN}}\gamma_2\left\{1+\frac{m_2}{\lambda};\left(\frac{S_1}{q}\right)^{\lambda}\right\}\right], \tag{6.28}$$

where ν is the fatigue cycle rate and γ_1 and γ_2 correspond to incomplete gamma functions. Since load predictions in offshore wind turbines involve significant uncertainties, the scale parameter q of long-term stress range distribution is also modeled as a random variable. When installing load monitoring systems, the collected data can be used to reduce the statistical uncertainty of q, subsequently updating the fatigue damage estimation $D_{SN}(t)$.

Fatigue damage, although its estimation uncertainty can be reduced indirectly through load monitoring, is not physically observable. Non-destructive inspections can only detect the presence, location and/or severity of fatigue cracks. A prior modeling of the fatigue deterioration cannot be combined with inspections within a Bayesian approach. In this regard, we calibrate a fracture mechanics (FM) model to the probabilistic SN-Miner's cumulative damage model, such that the FM model with the following limit state renders similar structural reliability as the SN model:

$$g_{FM}(t) = d_c - \left[d_{t-1}^{\frac{2-m}{2}} + \frac{2-m}{2}C_{FM}(Y\pi^{0.5}S_e)^m n\right]^{\frac{2}{2-m}} \tag{6.29}$$

where d_c is the critical crack size, often assumed to be the thickness of a structural component. The crack growth d_t is computed, according to Paris law originally proposed by Ditlevsen [119], as a function of the crack size at the previous time step d_{t-1}, Paris' law parameters C_{FM} and m, equivalent stress range S_e, stress intensity factor Y and number of cycles in one time step n. The uncertainty of the crack size can be updated through Bayesian inference using the crack observations collected at inspections or the load observations from monitoring. It is worth noting that if a fracture mechanics model has been used during the design phase, it can then be directly used for inspection and maintenance planning.

The equivalent stress range S_e is often considered time-invariant and is computed as the expected stress range of the two-parameter Weibull distribution:

$$S_e = q\Gamma(1+1/\lambda), \tag{6.30}$$

where q is the scale parameter and λ is the shape parameter, usually taken as 0.8 for offshore wind support structures. In this paper, we consider the equivalent stress range to be time-variant by introducing Gaussian noise to the temporal evolution of the scale parameter q. This incorporation of stochastic variation encapsulates the potential dynamic changes in offshore wind turbine behavior, accounting for dynamic fluctuations induced by scouring, rotor imbalance, among others. At each time step, the scale parameter experiences the perturbation of Gaussian noise, represented by $\epsilon_q \sim \mathcal{N}(1, \sigma_{\epsilon_q})$, and the fracture mechanics limit state in Equation (6.29) is reformulated as:

$$g_{FM}(t) = d_c - \left[d_{t-1}^{\frac{2-m}{2}} + \frac{2-m}{2} C_{FM} \{ Y \pi^{0.5} q \epsilon_q \Gamma(1 + 1/\lambda) \}^m n \right]^{\frac{2}{2-m}} \tag{6.31}$$

6.5.2 System modeling

The analyzed offshore sub-structure consists of three different types fatigue hotspots with the following characteristics:

- Hotspot above the water level (AW): This hotspot is assumed to be located above the waterline, and with a plate thickness of 20 mm. Due to the ease of accessibility, a fatigue design factor (FDF) of 1 is considered. The SN curves in air environmental conditions are used to compute the fatigue damage.
- Hotspot below the water level (BW): This structural detail is located in the splash zone, below the waterline, and with a plate thickness of 60 mm. Due to a slightly difficult accessibility, a fatigue design factor (FDF) of 2 is considered for this hotspot. In the splash zone, the corrosive environment accelerates fatigue deterioration, thus the SN reference curves in a corrosive environment with cathodic protection are selected for the fatigue damage computation.
- Hotspot below the mudline (MD): This structural detail is located below the mudline, and with a plate thickness of 60 mm. Being in an inaccessible zone, the structural detail cannot be maintained as well as inspections and monitoring are not possible. Therefore, a fatigue design factor (FDF) of 3 is assigned. The fatigue damage is estimated using SN curves for a corrosive environment with cathodic protection.

The deterioration parameters used for each type of component, both for Miner's rule fatigue model and the calibrated fracture mechanics model are listed in Table 6.1. All the continuous random variables are discretized, allowing the efficient implementation of deterioration transition and inference through dynamic Bayesian networks. The discretization scheme is presented in Table 6.2, listing the intervals and state space of the crack size d, deterioration rate τ, scale factor q, and the sensor status h.

The considered offshore wind turbine substructure consists of five levels- two above water, two below water and one below mudline. There are two hotspots at each level which are

Table 6.1. Random variables and deterministic parameters for modeling the fatigue deterioration.

Parameter	Distribution	Mean	Std
h	Deterministic	0.8	-
v (cycles/s)	Deterministic	0.16	-
m_1	Deterministic	3	-
m_2	Deterministic	5	-
Δ	Lognormal	1	0.3
d_0 (mm)	Exponential	0.11	-
Y	Lognormal	1	0.1
m_{FM} (mm)	Deterministic	3	-
AW hotspot			
$C_{1,SN}$*	Normal	12.164	0.2
$C_{2,SN}$*	Normal	15.606	0.2
q (MPa)	Trunc. Normal	10.21	2.55
ϵ_q	Normal	1	0.05
lnC_{FM}	Normal	-26.445	0.122
d_c (mm)	Deterministic	20	-
BW hotspot			
$C_{1,SN}$*	Normal	11.764	0.2
$C_{2,SN}$*	Normal	15.606	0.2
q (MPa)	Trunc. Normal	7.40	1.85
ϵ_q	Normal	1	0.05
lnC_{FM}	Normal	-26.043	0.403
d_c (mm)	Deterministic	60	-
MD hotspot			
$C_{1,SN}$*	Normal	11.764	0.2
$C_{2,SN}$*	Normal	15.606	0.2
q (MPa)	Trunc. Normal	6.74	1.68
ϵ_q	Normal	1	0.05
lnC_{FM}	Normal	-26.122	0.396
d_c (mm)	Deterministic	60	-

*Fully correlated.

considered to be in parallel connection and the five levels form a series system. In terms of implementation, we rely on efficient matricial algorithms proposed by Song et al. [183] to compute the system probability of failure $P_{f_{sys}}$ for k-out-of-N systems, i.e., k out of N components have to be working so that the system is functional. In this context, the whole support structure represents a 5-out-of-5 system where each level is a 1-out-of-2 system. Figure 6.2 illustrates the unconditional failure probabilities of each component and the system over the lifetime. The failure probability of the whole system is lower than the failure probability of some components due to the redundancies in the system modeling.

Table 6.2. Description of the discretization scheme for the crack size d, stress range scale parameter q, deterioration rate τ, sensor condition h, component failure state f, and system failure state f_{sys}.

Variable	Interval boundaries	States
d	$[0, d_0 : (d_c - d_0)/(\lvert S_d \rvert - 2) : d_c, \infty]$	30
q	$[0 : (\max(q) + 0.1)/(\lvert S_q \rvert - 1) : (\max(q) + 0.1), \infty]$	20
τ	$[0 : 1 : 21]$	21
h	[physical sensor : virtual sensor : no sensor]	3
f	[survival, failure]	2
f_{sys}	[survival, failure]	2

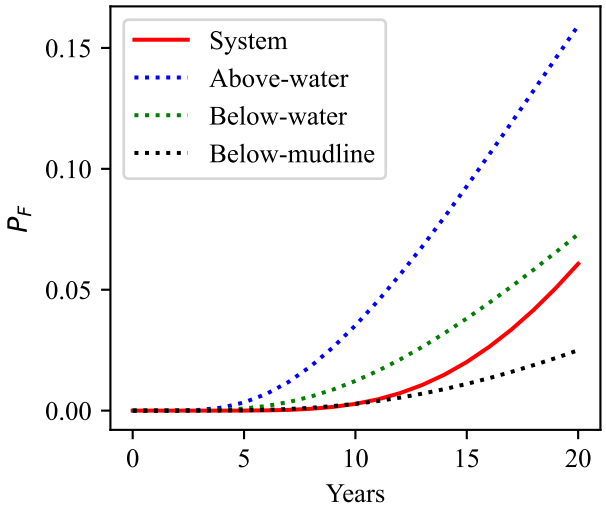

Fig. 6.2. Unconditional failure probability of each component and the support structure system.

6.5.3 Inspection and monitoring models

The measurement uncertainty associated to inspections is described by probability of detection curves, which can be found in offshore wind design standards [117]. Out of the three different types of hotspots, only the structural details above the waterline and those at the splash zone can undergo inspection. The eddy current (EC) inspection, a common inspection method for offshore wind support structures, is used as a reference inspection model and the probability of detection curve is defined as follows:

$$p(o_{d_t} \mid d_t) = 1 - \frac{1}{1 + (d_t/\chi)^b}, \tag{6.32}$$

where the parameters χ and b are 1.16 and 0.90 respectively for the EC inspection at the normal working conditions.

In addition to crack inspections, strain can also be monitored through operational sensors, providing statistical information on stress range distribution, denoted as q_t. If the sensors could measure q_t perfectly, then the state of q_t would be fully observable. However, due to

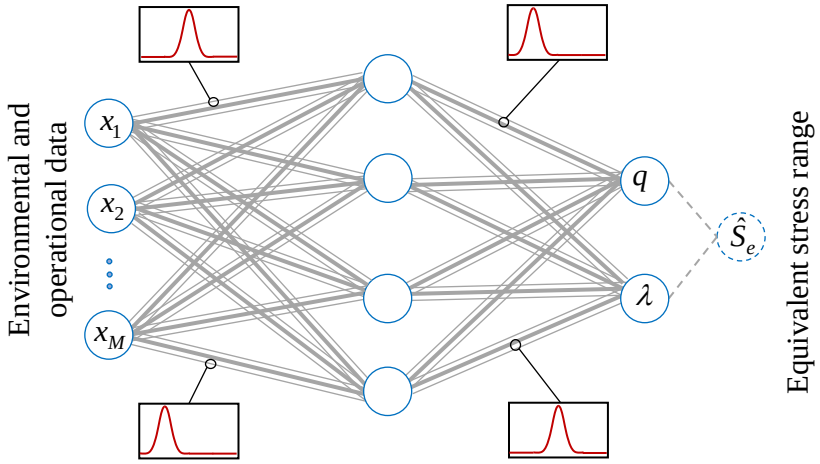

Fig. 6.3. Virtual sensing model using a probabilistic Bayesian neural network.

existing measurement errors and the fact that sensors do not directly measure strain at the precise fatigue hotspot location, we account for these uncertainties by introducing an unbiased Gaussian noise, characterized by a 7% coefficient of variation. Since the monitoring outcome is continuous, we use the same discretization scheme as q in Table 6.2. In total, there are a total of 40 observations available, consisting of 20 load observations for o_{q_t} and 2 crack observations for o_{d_t}.

Virtual sensing through Bayesian neural networks

In the numerical experiments, we adopt Bayesian neural networks as a conceptual model of virtual sensing. The BNN-based virtual sensing framework has been previously proposed for offshore wind farms [176], enabling the intrinsic modeling of aleatory uncertainty (emerging due to the random nature of the physical system) and epistemic uncertainty (arising due to the lack of knowledge of the system). Bayesian neural networks are demonstrated to be able to inform higher model uncertainty when receiving input variables that divert from the training data, especially being suitable as virtual sensors for offshore wind turbines being subject to structural dynamic variations over time.

The considered conceptual model consists of a Bayesian neural network which receives environmental and operational data in the form of short-term statistics and predicts the equivalent stress range for that period, by learning the underlying scale parameter q of the stress range distribution, shown in Figure 6.3. The environmental and operational data can be collected mainly from supervisory, control and data acquisition (SCADA) system and motion sensors such as accelerometers, which are generally present along the wind turbine's lifetime. The output layer captures the equivalent stress range parameterized by the Weibull scale parameter q and shape parameter λ. Whereas the shape parameter λ is deterministic with a fixed value of 0.8, the network learns the scale parameter q. The model uncertainty is represented by

the network's probabilistic weights, which are estimated during the training through Bayesian inference [162, 184].

Given that the weights and biases within the network are represented as probability distributions, the BNN produces stochastic realizations of the equivalent stress range $\hat{S}_e$ and the underlying scale parameter $\hat{q}$ during each forward execution. A sufficient number of forward simulations is required to estimate the expected prediction and associated uncertainty of q. In the posterior analysis, the short term predictions can be accumulated up to one year period (decision interval), obtaining the yearly observation of shape parameter o_{q_t} and $\hat{\varepsilon}_v$. In the pre-posterior analysis, $\hat{\varepsilon}_v$ is randomly withdrawn from the current probability distribution of virtual sensor's uncertainty $\varepsilon_{v,t}$, and the probability of $o_{q_t} \in \mathcal{O}_q$ depends on the current belief state of q and the uncertainty of the virtual sensor. The initial ε_v of virtual sensor just after (re-)training is assumed to follow a Gaussian distribution $\mathcal{N}(\mu_\varepsilon = 0.1, cov_\varepsilon = 0.25)$, referenced from the previously reported results [176].

6.5.4 Virtual monitoring integrated life-cycle management of offshore wind substructures

The decision-maker can choose from six possible actions for each accessible hotspot: (i) Do nothing (DN), (ii) Inspection (I), (iii) Install sensor (S), (iv) Inspection and install sensor (I-S), (v) Perfect repair (PR), (vi) Perfect repair and install sensor (PR-S). Actions (i), (ii), and (v) are frequently used in conventional I&M decision-making problem. On the other hand, actions (iii), (iv), and (vi) are related to monitoring decisions. Installing a sensor changes the sensor state from 'no sensor' to 'physical sensor' and the monitoring observation will be only collected in the next year. As already mentioned in the previous section, the lifetime of a physical sensor is assumed to be only one year after which the sensor state will transition to 'virtual sensor'. If a maintenance action (PR) is performed while 'virtual sensor' is active, the sensor state will go back to 'no sensor', assuming that the virtual monitoring model does not represent the structural component anymore. The load observations are collected if the agent is in the 'physical sensor' or 'virtual sensor' state.

Each action incurs an associated cost and the overall cost encompasses failure costs, inspection costs, repair costs, and sensor installation costs. The campaign cost is also included if at least one hotspot is inspected, monitored or maintained and C_{ins}, C_{ins} and C_{rep} are added per action per component respectively. The cost models for various actions and components are described in Table 6.3. Underwater inspection and maintenance actions are more costly than above water due to the added complications such as the need of divers, etc.

6.5.5 Results

As explained before, the life-cycle planning is performed with and without deploying virtual sensing models. For each experiment, the multi-agent reinforcement learning is trained using

Table 6.3. Cost model for the life-cycle management planning.

γ	0.95		
c_{fsys}	10000		
c_{camp}	3		
	AW hotspot	BW hotspot	MD hotspot
c_{ins}	1	3	-
c_{sens}	3	9	-
c_{rep}	50	100	-

five different training seeds with up to 2,050,000 episodes in each realization. The model weights are saved at every 20,000 episodes. The replay buffer contains the most recent 2000 episodes. Batches of 64 episodes are sampled from the replay buffer during training. All neural networks are trained using RMSprop with learning rate $5 \cdot 10^{-4}$ and the target networks are updated at every 200 training episodes. The architecture of the agent networks and mixing networks is the same as in the original work of QMIX [54]. We used the same hyperparameters in all experiments without fine-tuning.

The statistics (mean, min-max range) of training realizations are shown in Figure 6.4. In all the training realizations, the expected total cost converges after $1e^6$ episodes. The value of the standard deviation is still significantly high in both cases at the end of the training. Since we also consider load monitoring information, the observation space is higher than traditional I&M settings in which only two observation outcomes are modeled, leading to the high variance of life-cycle cost. The networks' parameters which give the highest expected rewards are chosen from the saved models by running 100 test realizations of the life-cycle management environment. Table 6.4 describes the expected costs of life-cycle planning, $\mathbb{E}(C_0)$ without virtual monitoring and $\mathbb{E}(C_1)$ with virtual monitoring. The value of virtual monitoring information is significant in this case being 9.4 monetary units and 7% of the original life-cycle cost.

In order to visualize MARL-devised strategies, policy realizations for each case are shown in Figure 6.5 and 6.6, along with the evolution of failure probability for each component and the whole support structure. In addition, action histograms for each component over 1000 life-cycle realizations are also presented in Figure 6.7. Since the AW hotspots have the highest probability of failure among all the components, and the cost of observation and maintenance actions are also significantly cheaper, the MARL-devised strategies prioritize, in both cases, inspecting and maintaining those components.

In the life-cycle management without virtual monitoring, first sensor installations occur at year 6 on all components above the water. It can be seen that the agents seek to minimize the campaign cost by grouping the observation actions at the same decision step. Two years later, one of the below water components is also instrumented with a load sensor. Taking advantage of the unavoidable campaign cost, inspections of some components above the water are also performed, by selecting those with higher load indication from the installed sensors in order to effectively control the system failure probability. Since the cost of perfect repair incurs high

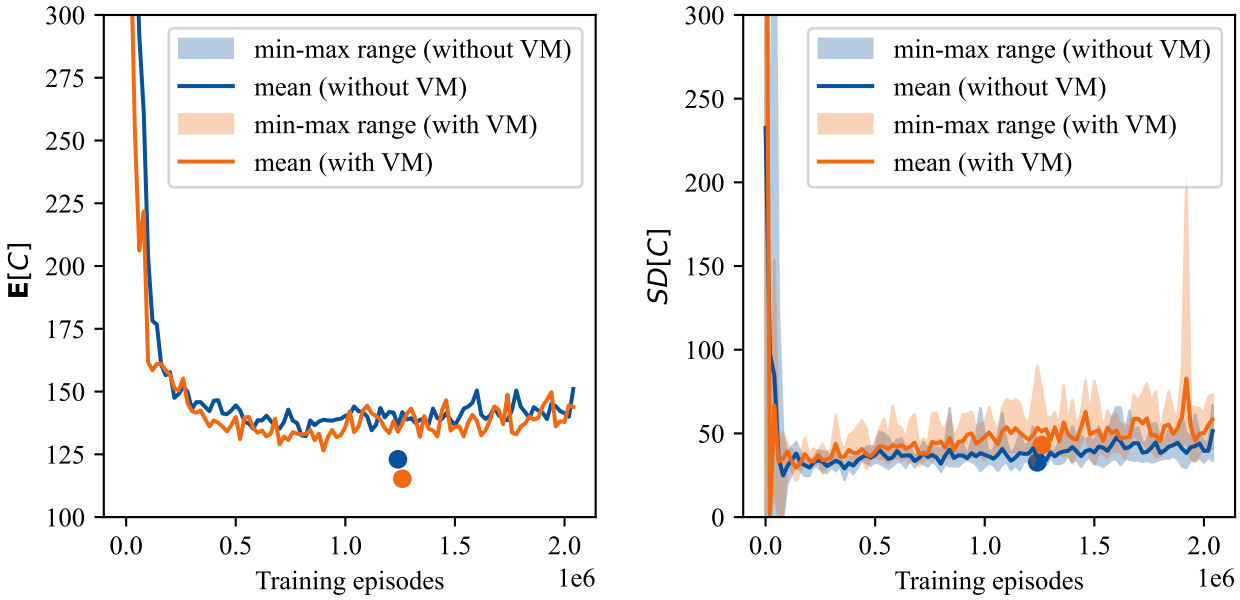

Fig. 6.4. The expected life-cycle cost (mean, min-max range) over training realizations. The solid lines and shaded colors represent the mean and min-max range of the expected life-cycle costs from different training seeds. The circular dots indicate the minimum expected life-cycle cost and its associated standard deviation for each case.

Table 6.4. Expected costs of life-cycle management planning and value of virtual health monitoring.

Life-cycle management case	$\mathbb{E}(C)$ (monetary units)	95% Confidence interval
Without virtual monitoring	133.6	2.35
With virtual monitoring	124.2	2.42
Value of virtual monitoring	9.4 (7%)	

costs, PR action is taken only in severe deterioration states, i.e., after consecutive detection indications or very high failure probability. Sensors are sometimes re-installed on the above water components before the end of the lifetime, as illustrated in Figure 6.5.

When virtual monitoring is included in the life-cycle management planning, we can see a different pattern of policies. The agents purely focus on the AW components, always assigning 'do nothing' action for the BW components. For all the AW components, sensors are installed earlier at year 4, thus maximizing the amount of information that can collected from virtual sensors. Inspections are then planned if the virtual monitoring yields high stress indications, and 'perfect repair' is executed following detection outcomes, as shown in Figure 6.6. It can be seen that the actions after year 4 significantly rely on the virtual sensing information. Sensor re-installation on the AW components is more frequent than in the previous case.

The percentage of each action taken over 1000 policy realizations is summarized in Figure 6.8, accounted for all accessible components. The number of inspections is reduced by using

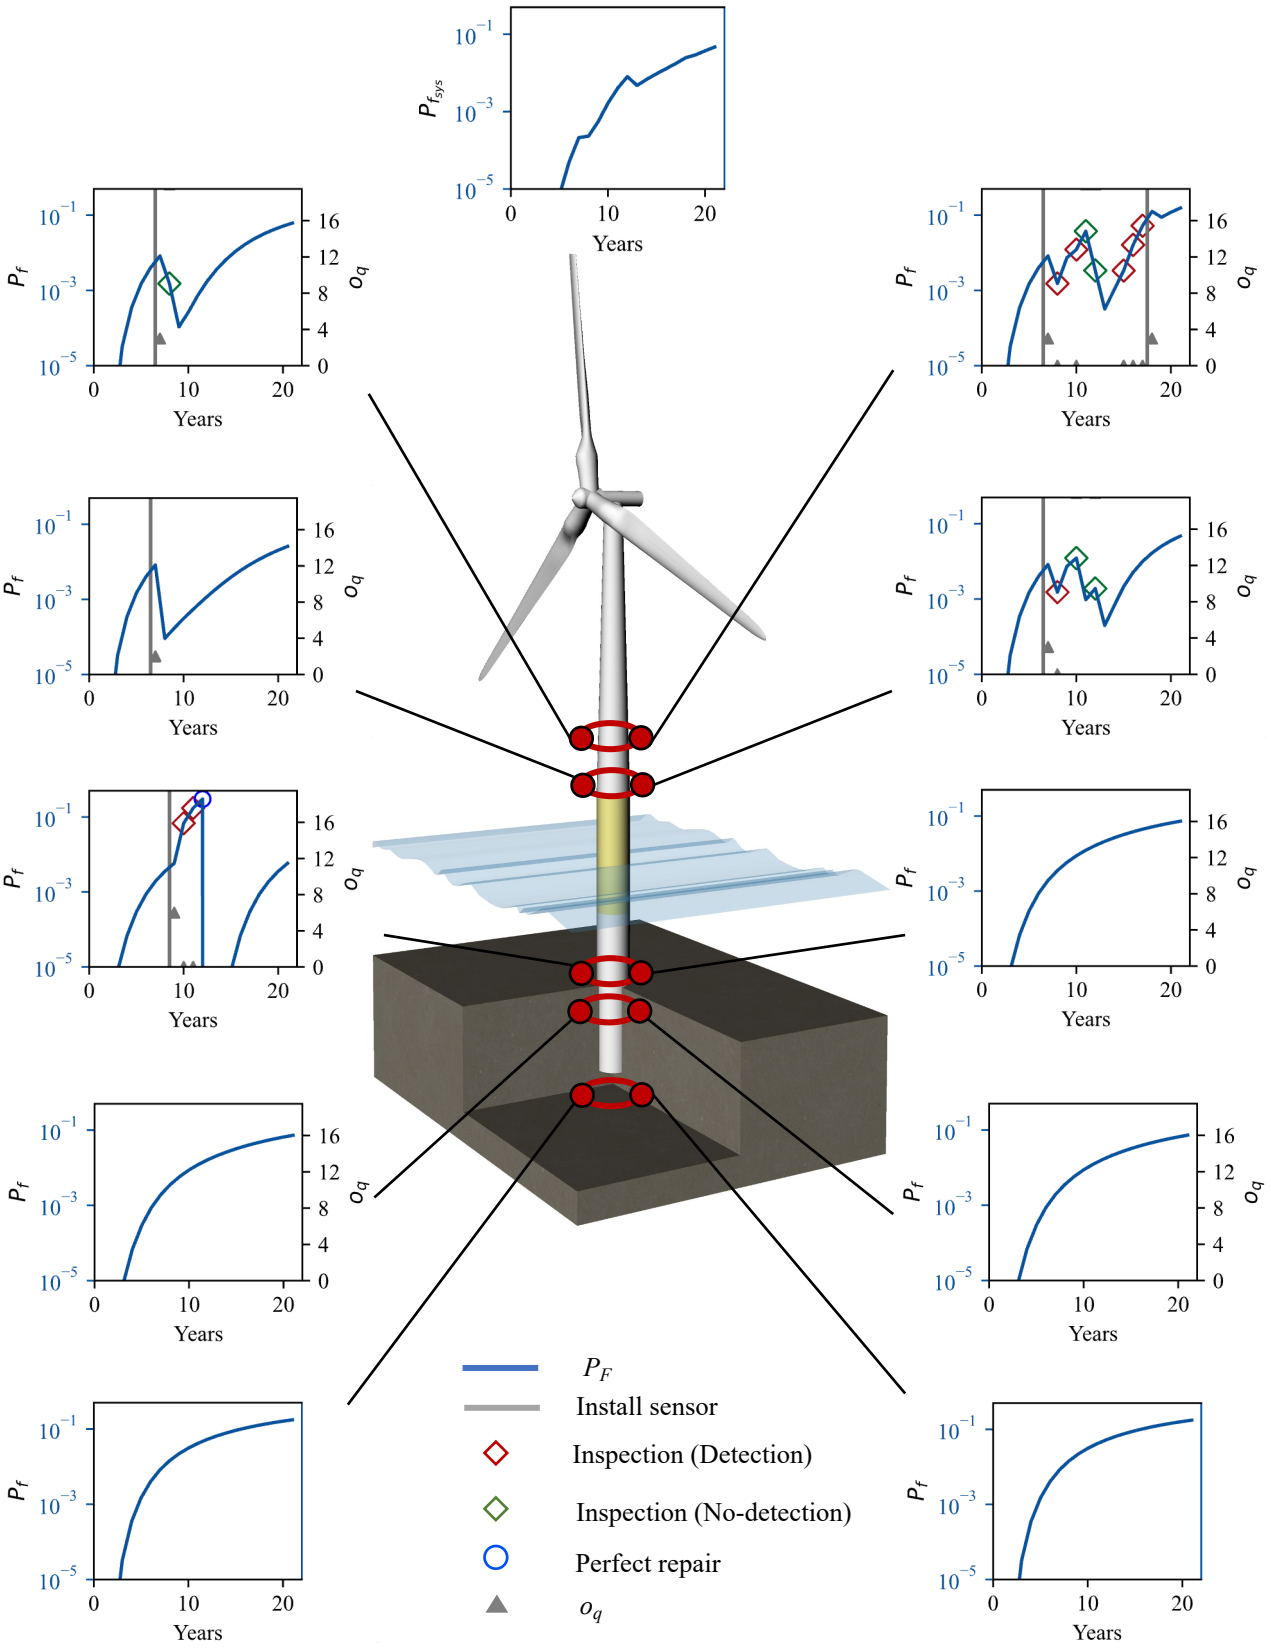

Fig. 6.5. Policy realizations of the components and system's failure probability in life-cycle management without virtual monitoring.

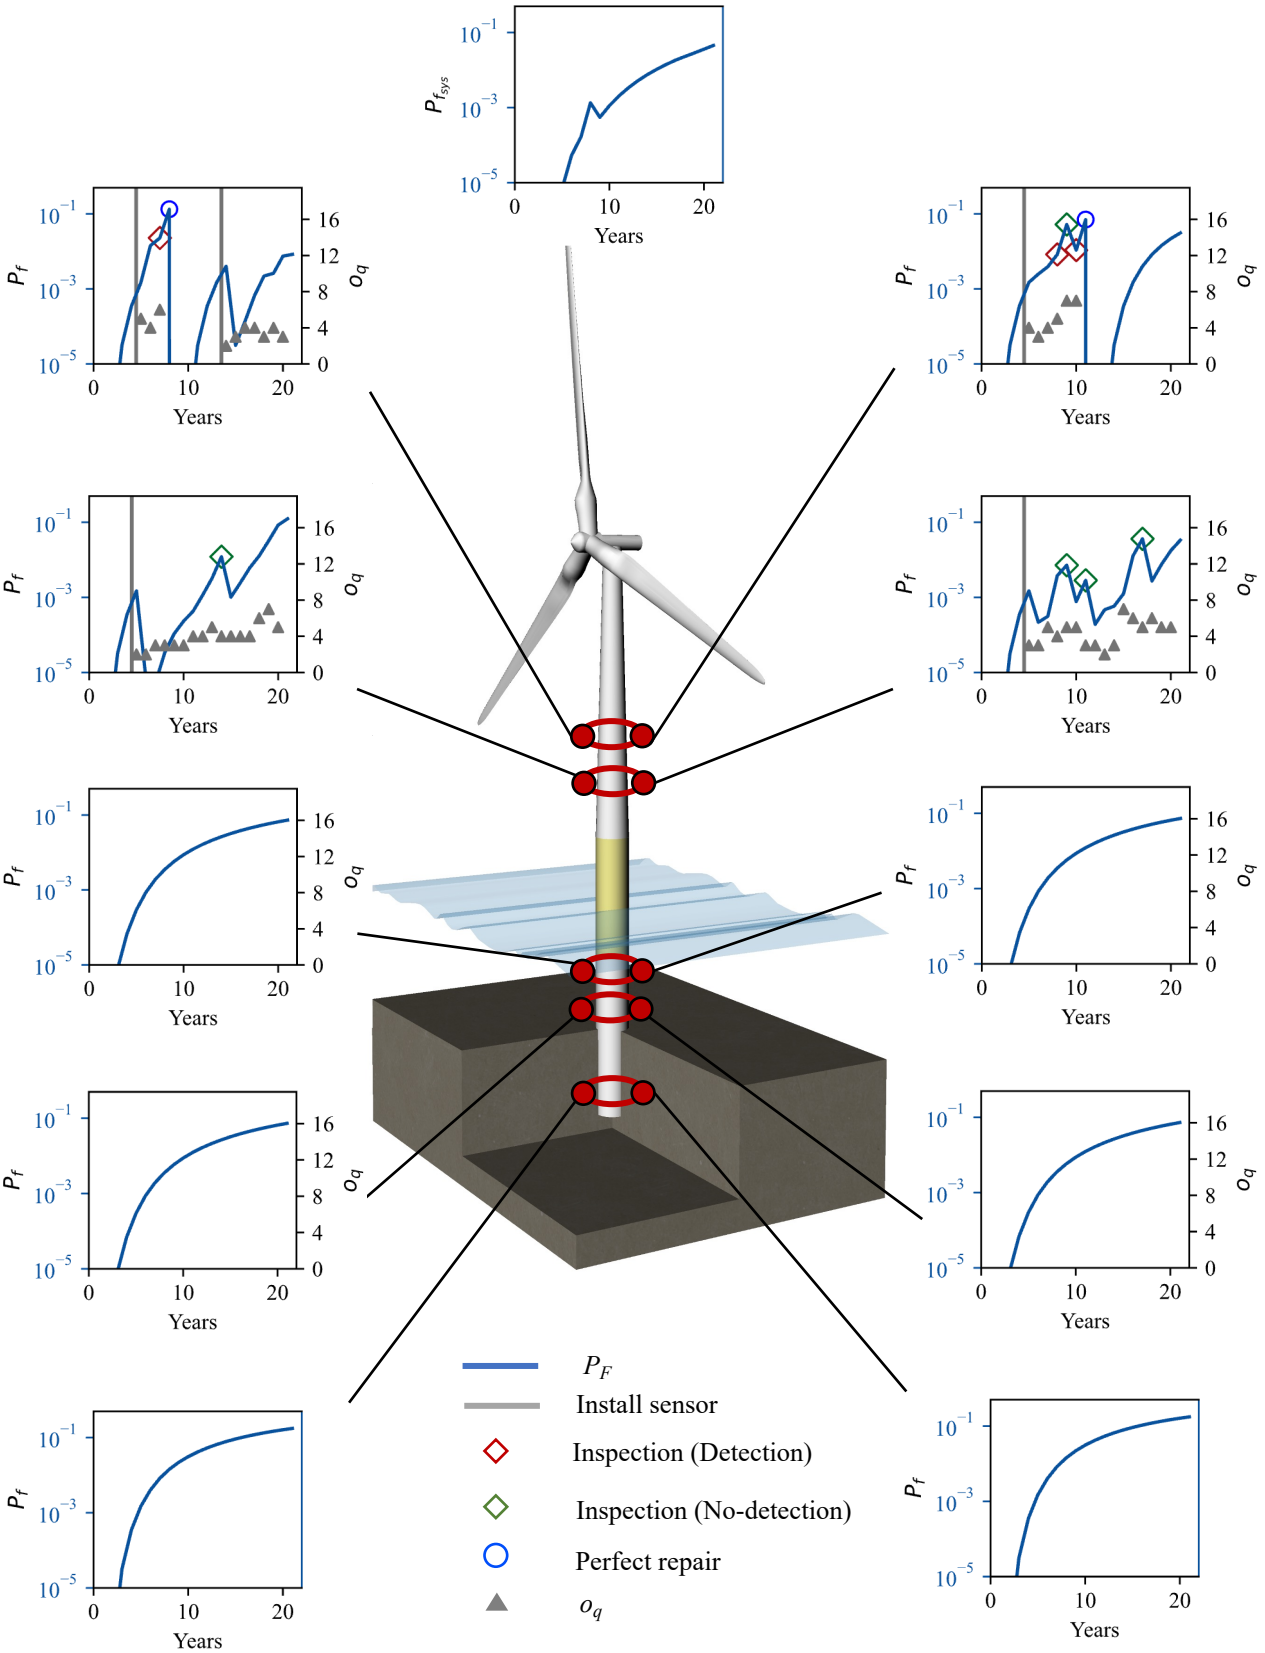

Fig. 6.6. Policy realizations of the components and system's failure probability in life-cycle management with virtual monitoring.

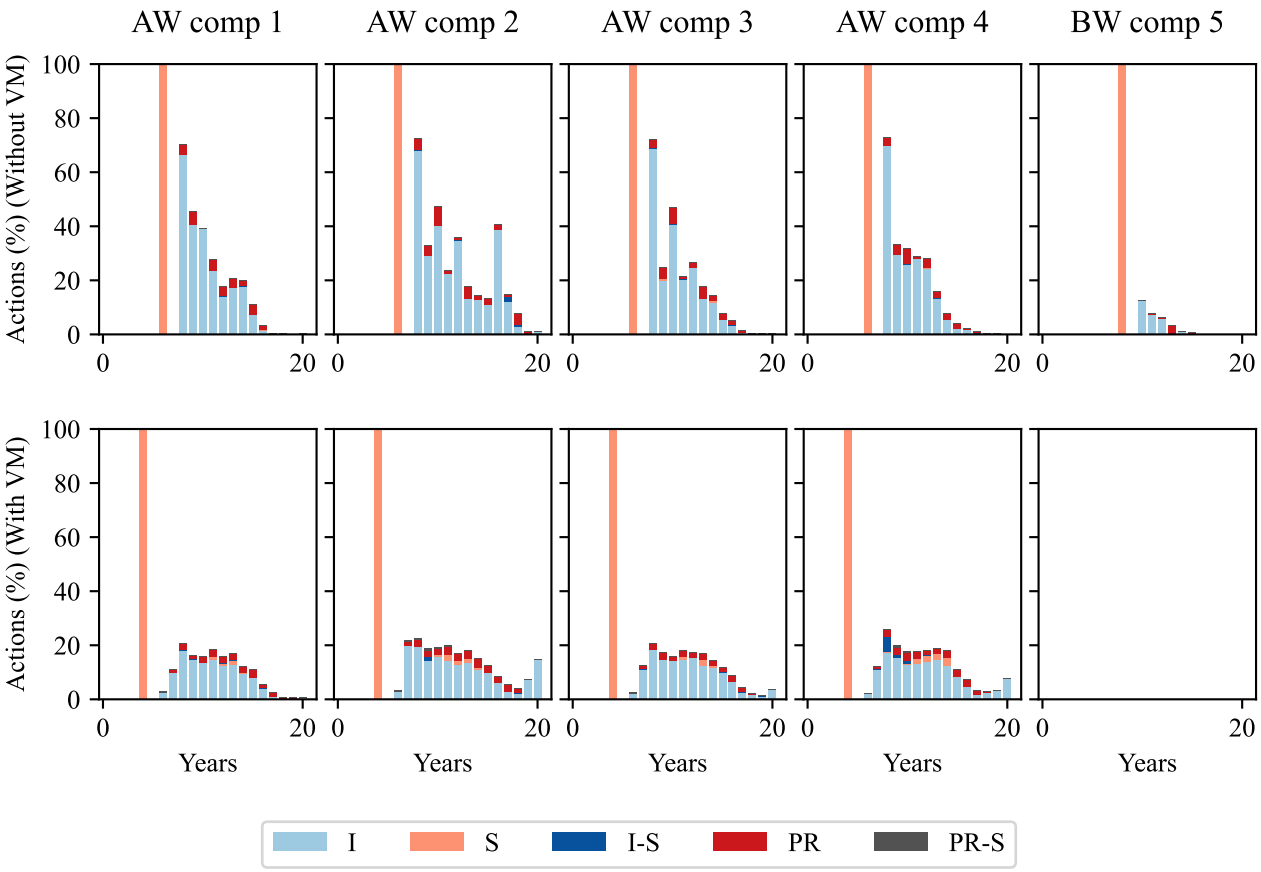

Fig. 6.7. Action histogram of each component in the life-cycle management. BW comps 6, 7, 8 are not included since the action is always 'do nothing' (DN).

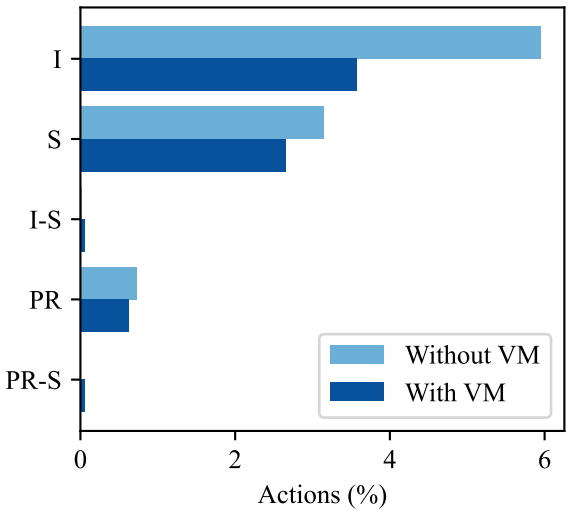

Fig. 6.8. Action histogram of all components in the life-cycle management.

virtual health monitoring, from around 6% to less than 4%, while also decreasing slightly the sensor installation and repair actions.

6.6 Concluding remarks

This paper introduces a life-cycle management framework for optimal monitoring, inspection, and maintenance planning of offshore wind structures subject to fatigue deterioration. Whereas previous studies assume pre-existence of physics-based/data-based digital twin for virtual sensing, the proposed framework relies on the pre-posterior analysis to optimally identify when to develop/re-calibrate the virtual monitoring model. The decision-making problem is formulated as a decentralized POMDP with high dimensional state, action, and observation spaces, and solved via a state-of-the-art multi agent reinforcement learning algorithm.

Monitoring observations are specified conditional on the sensor status, i.e., the information is only collected if there is a functional physical or virtual sensor. While a Bayesian neural network is showcased as a conceptual model for virtual monitoring to quantify predictive uncertainty, the presented framework is compatible with a broad range of uncertainty-aware probabilistic models. We also propose a stochastic uncertainty propagation of the virtual sensor, representing the time-varying nature of offshore wind turbine's structural dynamics. The proposed uncertainty propagation model is flexible and can be adapted for different initial uncertainties of virtual model and structural dynamics variations.

We applied the proposed framework in the life-cycle management planning of an offshore wind support structure consisting of multiple fatigue-critical components. It is demonstrated that the policies computed by the MARL algorithm can efficiently combine monitoring, inspection, and maintenance actions, providing optimal decision sequences that might be otherwise difficult to predict. The value of virtual monitoring is quantified by executing the prior and pre-posterior decision analyses. The results demonstrate that virtual monitoring information is effectively used to plan inspections and maintenance actions, resulting in the significant reduction of the expected life-cycle cost. Further research efforts are recommended toward the development and investigation of decision-making frameworks that include the management of the complete life-cycle from design to decommissioning. Life-cycle management of entire offshore wind farms informed by fleet-leader-based virtual monitoring models are also worth-exploring.

Authorship contribution statement

Hlaing, N.: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing - Original draft, Writing - Review & editing, Visualization.
Morato, P. G.: Conceptualization, Methodology, Software, Validation, Writing - Review & editing. **Rigo P.:** Supervision, Project administration, Funding acquisition.

Acknowledgements

This research is funded by the Belgian Energy Transition Fund (FPS Economy) through PhairywinD (www.phairywind.be) project. We further acknowledge the "Consortium des Équipements de Calcul Intensif", at the University of Liege (Belgium), for granting us access to high performance computing facilities.

Chapter 7

Conclusions and outlook

7.1 Concluding remarks

This thesis investigated the potential of leveraging data-driven techniques for virtual monitoring and effective management of offshore wind support structures. This work emphasizes the significant role of data-driven methodologies, particularly Bayesian neural networks (BNNs), partially observable Markov decision processes (POMDPs) and reinforcement learning, in optimizing maintenance strategies and managing uncertainties. This chapter provides a summary of the performed research along with the main findings and contributions drawn throughout the investigation.

As presented in Chapter 2, dynamic Bayesian networks (DBNs) can be adopted to model stochastic deterioration processes for offshore wind support structures. Even sophisticated failure criteria can be integrated with dynamic Bayesian networks, allowing efficient computation of failure probability and inference tasks. Life-cycle management problems can be formulated as partially observable Markov Decision Processes (POMDPs), with transition and observation models derived from DBNs. In medium to high dimensional state space settings, point-based solvers can provide POMDP solutions within a reasonable computational time. Despite still being uncommon among wind farm operators, POMDP solutions from point-based solvers offer rich interpretation aspects, as demonstrated in Chapter 3. Particularly, decision-makers can sequentially choose optimal actions only by tracking the belief states. Furthermore, when human factors and/or inevitable constraints intervene the POMDP-suggested strategies, analysis of alternative actions can be performed, devising subsequent strategies in order to minimize the associated economic regret.

Sensors can also be installed on offshore wind turbines in order to collect load information, however, there is a challenge that physical sensors are prone to damage in harsh marine environment and may become unreliable/not functional after a few years. When the installed sensors are no longer functional, data-driven tools can be used as virtual monitoring models by training artificial neural networks to learn complex relations from environment and operational conditions to the response of offshore wind turbines, as presented in Chapter 4. Fully probabilistic

virtual monitoring schemes can be developed through Bayesian neural networks, combining the strengths of deep learning and Bayesian inference principles (Chapter 5). Bayesian neural networks are particularly suited for stochastic modeling and uncertainty quantification in the predictions. Their capability to announce higher model uncertainty upon receiving out-of-training input data is advantageous for farm-wide monitoring applications where a fleet-leader-based model is directly deployed to other turbines.

In addition to inspection and monitoring data, information from virtual sensors can also be used to reduce epistemic uncertainties as well as to enable more informed maintenance decisions (Chapter 4 & 6). The benefit of a virtual monitoring model can be systematically quantified through a value of information analysis, computing the reduction of the expected life-cycle cost when the virtual monitoring model is deployed. In this regard, virtual monitoring model can also be incorporated within the sequential decision-making problem, by modeling monitoring observations conditional on the presence of physical or virtual sensors and following the integrated POMDP dynamics presented in Chapter 6.

The virtual monitoring integrated life-cycle management of offshore wind structures involves high dimensional state, action, and observation spaces. In this case, multi-agent reinforcement learning algorithms can be used to approximate POMDP solution, instead of point-based solvers. The decision-making problem then relies on the decentralized POMDP formulations, yet the value function or policies are parameterized with artificial neural networks. The policies obtained from the life-cycle management planning problem provide not only inspection and maintenance strategies but also virtual monitoring related decisions including sensor (re-)installation and re-calibration of the virtual model.

In particular, the original contributions can be shortlisted as:

i) modeling of stochastic fatigue failure criteria and deterioration models through dynamic Bayesian networks and partially observable Markov decision processes;

ii) interpretation and analysis of POMDP-based inspection and maintenance policies in order to accelerate implementation in the offshore wind energy industry;

iii) development of a fully probabilistic virtual monitoring model through combination of deep learning and Bayesian inference principles and application in a real-world case study, demonstrating uncertainty quantification capability and farm-wide applicability;

iv) modeling of a stochastic uncertainty propagation scheme for uncertainty-aware virtual monitoring methods;

v) development of an algorithmic framework for optimal monitoring, inspection, and maintenance planning of offshore wind structures using virtual sensors;

vi) application of the proposed framework for the life-cycle management of an offshore wind support structure and quantifying the value of virtual monitoring information.

7.2 Suggestions for further research

Further research directions are here suggested for future scientific exploration in the aspects of probabilistic virtual monitoring and integrated life-cycle management planning. In the probabilistic virtual monitoring context, deeper uncertainty quantification analysis, including the development of sophisticated principled metrics are strongly encouraged. The theoretical formulation of Bayesian neural networks and Bayesian inference forms a principled mathematical framework for probabilistic learning. However, simplifications are often imposed in the learning algorithms or during the training in order to approximate the Bayesian inference of high dimensional models within a reasonable computational time, e.g., assumption of posterior mean field weights, thus reducing the learning capacity. The investigation of efficient inference approximation methods, for instance, Hamiltonian Markov Chain Monte Carlo [185], stochastic gradient Langevin dynamics [186], are recommended. Thorough comparative analyses between Bayesian neural networks and other probabilistic models such as deep ensembles and Gaussian processes are also worth-conducting.

With respect to the scope of the life-cycle management planning, this work was mainly focused on the operational phase of already existing structures. This could be further extended to a complete life-cycle management from the design stage to decommissioning, considering not only gradual deterioration processes, but also the potential occurrence of extreme events. Whereas this work mainly relies on data-based models for virtual monitoring, it would also be worth-exploring the fusion of information from multiple virtual monitoring systems in the decision-making. For instance, when the management of an entire wind farm is carried out, information over the whole support structure of a turbine can be obtained from its physics-based virtual model [187] and farm-wide information can be retrieved from a data-driven probabilistic model [176].

References

[1] BS7910, "Guide to methods for assessing the acceptability of flaws in metallic structures," British Standards, BSI Knowledge, 389 Chiswick High Rd, Chiswick, London W4 4AL, United Kingdom, 2019.

[2] JCSS, "JCSS probabilistic model code," tech. rep., Joint Committee on Structural Safety, April 2011.

[3] H. Y. Chung, L. Manuel, and K. H. Frank, "Optimal inspection scheduling of steel bridges using nondestructive testing techniques.," *Journal of Bridge Engineering*, vol. 11, no. 3, pp. 305–319, 2006.

[4] D. Pérez-Campuzano, E. Gómez de las Heras-Carbonell, C. Gallego-Castillo, and A. Cuerva, "Modelling damage equivalent loads in wind turbines from general operational signals: Exploration of relevant input selection methods using aeroelastic simulations," *Wind Energy*, vol. 21, no. 6, pp. 441–459, 2018.

[5] H. Raiffa and R. Schlaifer, *Applied Statistical Decision Theory.* Harvard University Graduate School of Business Administration (Division of Research), 1961.

[6] H. Madsen, J. D. Sørensen, and R. Olesen, "Optimal inspection planning for fatigue damage of offshore structures," in *Proceedings of the 5th International Conference on Structural Safety and Reliability: Structural Safety and Reliability*, pp. 2099–2106, American Society of Civil Engineers, 1990.

[7] M. H. Faber, "Risk-based inspection: The framework," *Structural Engineering International*, vol. 12, no. 3, pp. 186–195, 2002.

[8] J. J. Nielsen and J. D. Sørensen, "Risk-based operation and maintenance planning for offshore wind turbines," in *Reliability and Optimization of Structural Systems*, 2010.

[9] J. Goyet, D. Straub, and M. H. Faber, "Risk-based inspection planning of offshore installations," *Structural Engineering International*, vol. 12, pp. 200–208, 8 2002.

[10] J. G. Rangel-Ramírez and J. D. Sørensen, "Risk-based inspection planning optimisation of offshore wind turbines," *Structure and Infrastructure Engineering*, vol. 8, no. 5, pp. 473–481, 2012.

[11] D. Straub, "Stochastic modeling of deterioration processes through dynamic Bayesian networks," *Journal of Engineering Mechanics*, vol. 135, no. 10, pp. 1089–1099, 2009.

[12] J. Luque and D. Straub, "Risk-based optimal inspection strategies for structural systems using dynamic Bayesian networks," *Structural Safety*, vol. 76, pp. 68–80, 1 2019.

[13] E. Bismut and D. Straub, "Optimal adaptive inspection and maintenance planning for deteriorating structural systems," *Reliability Engineering & System Safety*, vol. 215, p. 107891, 2021.

[14] D. Y. Yang and D. M. Frangopol, "Probabilistic optimization framework for inspection/repair planning of fatigue-critical details using dynamic Bayesian networks," *Computers & Structures*, vol. 198, pp. 40–50, 2018.

[15] I. Tien and A. Der Kiureghian, "Reliability assessment of critical infrastructure using Bayesian networks," *Journal of Infrastructure Systems*, vol. 23, no. 4, p. 04017025, 2017.

[16] D. Straub and I. Papaioannou, "Bayesian updating with structural reliability methods," *Journal of Engineering Mechanics*, vol. 141, no. 3, p. 34, 2015.

[17] J. Luque and D. Straub, "Reliability analysis and updating of deteriorating systems with dynamic Bayesian networks," *Structural Safety*, vol. 62, pp. 34–46, 2016.

[18] R. B. Corotis, J. H. Ellis, and M. Jiang, "Modeling of risk-based inspection, maintenance and life-cycle cost with partially observable Markov decision processes.," *Structure and Infrastructure Engineering*, vol. 1, no. 1, pp. 75–84, 2005.

[19] M. Memarzadeh and M. Pozzi, "Integrated inspection scheduling and maintenance planning for infrastructure systems.," *Computer-Aided Civil and Infrastructure Engineering*, vol. 31, no. 6, pp. 403–415, 2016.

[20] C. A. Robelin and S. Madanat, "History-dependent bridge deck maintenance and replacement optimization with Markov decision processes.," *Journal of Infrastructure Systems*, vol. 13, September 2007.

[21] G. Shani, J. Pineau, and R. Kaplow, "A survey of point-based POMDP solvers," *Autonomous Agents and Multi-Agent Systems*, vol. 27, no. 1, pp. 1–51, 2013.

[22] H. Kurniawati, D. Hsu, and W. Sun Lee, "SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces," in *Proceedings of Robotics: Science and Systems IV*, (Zurich, Switzerland), 2008.

[23] T. Smith and R. Simmons, "Focused real-time dynamic programming for MDPs: Squeezing more out of a heuristic," *Proceedings of the National Conference on Artificial Intelligence*, vol. 2, no. January, pp. 1227–1232, 2006.

[24] M. T. J. Spaan and N. Vlassis, "Perseus: Randomized point-based value iteration for POMDPs," *Journal of Artificial Intelligence Research*, vol. 24, pp. 195–220, 2005.

[25] M. Memarzadeh, M. Pozzi, and J. Z. Kolter, "Optimal planning and learning in uncertain environments for the management of wind farms," *Journal of Computing in Civil Engineering*, vol. 29, p. 04014076, 9 2015.

[26] K. G. Papakonstantinou and M. Shinozuka, "Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part I: Theory," *Reliability Engineering and System Safety*, vol. 130, pp. 214–224, 2014.

[27] K. G. Papakonstantinou and M. Shinozuka, "Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part II: POMDP implementation," *Reliability Engineering and System Safety*, vol. 130, pp. 214–224, 2014.

[28] D. Y. Yang and D. M. Frangopol, "Risk-based inspection planning of deteriorating structures," *Structure and Infrastructure Engineering*, vol. 18, no. 1, pp. 109–128, 2022.

[29] P. G. Morato, K. G. Papakonstantinou, C. P. Andriotis, J. S. Nielsen, and P. Rigo, "Optimal inspection and maintenance planning for deteriorating structural components through dynamic Bayesian networks and Markov decision processes," *Structural Safety*, vol. 94, p. 102140, 2022.

[30] J. J. Nielsen and J. D. Sørensen, "On risk-based operation and maintenance of offshore wind turbine components," *Reliability Engineering and System Safety*, vol. 96, no. 1, pp. 218–229, 2011.

[31] J. S. Nielsen and J. D. Sørensen, "Methods for risk-based planning of O&M of wind turbines," *Energies*, vol. 7, pp. 6645–6664, 10 2014.

[32] E. Bismut, J. Luque, and D. Straub, "Optimal prioritization of inspections in structural systems considering component interactions and interdependence," in *Proceedings of the 12th International Conference on Structural Safety & Reliability ICOSSAR 2017, Vienna, Austria*, 2017.

[33] E. Bismut and D. Straub, "Adaptive direct policy search for inspection and maintenance planning in structural systems," in *Proceedings of the 6th International Symposium on Life-Cycle Civil Engineering (IALCCE)*, 2018.

[34] J. S. Nielsen and J. D. Sørensen, "Computational framework for risk-based planning of inspections, maintenance and condition monitoring using discrete Bayesian networks," *Structure and Infrastructure Engineering*, vol. 14, no. 8, pp. 1082–1094, 2018.

[35] K. G. Papakonstantinou, C. P. Andriotis, and M. Shinozuka, "POMDP and MOMDP solutions for structural life-cycle cost minimization under partial and mixed observability," *Structure and Infrastructure Engineering*, vol. 14, no. 7, pp. 869–882, 2018.

[36] P. G. Morato, J. S. Nielsen, A. Q. Mai, and P. Rigo, "POMDP based maintenance optimization of offshore wind substructures including monitoring," in *13th International Conference on Applications of Statistics and Probability in Civil Engineering, ICASP 2019*, pp. 270–277, 2019.

[37] M. Memarzadeh, M. Pozzi, and J. Z. Kolter, "Hierarchical modeling of systems with similar components: A framework for adaptive monitoring and control," *Reliability Engineering & System Safety*, vol. 153, pp. 159–169, 2016.

[38] C. P. Andriotis and K. G. Papakonstantinou, "Managing engineering systems with large state and action spaces through deep reinforcement learning," *Reliability Engineering & System Safety*, vol. 191, p. 106483, 2019.

[39] W. Zhu, B. Castanier, and B. Bettayeb, "A dynamic programming-based maintenance model of offshore wind turbine considering logistic delay and weather condition," *Reliability Engineering & System Safety*, vol. 190, p. 106512, 2019.

[40] C. P. Andriotis and K. G. Papakonstantinou, "Deep reinforcement learning driven inspection and maintenance planning under incomplete information and constraints," *Reliability Engineering & System Safety*, vol. 212, p. 107551, 2021.

[41] N. Zhang and W. Si, "Deep reinforcement learning for condition-based maintenance planning of multi-component systems under dependent competing risks," *Reliability Engineering & System Safety*, vol. 203, p. 107094, 2020.

[42] A. Kamariotis, D. Straub, and E. Chatzi, "Optimal maintenance decisions supported by SHM: A benchmark study," in *Life-Cycle Civil Engineering: Innovation, Theory and Practice*, pp. 679–686, CRC Press, 2021.

[43] P. G. Morato, C. P. Andriotis, K. G. Papakonstantinou, and P. Rigo, "Inference and dynamic decision-making for deteriorating systems with probabilistic dependencies through Bayesian networks and deep reinforcement learning," *Reliability Engineering & System Safety*, vol. 235, p. 109144, 2023.

[44] P. G. Morato, K. G. Papakonstantinou, C. P. Andriotis, and P. Rigo, "Managing offshore wind turbines through Markov decision processes and dynamic Bayesian networks," in *International Conference on Structural Safety and Reliability*, 2022.

[45] R. Mohammadi and Q. He, "A deep reinforcement learning approach for rail renewal and maintenance planning," *Reliability Engineering & System Safety*, vol. 225, p. 108615, 2022.

[46] J. Lee and M. Mitici, "Deep reinforcement learning for predictive aircraft maintenance using probabilistic remaining-useful-life prognostics," *Reliability Engineering & System Safety*, vol. 230, p. 108908, 2023.

[47] A. Kamariotis, E. Chatzi, and D. Straub, "A framework for quantifying the value of vibration-based structural health monitoring," *Mechanical Systems and Signal Processing*, vol. 184, p. 109708, 2023.

[48] H. Li and C. Guedes Soares, "Reliability and opportunistic maintenance of floating offshore wind turbines," *Advances in Reliability and Maintainability Methods and Engineering Applications: Essays in Honor of Professor Hong-Zhong Huang on his 60th Birthday*, pp. 331–354, 2023.

[49] P. Leroy, P. G. Morato, J. Pisane, A. Kolios, and D. Ernst, "IMP-MARL: A Suite of environments for large-scale infrastructure management planning via MARL," in *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023.

[50] G. Arcieri, C. Hoelzl, O. Schwery, D. Straub, K. G. Papakonstantinou, and E. Chatzi, "Bridging pomdps and bayesian decision making for robust maintenance planning under model uncertainty: An application to railway systems," *Reliability Engineering & System Safety*, vol. 239, p. 109496, 2023.

[51] P. Bhustali and C. Andriotis, "Assessing the optimality of decentralized inspection and maintenance policies for stochastically degrading engineering systems," in *BNAIC/BeNeLearn 2023: Joint International Scientific Conferences on AI and Machine Learning*, 2023.

[52] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.

[53] T. Ming, "Multi-agent reinforcement learning: Independent versus cooperative agents," in *Proceedings of the Tenth International Conference on Machine Learning (ICML 1993), San Francisco, CA, USA*, pp. 330–337, 1993.

[54] T. Rashid, M. Samvelyan, C. Schroeder, G. Farquhar, J. Foerster, and S. Whiteson, "QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning," in *International Conference on Machine Learning*, pp. 4295–4304, PMLR, 2018.

[55] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch, "Multi-agent actor-critic for mixed cooperative-competitive environments," in *Advances in Neural Information Processing Systems*, vol. 30, Curran Associates, Inc., 2017.

[56] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[57] C. Yu, M. Fang, Y. Yang, X. Zhang, T. Zhou, and L. Wang, "The surprising effectiveness of PPO in cooperative multi-agent games," *arXiv preprint arXiv:2103.01955*, 2021.

[58] S. Yoon, "Virtual sensing in intelligent buildings and digitalization," *Automation in Construction*, vol. 143, p. 104578, 2022.

[59] S. Kabadayi, A. Pridgen, and C. Julien, "Virtual sensors: Abstracting data from physical sensors," in *2006 International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM'06)*, pp. 6–pp, IEEE, 2006.

[60] D. Martin, N. Kühl, and G. Satzger, "Virtual sensors," *Business & Information Systems Engineering*, vol. 63, p. 315–323, 06 2021.

[61] K. Maes, A. Iliopoulos, W. Weijtjens, C. Devriendt, and G. Lombaert, "Dynamic strain estimation for fatigue assessment of an offshore monopile wind turbine using filtering and modal expansion algorithms," *Mechanical Systems and Signal Processing*, vol. 76-77, pp. 592–611, 2016.

[62] N. Noppe, A. Iliopoulos, W. Weijtjens, and C. Devriendt, "Full load estimation of an offshore wind turbine based on SCADA and accelerometer data," *Journal of Physics: Conference Series*, vol. 753, p. 072025, September 2016.

[63] K. Tatsis, V. Dertimanis, I. Abdallah, and E. Chatzi, "A substructure approach for fatigue assessment on wind turbine support structures using output-only measurements," *Procedia engineering*, vol. 199, pp. 1044–1049, 2017.

[64] H. C. Kim, M. H. Kim, and E. C. Do, "Structural health monitoring of towers and blades for floating offshore wind turbines using operational modal analysis and modal properties with numerical-sensor signals," *Ocean Engineering*, vol. 188, p. 106226, 2019.

[65] J. Kullaa, "Bayesian virtual sensing in structural dynamics," *Mechanical Systems and Signal Processing*, vol. 115, pp. 497–513, 2019.

[66] Y. Xu, G. Nikitas, T. Zhang, Q. Han, M. Chryssanthopoulos, S. Bhattacharya, and Y. Wang, "Support condition monitoring of offshore wind turbines using model updating techniques," *Structural Health Monitoring*, vol. 19, no. 4, pp. 1017–1031, 2020.

[67] F. de Nolasco Santos, N. Noppe, W. Weijtjens, and C. Devriendt, "SCADA-based neural network thrust load model for fatigue assessment: Cross validation with in-situ measurements," *Journal of Physics: Conference Series*, vol. 1618, p. 022020, September 2020.

[68] M. Henkel, J. Häfele, W. Weijtjens, C. Devriendt, C. Gebhardt, and R. Rolfes, "Strain estimation for offshore wind turbines with jacket substructures using dual-band modal expansion," *Marine Structures*, vol. 71, p. 102731, 2020.

[69] M.-S. Nabiyan, F. Khoshnoudian, B. Moaveni, and H. Ebrahimian, "Mechanics-based model updating for identification and virtual sensing of an offshore wind turbine using sparse measurements," *Structural Control and Health Monitoring*, vol. 28, no. 2, p. e2647, 2021.

[70] N. Dimitrov and T. Göçmen, "Virtual sensors for wind turbines with machine learning-based time series models," *Wind Energy*, vol. 25, no. 9, pp. 1626–1645, 2022.

[71] F. Pimenta, J. Pacheco, S. Pereira, and F. Magalhães, "Reconstructing the bending moments time history of wind turbine tower from acceleration measurements using Gaussian processes," *Journal of Physics: Conference Series*, vol. 2265, p. 032080, May 2022.

[72] F. de Nolasco Santos, N. Noppe, W. Weijtjens, and C. Devriendt, "Data-driven farm-wide fatigue estimation on jacket-foundation OWTs for multiple SHM setups," *Wind Energy Science*, vol. 7, no. 1, pp. 299–321, 2022.

[73] M. Tarpø, S. Amador, E. Katsanos, M. Skog, J. Gjødvad, and R. Brincker, "Data-driven virtual sensing and dynamic strain estimation for fatigue analysis of offshore wind turbine using principal component analysis," *Wind Energy*, vol. 25, no. 3, pp. 505–516, 2022.

[74] S. Vettori, E. Di Lorenzo, B. Peeters, and E. Chatzi, "Virtual sensing for wind turbine blade full field response estimation in operational modal analysis," in *Model Validation and Uncertainty Quantification, Volume 3: Proceedings of the 39th IMAC, A Conference and Exposition on Structural Dynamics 2021*, pp. 49–52, Springer, 2022.

[75] F. C. Mehlan, A. R. Nejad, and Z. Gao, "Digital twin based virtual sensor for online fatigue damage monitoring in offshore wind turbine drivetrains," *Journal of Offshore Mechanics and Arctic Engineering*, vol. 144, p. 060901, 10 2022.

[76] X. Li and W. Zhang, "Physics-informed deep learning model in wind turbine response prediction," *Renewable Energy*, vol. 185, pp. 932–944, 2022.

[77] Z. Zhang, C. Peng, G. Wang, Z. Ju, and L. Ma, "Virtual sensing based on hierarchical Bayesian modeling framework using a Laplace-based Gibbs sampler," *Applied Mathematical Modelling*, vol. 124, pp. 167–191, 2023.

[78] S. Vettori, E. Di Lorenzo, B. Peeters, M. Luczak, and E. Chatzi, "An adaptive-noise augmented kalman filter approach for input-state estimation in structural dynamics," *Mechanical Systems and Signal Processing*, vol. 184, p. 109654, 2023.

[79] J. Zou, E. M. Lourens, and A. Cicirello, "Virtual sensing of subsoil strain response in monopile-based offshore wind turbines via Gaussian process latent force models," *Mechanical Systems and Signal Processing*, vol. 200, p. 110488, 2023.

[80] A. Iliopoulos, W. Weijtjens, D. Van Hemelrijck, and C. Devriendt, "Fatigue assessment of offshore wind turbines on monopile foundations using multi-band modal expansion," *Wind Energy*, 2017.

[81] N. Cosack, *Fatigue load monitoring with standard wind turbine signals*. PhD thesis, Universitat Stuttgart, 2010.

[82] D. Singh, R. P. Dwight, K. Laugesen, L. Beaudet, and A. Viré, "Probabilistic surrogate modeling of offshore wind-turbine loads with chained Gaussian processes," *Journal of Physics: Conference Series*, vol. 2265, p. 032070, May 2022.

[83] Z. Zhang and C. Sun, "Structural damage identification via physics-guided machine learning: A methodology integrating pattern recognition with finite element model updating," *Structural Health Monitoring*, vol. 20, no. 4, pp. 1675–1688, 2021.

[84] C. M. Atkinson, T. W. Long, and E. L. Hanzevack, "Virtual sensing: A neural network-based intelligent performance and emissions prediction system for on-board diagnostics and engine control," tech. rep., SAE Technical Paper, 1998.

[85] J. Kim, J. Park, S. Shin, Y. Lee, K. Min, S. Lee, and M. Kim, "Prediction of engine nox for virtual sensor using deep neural network and genetic algorithm," *Oil & Gas Science and Technology–Revue d'IFP Energies nouvelles*, vol. 76, p. 72, 2021.

[86] S. B. Sun, Y. Y. He, S. D. Zhou, and Z. J. Yue, "A data-driven response virtual sensor technique with partial vibration measurements using convolutional neural network," *Sensors*, vol. 17, no. 12, p. 2888, 2017.

[87] 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," *Structural Health Monitoring Based on Data Science Techniques*, pp. 347–367, 2022.

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

[89] S. J. Gibson, T. J. Rogers, and E. J. Cross, "Integrating physical knowledge into Gaussian process regression models for probabilistic fatigue assessment," in *European Workshop on Structural Health Monitoring*, pp. 472–481, Springer, 2022.

[90] IEC61400-3-1, "Wind energy generation systems - Part 3-1: Design requirements for fixed offshore wind turbines," international standard, International Electrotechnical Commission, CH-1211 Geneva 20, Switzerland, 2019.

[91] DNV-OS-J101, "Design of offshore wind turbine structures.," offshore standard, DNV, 1363 Høvik, Norway, 2010.

[92] M. C. Kennedy and A. O'Hagan, "Bayesian calibration of computer models," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 63, no. 3, pp. 425–464, 2001.

[93] C. Devriendt, F. Magalhães, W. Weijtjens, G. De Sitter, Á. Cunha, and P. Guillaume, "Structural health monitoring of offshore wind turbines using automated operational modal analysis," *Structural Health Monitoring*, vol. 13, pp. 644–659, November 2014.

[94] W. Weijtjens, N. Noppe, T. Verbelen, and C. Devriendt, "Fleet-wise structural health monitoring of (offshore) wind turbine foundations," in *Proceedings of the 8th European Workshop on Structural Health Monitoring, Bilbao, Spain*, pp. 5–8, 2016.

[95] B. Piascik, J. Vickers, D. Lowry, S. Scotti, J. Stewart, and A. Calomino, "Materials, structures, mechanical systems, and manufacturing roadmap," *NASA TA*, pp. 12–2, 2012.

[96] P. G. Morato, *Optimal Inspection and Maintenance Planning for Deteriorating Structures via Markov Decision Processes and Deep Reinforcement Learning: Application to Offshore Wind Substructures.* PhD thesis, University of Liege, Liege, September 2021.

[97] C. Agrell, K. Rognlien Dahl, and A. Hafver, "Optimal sequential decision making with probabilistic digital twins: Theoretical foundations," *SN Applied Sciences*, vol. 5, no. 4, p. 114, 2023.

[98] J. Xia and G. Zou, "Operation and maintenance optimization of offshore wind farms based on digital twin: A review," *Ocean Engineering*, vol. 268, p. 113322, 2023.

[99] J. Jorgensen, M. Hodkiewicz, E. Cripps, and G. M. Hassan, "Requirements for the application of the digital twin paradigm to offshore wind turbine structures for uncertain fatigue analysis," *Computers in Industry*, vol. 145, p. 103806, 2023.

[100] M. Fujita, G. Schall, and R. Rackwitz, "Adaptive reliability-based inspection strategies for structures subject to fatigue.," in *Proceedings of the 5th ICOSSAR*, vol. 2, (United States), pp. 1619–1626, 1989.

[101] J. Goyet, A. Maroini, M. H. Faber, and J. C. Paygnard, "Optimal inspection and repair planning: Case studies using IMREL software.," in *Proceedings of the 13th international conference on offshore mechanics and Arctic engineering OMAE-13*, vol. 2, (United States), pp. 325–333, 1994.

[102] M. H. Faber, J. D. Sørensen, and I. Kroon, "Optimal inspection strategies for offshore structural systems.," in *Proceedings of the 11th international conference on offshore mechanics and Arctic engineering OMAE-11*, vol. 2, (Canada), pp. 145–152, 1992.

[103] D. Straub, *Generic Approaches to Risk Based Inspection Planning for Steel Structures*. PhD thesis, Swiss Federal Institute of Technology Zürich (ETH), 2004.

[104] D. Straub and M. H. Faber, "Computational aspects of risk-based inspection planning.," *Computer-Aided Civil and Infrastructure Engineering*, vol. 21, no. 3, pp. 179–192, 2006.

[105] D. Straub and M. H. Faber, "Risk based inspection planning for structural systems," *Structural Safety*, vol. 27, no. 4, pp. 335–355, 2005.

[106] T. Moan, "Reliability-based management of inspection, maintenance and repair of offshore structures.," *Structure and Infrastructure Engineering*, vol. 1, no. 1, pp. 33–62, 2005.

[107] S. M. Soliman, D. M. Frangopol, and A. Mondoro, "A probabilistic approach for optimizing inspection, monitoring, and maintenance actions against fatigue of critical ship details," *Structural Safety*, vol. 60, pp. 91–101, 2016.

[108] D. Frangopol and S. Kim, *Life-Cycle of Structures Under Uncertainty: Emphasis on Fatigue-Sensitive Civil and Marine Structures*. Boca Raton: CRC Press, 2019.

[109] S. Kim, B. Ge, and D. M. Frangopol, "Effective optimum maintenance planning with updating based on inspection information for fatigue-sensitive structures," *Probabilistic Engineering Mechanics*, vol. 58, p. 103003, 2019.

[110] K. G. Papakonstantinou, C. P. Andriotis, and M. Shinozuka, "Point-based POMDP solvers for life-cycle cost minimization of deteriorating structures.," in *Proceedings of 5th International Symposium on Life-Cycle Engineering IALCCE-2016*, (Netherlands), pp. 427–434, 2017.

[111] A. R. Dowling and C. H. A. Townley, "The effects of defects on structural failure: A two-criteria approach.," *International Journal of Pressure Vessels and Piping*, vol. 3, pp. 77–137, 1975.

[112] R. P. Harrison, I. Milne, and T. G. F. Gray, "Assessment of defects: The C.E.G.B. approach [and discussion].," *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 299, no. 1446, pp. 145–153, 1981.

[113] A. Fajuyigbe and F. Brennan, "Fitness-for-purpose assessment of cracked offshore wind turbine monopile.," *Marine Structures*, vol. 77, 2021.

[114] A. Q. Mai, J. D. Sørensen, and P. Rigo, "Updating failure probability of a welded joint in offshore wind turbine substructures.," in *Proceedings of The 35th International Conference on Ocean, Offshore and Arctic Engineering Conference*, vol. 3, (South Korea), 2016.

[115] N. Hlaing, P. G. Morato, P. Rigo, P. Amirafshari, A. Kolios, and J. S. Nielsen, "The effect of failure criteria on risk-based inspection planning of offshore wind support structures," in *Life-Cycle Civil Engineering: Innovation, Theory and Practice - Proceedings of the 7th International Symposium on Life-Cycle Civil Engineering, IALCCE 2020*, pp. 146–153, CRC Press, 2020.

[116] P. Amirafshari, F. Brennan, and A. Kolios, "A fracture mechanics framework for optimising design and inspection of offshore wind turbine support structures against fatigue failure.," *Wind Energy Science*, vol. 6, no. 3, pp. 677–699, 2021.

[117] DNV-RP-C210, "Probabilistic methods for planning of inspection for fatigue cracks in offshore structures.," recommended practice, DNV, Veritasveien 1, 1363 Høvik, Norway, 2019.

[118] P. Paris and F. Erdogan, "A critical analysis of crack propagation laws.," *Journal of Basic Engineering*, vol. 85, no. 4, pp. 528–533, 1963.

[119] O. Ditlevsen and H. O. Madsen, *Structural Reliability Methods*. Department of Mechanical Engineering. Technical University of Denmark., 2007.

[120] T. L. Anderson, *Fracture Mechanics: Fundamentals and Applications.* Taylor and Francis Group, June 2005.

[121] O. Dijkstra, "A fracture mechanics approach to the assessment of the remaining fatigue life of defective welded joints," in *Proceedings of International Association for Bridge and Structural Engineers (IABSE) Workshop*, (Lausanne), April 1990.

[122] A. P. Berens and P. W. Hovey, "AFWAL-TR-81-4160 Evaluation of NDE characterization.," tech. rep., University of Dayton Research Institute, Dayton, Ohio 45649, 1981.

[123] N. A. Macmillan and C. D. Creelman, *Detection Theory: A User's Guide.* Taylor & Francis, second ed., 2004.

[124] H. P. Hong, "Reliability analysis with nondestructive inspection.," *Structural Safety*, vol. 19, no. 4, pp. 383–395, 1997.

[125] I. Lotsberg, G. Sigurdsson, A. Fjeldstad, and T. Moan, "Probabilistic methods for planning of inspection for fatigue cracks in offshore structures.," *Marine Structures*, vol. 46, pp. 167–192, 2016.

[126] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988.

[127] J. Zhu and M. Collette, "A dynamic discretization method for reliability inference in dynamic Bayesian networks.," *Reliability Engineering & System Safety*, vol. 138, pp. 242–252, 2015.

[128] DNV-ST-0126, "Support structures for wind turbines.," standard, DNV, Veritasveien 1, 1363 Høvik, Norway, 2018.

[129] DNV-RP-C203, "Fatigue design of offshore steel structures.," recommended practice, DNV, Veritasveien 1, 1363 Høvik, Norway, 2016.

[130] V. Igwemezie, A. Mehmanparast, and A. Kolios, "Materials selection for XL wind turbine support structures: A corrosion-fatigue perspective.," *Marine Structures*, vol. 61, pp. 381–397, 2018.

[131] K. Wallin, *Fracture Toughness of Engineering Materials: Estimation and Application.* United Kingdom: EMAS Publishing, 2011.

[132] J. Newman and I. Raju, "An empirical stress-intensity factor equation for the surface crack.," *Engineering Fracture Mechanics*, vol. 15, no. 1, pp. 185–192, 1981.

[133] N. Hlaing, P. G. Morato, J. S. Nielsen, P. Amirafshari, A. Kolios, and P. Rigo, "Inspection and maintenance planning for offshore wind structural components: Integrating fatigue failure criteria with Bayesian networks and Markov decision processes.," *Structure and Infrastructure Engineering*, vol. 18, no. 7, pp. 983–1001, 2022.

[134] C. Hübler, W. Weijtjens, R. Rolfes, and C. Devriendt, "Reliability analysis of fatigue damage extrapolations of wind turbines using offshore strain measurements," *Journal of Physics: Conference Series*, vol. 1037, pp. 032–035, 2018.

[135] Q. A. Mai, W. Weijtjens, C. Devriendt, P. G. Morato, P. Rigo, and J. D. Sørensen, "Prediction of remaining fatigue life of welded joints in wind turbine support structures considering strain measurement and a joint distribution of oceanographic data," *Marine Structures*, vol. 66, pp. 307–322, 2019.

[136] J. Jonkman, S. Butterfield, W. Musial, and G. Scott, "Definition of a 5-MW Reference Wind Turbine for Offshore System Development ," tech. rep., National Renewable Energy Laboratory, 1617 Cole Boulevard, Golden, Colorado 80401-3393, 2009.

[137] J. Jonkman and W. Musial, "Offshore code comparison collaboration (OC3) for IEA Task 23 offshore wind technology and deployment," tech. rep., National Renewable Energy Laboratory, 1617 Cole Boulevard, Golden, Colorado 80401-3393, 2010.

[138] T. Fischer, W. de Vries, and B. Schmidt, "Upwind design dasis (WP4: Offshore foundations and support structures) ," tech. rep., Endowed Chair of Wind Energy (SWE) at the Institute of Aircraft Design, Universität Stuttgart, Allmandring 5B, 70550 Stuttgart, Germany, 2010.

[139] BS61400-3, "Wind Turbines - Part 3: Design requirements for offshore wind turbines," British Standards, BSI Knowledge, 389 Chiswick High Rd, Chiswick, London W4 4AL, United Kingdom, 2009.

[140] F. Chollet *et al.*, "Keras," 2015. Available at https://github.com/fchollet/keras.

[141] M. Farhan, R. Schneider, and S. Thöns, "Predictive information and maintenance optimization based on decision theory: A case study considering a welded joint in an offshore wind turbine support structure," *Structural Health Monitoring*, vol. 21, no. 1, pp. 185–207, 2022.

[142] H. Zhao, F. Xu, B. Liang, J. Zhang, and P. Song, "A condition-based opportunistic maintenance strategy for multi-component system," *Structural Health Monitoring*, vol. 18, no. 1, pp. 270–283, 2019.

[143] J. S. Nielsen, "Value of information of structural health monitoring with temporally dependent observations," *Structural Health Monitoring*, vol. 21, no. 1, pp. 165–184, 2022.

[144] M. Schedat, T. Faber, and A. Sivanesan, "Structural health monitoring concept to predict the remaining lifetime of the wind turbine structure," in *2016 International Conference on the Domestic Use of Energy (DUE)*, pp. 1–5, IEEE, 2016.

[145] J. Pacheco, F. Pimenta, S. Pereira, A. Cunha, and F. Magalhães, "Fatigue assessment of wind turbine towers: Review of processing strategies with illustrative case study," *Energies*, vol. 15, no. 13, p. 4782, 2022.

[146] K. Smarsly, D. Hartmann, and K. H. Law, "A computational framework for life-cycle management of wind turbines incorporating structural health monitoring," *Structural Health Monitoring*, vol. 12, no. 4, pp. 359–376, 2013.

[147] N. Hlaing, P. G. Morato, and P. Rigo, "Probabilistic virtual load monitoring of offshore wind substructures: A supervised learning approach," in *The Proceedings of the 32nd International Ocean and Polar Engineering Conference, ISOPE-2022*, vol. 4, pp. 3137–3144, 2022.

[148] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics*, vol. 3, pp. 422–440, 2021.

[149] J. Zou, A. Cicirello, A. Iliopoulos, and E. M. Lourens, "Gaussian process latent force models for virtual sensing in a monopile-based offshore wind turbine," in *EWSHM 2022: European Workshop on Structural Health Monitoring.*, pp. 290–298, Springer International Publishing, 2023.

[150] N. Noppe, C. Hübler, C. Devriendt, and W. Weijtjens, "Validated extrapolation of measured damage within an offshore wind farm using instrumented fleet leaders," *Journal of Physics: Conference Series*, vol. 1618, p. 022005, sep 2020.

[151] W. Weijtens, N. Noppe, T. Verbelen, A. Iliopoulos, and C. Devriendt, "Offshore wind turbine foundation monitoring, extrapolating fatigue measurements from fleet leaders to the entire wind farm," *Journal of Physics: Conference Series*, vol. 753, p. 092018, September 2016.

[152] R. Mbuvha, M. Jonsson, N. Ehn, and P. Herman, "Bayesian neural networks for one-hour ahead wind power forecasting," in *2017 IEEE 6th International Conference on Renewable Energy Research and Applications (ICRERA)*, pp. 591–596, IEEE, 2017.

[153] Y. Wang, Q. Hu, D. Meng, and P. Zhu, "Deterministic and probabilistic wind power forecasting using a variational Bayesian-based adaptive robust multi-kernel regression model," *Applied Energy*, vol. 208, pp. 1097–1112, 2017.

[154] R. Mbuvha, W. T. Mongwe, and T. Marwala, "Separable shadow Hamiltonian hybrid Monte Carlo for Bayesian neural network inference in wind speed forecasting," *Energy and AI*, vol. 6, pp. 100–108, 2021.

[155] L. V. Jospin, H. Laga, F. Boussaid, W. Buntine, and M. Bennamoun, "Hands-on Bayesian neural networks — A tutorial for deep learning users," *IEEE Computational Intelligence Magazine*, vol. 17, no. 2, pp. 29–48, 2022.

[156] H. Ritter, A. Botev, and D. Barber, "A scalable Laplace approximation for neural networks," in *ICLR 2018- Conference Track Proceedings: 6th International Conference on Learning Representations*, 2018.

[157] K. Osawa, S. Swaroop, M. E. E. Khan, A. Jain, R. Eschenhagen, R. E. Turner, and R. Yokota, "Practical deep learning with Bayesian principles," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[158] F. Wenzel, K. Roth, B. Veeling, J. Swiatkowski, L. Tran, S. Mandt, J. Snoek, T. Salimans, R. Jenatton, and S. Nowozin, "How good is the Bayes posterior in deep neural networks really?," in *International Conference on Machine Learning*, pp. 10248–10259, PMLR, 2020.

[159] 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*, vol. 21, no. 6, pp. 1087–1092, 1953.

[160] W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika*, vol. 57, pp. 97–109, 04 1970.

[161] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, vol. 2. Springer, August 1999.

[162] A. Graves, "Practical variational inference for neural networks," in *Advances in Neural Information Processing Systems*, vol. 24, pp. 2348–2356, 2011.

[163] S. Kullback and R. A. Leibler, "On information and sufficiency.," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79 – 86, 1951.

[164] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning: with Applications in R*. Springer, 2013.

[165] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine learning*, vol. 8, pp. 229–256, 1992.

[166] L. Richter, A. Boustati, N. Nüsken, F. Ruiz, and O. D. Akyildiz, "VarGrad: A low-variance gradient estimator for variational inference," *Advances in Neural Information Processing Systems*, vol. 33, pp. 13481–13492, 2020.

[167] Y. Bengio, N. Léonard, and A. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," *arXiv preprint arXiv:1308.3432*, 2013.

[168] S. Depeweg, J. M. Hernandez Lobato, S. Udluft, and T. Runkler, "Sensitivity analysis for predictive uncertainty in Bayesian neural networks," in *ESANN 2018 - Proceedings European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, (Bruges (Belgium)), pp. 279–284, 25-27 April 2018.

[169] S. Depeweg, J. M. Hernandez Lobato, F. Doshi Velez, and S. Udluft, "Decomposition of uncertainty in Bayesian deep learning for efficient and risk-sensitive learning," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80, pp. 1184–1193, PMLR, 10-15 July 2018.

[170] P. Wang, N. C. Bouaynaya, L. Mihaylova, J. Wang, Q. Zhang, and R. He, "Bayesian neural networks uncertainty quantification with cubature rules," in *2020 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–7, 2020.

[171] Agentschap Maritieme Dienstverlening en Kust, "Meetnet Vlaamse Banken - Westhinder Buoy." https://meetnetvlaamsebanken.be/, 2021. Accessed: 05/11/2021.

[172] Y. Wen, P. Vicol, J. Ba, D. Tran, and R. Grosse, "Flipout: Efficient pseudo-independent weight perturbations on mini-batches," in *ICLR 2018- Conference Track Proceedings: 6th International Conference on Learning Representations*, 2018.

[173] W. Ostachowicz, R. Soman, and P. Malinowski, "Optimization of sensor placement for structural health monitoring: A review," *Structural Health Monitoring*, vol. 18, no. 3, pp. 963–988, 2019.

[174] D. Straub, "Value of information analysis with structural reliability methods," *Structural Safety*, vol. 49, pp. 75–85, 2014. Special Issue In Honor of Professor Wilson H. Tang.

[175] C. P. Andriotis, K. G. Papakonstantinou, and E. N. Chatzi, "Value of structural health information in partially observable stochastic environments," *Structural Safety*, vol. 93, p. 102072, 2021.

[176] N. Hlaing, P. G. Morato, F. de Nolasco Santos, W. Weijtjens, C. Devriendt, and P. Rigo, "Farm-wide virtual load monitoring for offshore wind structures via Bayesian neural networks," *Structural Health Monitoring*, vol. 0, pp. 0–0, 8 2023.

[177] G. Stefanou, "The stochastic finite element method: Past, present and future," *Computer Methods in Applied Mechanics and Engineering*, vol. 198, no. 9, pp. 1031–1051, 2009.

[178] B. Grimstad, M. Hotvedt, A. T. Sandnes, O. Kolbjørnsen, and L. S. Imsland, "Bayesian neural networks for virtual flow metering: An empirical study," *Applied Soft Computing*, vol. 112, p. 107776, 2021.

[179] S. Madanat and M. Ben-Akiva, "Optimal inspection and repair policies for infrastructure facilities," *Transportation science*, vol. 28, no. 1, pp. 55–62, 1994.

[180] A. Du and A. Ghavidel, "Parameterized deep reinforcement learning-enabled maintenance decision-support and life-cycle risk assessment for highway bridge portfolios," *Structural Safety*, vol. 97, p. 102221, 2022.

[181] X. Lei, Y. Xia, L. Deng, and L. Sun, "A deep reinforcement learning framework for life-cycle maintenance planning of regional deteriorating bridges using inspection data," *Structural and Multidisciplinary Optimization*, vol. 65, no. 5, p. 149, 2022.

[182] E. Bismut and D. Straub, "A unifying review of NDE models towards optimal decision support," *Structural Safety*, vol. 97, p. 102213, 2022.

[183] J. Song and W. H. Kang, "System reliability and sensitivity under statistical dependence by matrix-based system reliability method," *Structural Safety*, vol. 31, no. 2, pp. 148–156, 2009.

[184] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *international conference on machine learning*, pp. 1050–1059, PMLR, 2016.

[185] K. M. Hanson, "Markov Chain Monte Carlo posterior sampling with the Hamiltonian method," in *Medical Imaging 2001: Image Processing*, vol. 4322, pp. 456–467, SPIE, 2001.

[186] M. Welling and Y. W. Teh, "Bayesian learning via stochastic gradient langevin dynamics," in *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 681–688, 2011.

[187] M. Henkel, W. Weijtjens, and C. Devriendt, "Fatigue stress estimation for submerged and sub-soil welds of offshore wind turbines on monopiles using modal expansion," *Energies*, vol. 14, no. 22, p. 7576, 2021.

Appendix A

Network architectures and training environments

Table A1. Comparison between deterministic and Bayesian neural networks for sensor configurations "SCADA + wave, SCADA + wave + accelerometer (LAT-017), SCADA + wave + accelerometer (LAT-038), SCADA + wave + accelerometer (LAT-077), SCADA + wave + accelerometers (LAT-017, 038)".

	Deterministic NN	Bayesian NN
Optimizer	Adamax (lr=0.001)	Adam (lr=0.0003)
Batch size	32	1024
No of episodes	200	2000
Early stopping	Validation loss (Patience = 5)	Training loss (Patience = 30)
No of neurons:		
Hidden layers	64, 128, 64	31, 64, 32
Output layer	2	4
Distribution layer	-	2

Table A2. Comparison between deterministic and Bayesian neural networks for sensor configuration "SCADA + wave + accelerometers (LAT-017, 038, 077)".

	Deterministic NN	Bayesian NN
Optimizer	Adamax (lr=0.001)	Adam (lr=0.00035)
Batch size	32	1024
No of episodes	200	2000
Early stopping	Validation loss (Patience = 5)	Training loss (Patience = 30)
No of neurons:		
Hidden layers	64, 128, 64	31, 64, 32
Output layer	2	4
Distribution layer	-	2

Table A3. Comparison between deterministic and Bayesian neural networks for sensor configurations "SCADA, SCADA + accelerometer (LAT-017), SCADA + accelerometer (LAT-038), SCADA + accelerometer (LAT-077)".

	Deterministic NN	Bayesian NN
Optimizer	Adamax (lr=0.001)	Adam (lr=0.0002)
Batch size	32	1024
No of episodes	200	2000
Early stopping	Validation loss (Patience = 5)	Training loss (Patience = 30)
No of neurons:		
Hidden layers	64, 128, 64	32, 64, 32
Output layer	2	4
Distribution layer	-	2

Table A4. Comparison between deterministic and Bayesian neural networks for sensor configuration "SCADA + accelerometers (LAT-017, 038), SCADA + accelerometers (LAT-017, 038, 077)".

	Deterministic NN	Bayesian NN
Optimizer	Adamax (lr=0.001)	Adam (lr=0.0003)
Batch size	32	1024
No of episodes	200	2000
Early stopping	Validation loss (Patience = 5)	Training loss (Patience = 30)
No of neurons:		
Hidden layers	64, 128, 64	32, 64, 32
Output layer	2	4
Distribution layer	-	2

Appendix B

Probability distributions of each input variable for the analyzed turbines

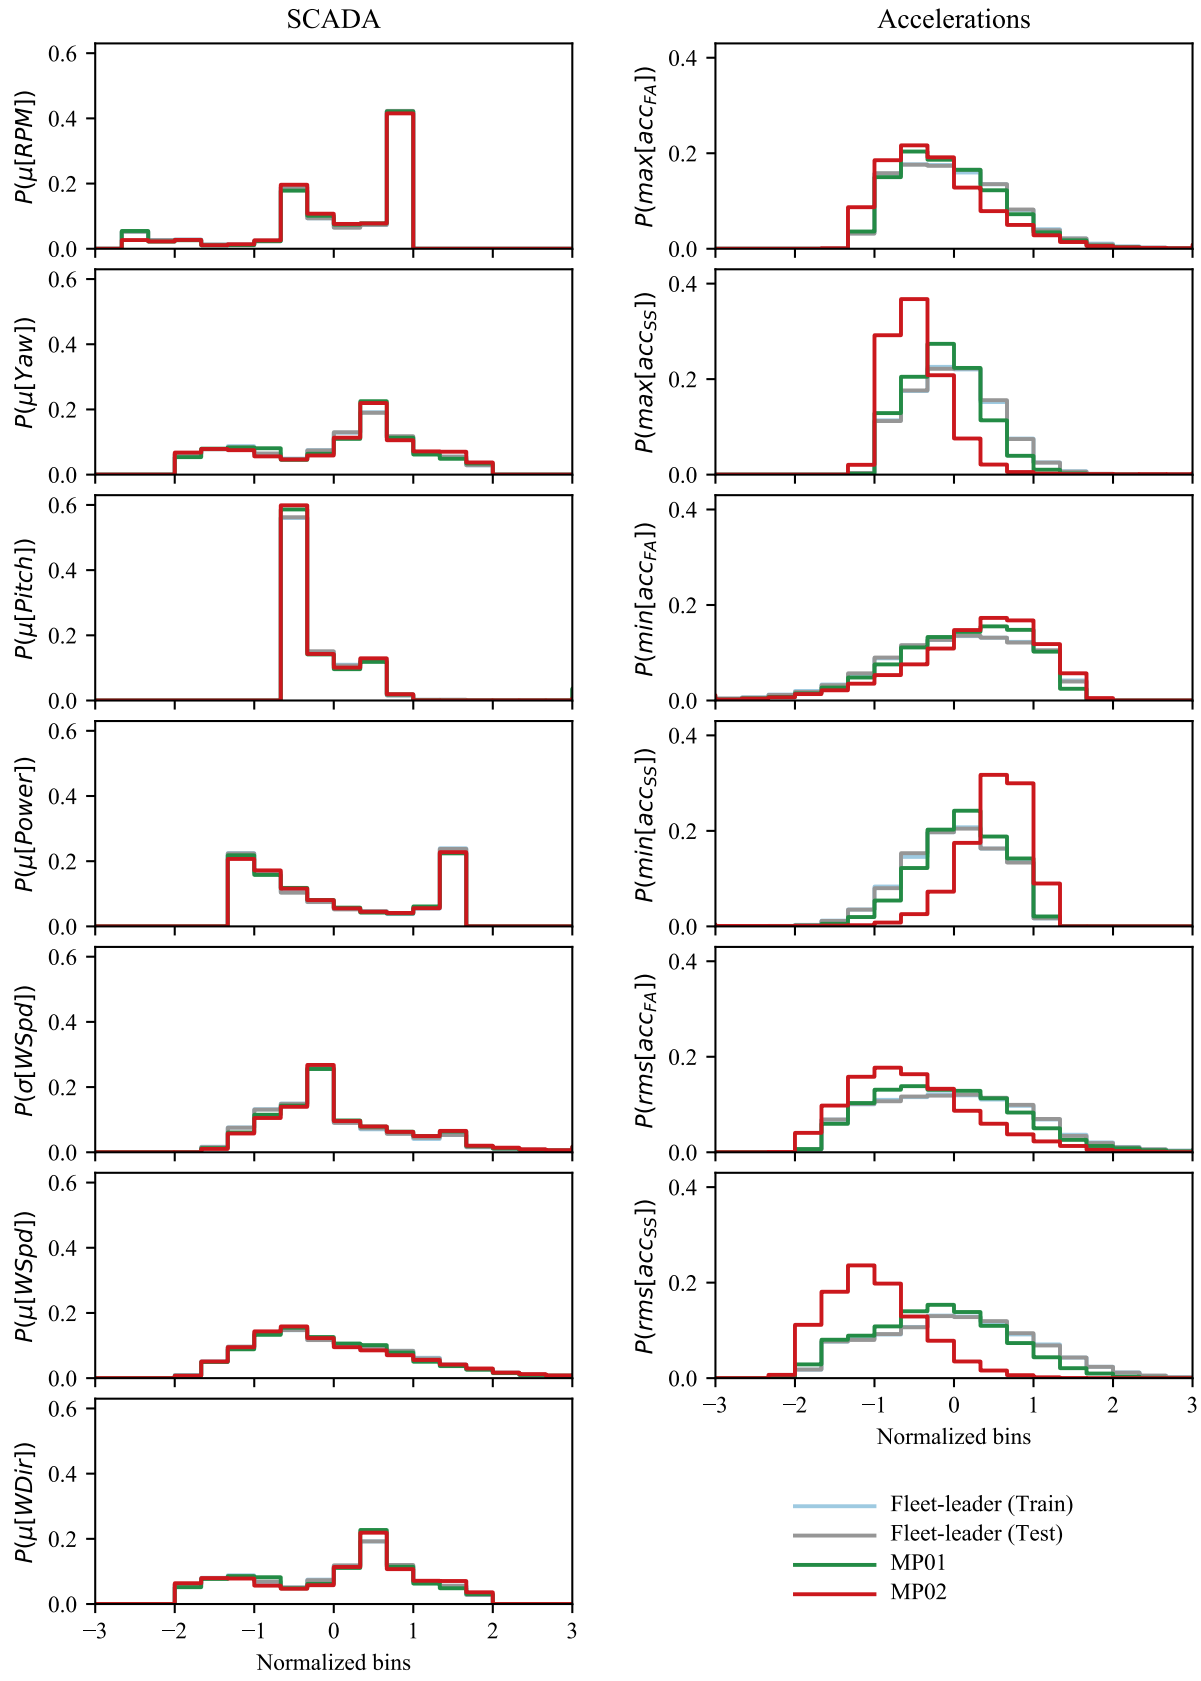

Fig. B1. Probability distributions of 10-minute SCADA and acceleration statistics for the fleet-leader, MP01 and MP02 turbines. The X-axes are normalized due to data confidentiality concerns.