



Delft University of Technology

Document Version

Final published version

Citation (APA)

Borsotti, M. (2026). *Towards Learning-Based Decision-Making for Maintenance of Offshore Wind Farms*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:186cd10d-f911-45ae-ba96-cfa1f542638d>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Towards Learning-Based Decision-Making for Maintenance of Offshore Wind Farms

Marco Borsotti



Towards Learning-Based Decision-Making for Maintenance of Offshore Wind Farms

Marco BORSOTTI

Towards Learning-Based Decision-Making for Maintenance of Offshore Wind Farms

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus, Prof. dr. ir. H. Bijl
chair of the Board for Doctorates
to be defended publicly on
Tuesday 6, October 2026 at 10:00

by

Marco BORSOTTI

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus
Prof. dr. R. R. Negenborn
Dr. X. Jiang

Chairperson
Delft University of Technology, promotor
Delft University of Technology, copromotor

Independent members:

Prof. dr. S. J. Watson
Prof. dr. M. T. J. Spaan
Prof. dr. T. Vanelander
Prof. dr. R. D. van der Mei
Dr. M. A. Mitici

Delft University of Technology
Delft University of Technology
University of Antwerp
Vrije Universiteit Amsterdam
Utrecht University



This research has been funded by NWO through the Holi-DOCTOR project (HOListic framework for DiagnOstiCs and moniTORing of Wind turbine blades), grant number: KICH1.ED02.20.004.

TRAIL Thesis Series no. T2026/26, The Netherlands Research School TRAIL

TRAIL
P.O. BOX 5017
2600 GA Delft
The Netherlands
E-mail: info@rsTRAIL.nl

ISBN: 978-90-5584-401-2

Copyright © 2026 by Marco BORSOTTI

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission of the author.

Printed in the Netherlands

In multa sapientia multa sit indignatio; et qui addit scientiam, addit et laborem.

Ecclesiastes 1:18 (Vulgate)

Preface

Shortly after graduating, I began working in the maintenance of renewable energy systems. I started with solar photovoltaic plants, and later moved into the operation and maintenance of onshore wind turbines. I still remember a conversation with my former boss, a good person and accomplished engineer. He once told me: “Once you have been here long enough, you will be able to recognise when a fault is about to occur just by looking at the turbine’s power curve.” And I, being only occasionally modest, thought: “Gotcha.” Four years later, after an entire PhD on maintenance planning, prognostics, uncertainty, and decision-making, I can safely say that I did not, in fact, “gotcha” at all.

That sentence stayed with me because it reflected something I saw repeatedly in practice: much of maintenance relies on knowledge that is difficult to formalise. Experienced operators and technicians develop an intuition for their assets over years of observation. They recognise subtle changes in behaviour, notice patterns that do not appear in manuals, and often have a surprisingly good sense of when something is not quite right. This in-house knowledge is incredibly valuable, but it is also difficult to transfer, scale, or systematically incorporate into decision-making processes. As I became more involved in maintenance planning and asset management, I realised that the challenge was not only to detect faults, but also to support the decisions that follow. In many cases, the answers are uncertain, and the consequences of being wrong can be expensive. In a way, this thesis is an attempt to increase the chances that the not-so-modest version of Marco could actually spot a failure coming and, more importantly, know what to do about it. The goal is not to replace experience, but to complement it, to make valuable operational knowledge more actionable, more transparent, and more effective in supporting timely maintenance decisions.

Around that time, I also knew that I wanted to work with offshore wind. I believed, and still believe, that offshore wind will be an essential part of our energy future, and I wanted to contribute to it. I applied for two positions that would have taken me in that direction: one at TU Delft and one in industry. Both selection processes reached their final stages, and for a while it seemed that I might have been able to choose between two very different ways of entering the offshore wind sector. A couple of weeks before the final decision, I had a serious car accident. I injured my spine and experienced short-term memory problems for several weeks. It was a strange and disorienting period, exactly when I was supposed to make an important decision about my future, I did not fully feel in control of my own ability to remember, organise, or evaluate things. During those weeks, I chose to start this PhD.

What followed was not always linear. There were models that worked only after failing many times, assumptions that had to be revised, papers that changed shape repeatedly, and research questions that became clearer only after trying to answer them. In this process, Rudy and Xiaoli played a fundamental role. They did not only supervise the research presented in this thesis; they helped shape the professional I became while writing it. I entered the PhD with a very practical

mindset, led by my inclination and experience in industry. My instinct was often to try something, see whether it worked, and then adjust. In many ways, this trial-and-error approach was useful. It kept me close to the problem, to implementation, and to the practical relevance of the work. Rudy and Xiaoli taught me that doing research means more than finding something that works. They pushed me to move from “this seems reasonable” to “this can be defended”. Their guidance helped me become more rigorous without losing the practical motivation that brought me to this topic in the first place. They gave me the freedom to develop my own ideas, while also teaching me how to turn those ideas into something that could be questioned, tested, improved, and communicated. For that, I am deeply grateful.

At the same time, I was fortunate to be surrounded by people who made the PhD experience richer and warmer. Among them, Donatella became an invaluable ally. Her kindness, passion, and support gave me more than practical help. She brought energy, curiosity, and generosity into the project, and many moments that could have felt difficult or frustrating became lighter because she was there. Our collaboration also grew beyond the boundaries of the project. We started side initiatives together, such as the LegOW game, with Ali and Brian, and those experiences reminded me that research can also be playful, creative, and genuinely fun. Donatella is not only someone I could rely on for support, but also a friendly face to turn to for guidance, encouragement, and perspective. She helped make this journey more interesting, more human, and, in many ways, more enjoyable. For that, too, I am deeply grateful.

More broadly, I am grateful to all the project members who contributed to making this research feel part of something larger: Daniele, Davide, Francesco, Helena, Mattijs, Muyao, Roger, Simon, Sumit, and Yanan. Having my work embedded within a broader project gave it perspective and a sense of belonging. In the context of this project, I have been entrusted with responsibilities that I enjoyed deeply and that helped me grow beyond the boundaries of my own research. One of the most meaningful examples was the organisation of two editions of the WindED summer school, for which we obtained a grant in 2025. Organising it was undoubtedly one of the highlights of these four years. The experience allowed me to bring together a group of speakers with exceptional knowledge and expertise, to expand my professional network, and to develop organisational skills that I had not previously had the chance to test in such a direct way. More importantly, it helped me satisfy my interest in dissemination. I was not only producing research or presenting my own work; I was helping shape what a room full of people would learn, discuss, and take away with them. For that, too, I am deeply grateful.

This experience also connects to another part of the PhD that became increasingly important to me: teaching. Thanks to Xiaoli and Donatella, who trusted me with lectures that gradually became longer, more complex, and more independent, I had the opportunity to discover how much I enjoy being in the classroom. Over time, I was given increasing freedom not only in how to teach, but also in what to discuss, how to structure the material, and how to engage students. Working with students has been one of the greatest sources of joy during these years, and one of the reasons I wanted to start a PhD in the first place. At the same time, teaching and supervision came with a non-negligible effort. I had to learn how to teach, how to explain things clearly, how to guide without taking over, and how to adapt to students with different backgrounds, motivations, and levels of confidence. This was especially meaningful to me because I have been, at different moments in my life, both a very good student and a struggling one. That made me particularly aware of how much the right support, the right explanation, or the right encouragement can change the way someone experiences learning. Whether I was leading groups through their projects, guiding a class through assignments,

or supervising students during their master's theses, I found the same underlying satisfaction in helping people move from an initial idea to something they could be proud of. Those moments were among the most gratifying experiences of my PhD.

At a time when the loneliness of a job mostly carried out behind the screen of a laptop had started to weigh on me, I found in Rafa a like-minded, bright colleague and friend. After hours spent discussing, researching, and developing ideas (with the help of a few beers), we built the embryo version of the model that would later become the final part of my research, and eventually published that work together. That collaboration meant a great deal to me. It reminded me how much I value working in a collaborative environment, challenging ideas together, improving, and sharing. It also reminded me of the joy of trusting someone enough to share part of the work with them, knowing that they will care about it with the same seriousness, curiosity, and commitment. Working with Rafa made that final stage of the PhD feel less solitary, and much more alive. For that, too, I am deeply grateful.

These four years have not only been my PhD. They have also been four years of life, with all the ups and downs that come with it. One of the most difficult parts has been being away from my home country, from my family, and from many of the people I love. I missed a lot. I missed birthdays, chill nights out, family dinners, small ordinary moments, and important occasions. I missed moments to celebrate, but also moments when I would have wanted to be close to people who may have needed me. That distance has not always been easy to accept. I know that I will have to work hard to recover some of the time that was lost, and I am eager to do so.

Mamma, Papà, Luca, Nonna, you supported me throughout this journey. You see me now as a grown man, and I hope I can make you proud.

Mamma, you gave me strength when I needed it. You came here to support me, to spend time with me, and to have our night-long conversations about everything and anything. With your beautiful smile, you have always given me hope. As you often say, "things always work out in the end, and better than you expected." I have carried that sentence with me more times than you know, and I have tried to share it with others.

Papà, I look up to you, there is so much that I have learnt from you and probably many lessons that I should have been more ready to receive when they were first offered. One of my earliest memories is of me "playing" piano on your legs. Since then, music has filled almost every moment of my life, and it has accompanied me through this PhD as well. Every time I sing or touch keys I feel blessed for having you in my life.

Luca, what we have is a bond of blood, but also something deeper than that. I wish I could have spent more time with you during these years, and I hope we will make up for it. I am so proud of the man you are becoming. You are funny, intelligent, and incredibly strong-willed. I admire you so much. It's funny how different we are, yet you are my brother, and I love you.

Nonna, ti amo. Grazie per tutto il supporto e per l'amore incondizionato che mi hai sempre dato. Le tue parole mi hanno accudito e protetto in più occasioni, anche da lontano. Se porto a termine questo dottorato, è anche per vedere l'orgoglio nei tuoi occhi.

In alphabetical order:

Ale, I cannot imagine growing up without you by my side. You grounded me, pushed me, challenged me, and helped me grow in more ways than I can count. Seeing how accomplished you are now makes me immensely proud. We have literally shed blood, sweat, and tears on those boats,

and now that you are old and weak, I suppose I am the one who has to row for both of us. That is fine, my friend. I will allow it.

Bea, with you I explored everything creative from the time we were young. Looking at you now, building a life out of creativity, makes me incredibly happy. I think back to our recordings, our videos, and all the strange things we did, and I feel lucky to have been even a tiny part of your journey. You are a Diva, and I will always be your friend and number one fan.

Budé, no laugh is more contagious than yours. I am grateful for all the time we spent together, for all the fun, and for the confidence you helped me build. I owe you more than you probably realise. It is crazy to think of this PhD as, in some way, an extension of what we started “studying” together back in 2013. Who. Would. Have. Thought? None of this would have been possible without you.

In different ways, the three of you taught me what I carried into this PhD: the effort required to pursue a goal, the creativity to improve my work and to find refuge from it, and the ability to keep smiling and finding fun along the way. With the three of you I do not have a bond of blood, but we made our own.

Mi Amor, this is really dedicated to you, to the life we want to build together. Since I met you, many things have slowly fallen into place. With you I have come to understand myself more clearly, the parts I like, the parts I still need to work on, and the person I want to become. A meaningful part of this PhD also took shape while I was with you in Colombia, in between days together and the final rounds of review of my last two papers. In that period, the thesis was becoming more real, but I was also beginning to see more clearly the life I want after it. Now I have someone with whom I get to share the doubts, the fears, the milestones and the successes. That has changed everything. This thesis is dedicated to you, to our future, and to the doors I hope it will open for us.

For this, I am many things now, I am an engineer, a researcher, a colleague, a friend, a partner, a brother and a son, and most of all, I am *lucky*.

Contents

Preface	vii
1 Introduction	1
1.1 Background and research motivation	1
1.2 Research questions and approach	7
1.3 Contributions	9
1.4 Thesis outline	9
2 Literature Review	11
2.1 Introduction	12
2.1.1 Maintenance strategies	12
2.1.2 Decision-making framework for OWF O&M	13
2.2 Review methodology	14
2.3 Strategic decision-making	16
2.3.1 Strategy evaluation	17
2.3.2 Policy design	22
2.4 Tactical decision-making	24
2.4.1 Fleet selection	24
2.4.2 Spare part management	25
2.5 Operational decision-making	28
2.5.1 Diagnostics and prognostics	28
2.5.2 Routing and scheduling	30
2.6 Gaps and limitations	33
2.6.1 Prognostic-driven scheduling	34
2.6.2 Considering the interconnections	35
2.6.3 Environmentally-inclusive decision-making	37
2.6.4 Lack of real-life data	38
2.7 Conclusions	39

3	OWF Characteristics & Dynamics	41
3.1	Requirements of the modeling environment	42
3.2	Wind farm configuration and spatial layout	45
3.2.1	Power production model	45
3.2.2	Modular layout generation	47
3.3	Critical components set	48
3.4	Transportation and access assets	52
3.4.1	Crew transfer vessel	53
3.4.2	Service operation vessel	54
3.4.3	Jack-up vessel	54
3.4.4	Helicopter	54
3.4.5	Spare parts management	55
3.5	Weather data and synthetic metocean dataset generation	55
3.5.1	Historical metocean data source	55
3.5.2	Vessel-specific operability states	56
3.5.3	Month-dependent Markov-chain estimation	56
3.5.4	Synthetic operability series generation	57
3.5.5	Validation of the synthetic operability dataset	58
3.6	Synthetic RUL prediction generator	59
3.6.1	Latent true RUL generation	62
3.6.2	Time-varying uncertainty model	63
3.6.3	Prediction model	66
3.7	Maintenance actions	68
3.7.1	Virtual-age formulation of maintenance effects	70
3.7.2	Effect on the predictive distribution	71
3.7.3	Action-specific resource requirements	72
3.8	Benchmark implementation	74
3.8.1	Modeling scope and abstractions	75
3.9	Conclusions	77
4	Optimization-Based O&M Planning	79
4.1	Introduction	80
4.2	Literature review	81
4.3	Methodology	84

4.3.1	RUL prognostics input to the MILP	84
4.3.2	Opportunistic maintenance	85
4.3.3	Maintenance model	86
4.3.4	Stochastic optimization formulation	89
4.3.5	MPC formulation	92
4.3.6	Diagnosis of rolling-horizon limitations	94
4.4	Case study	96
4.5	Results	99
4.6	Conclusions	103
5	From Optimization to Learning-Based O&M Planning	105
5.1	Introduction	106
5.2	Single-agent DRL approaches for offshore wind O&M	109
5.2.1	Deep Q-networks	111
5.2.2	Policy gradient methods	111
5.2.3	Actor–critic and other methods	112
5.3	Multi-agent DRL in maintenance planning	113
5.4	Problem formulation	115
5.4.1	Markov decision process	115
5.4.2	Partially observable Markov decision process	116
5.4.3	Graph and network-based representations	117
5.4.4	Hierarchical and interpretable models	117
5.4.5	Decision frameworks and objectives	118
5.5	Applications of DRL approaches for offshore wind O&M	119
5.5.1	Wind farm aerodynamics	119
5.5.2	Weather and sea state	120
5.5.3	Logistics	120
5.5.4	Prognostics and health management (PHM) data	120
5.5.5	Economic factors	121
5.6	Discussion	124
5.6.1	Cost and downtime reduction	124
5.6.2	Predictive (PHM-based) strategies	124
5.6.3	Opportunistic maintenance	125
5.6.4	Reliability and safety	125

5.6.5	Computational feasibility	125
5.6.6	Simulation-based studies vs. real-world applications	126
5.6.7	Multi-level repairs	127
5.7	Research gap	128
5.8	Conclusions	129
6	Learning-Based O&M Planning	133
6.1	Literature review	134
6.2	Methodology	136
6.3	Experimental setup	142
6.4	Results and comparative study	144
6.4.1	Training	144
6.4.2	Model comparison	144
6.5	Conclusions	151
7	Conclusions and Future Research	155
7.1	Answers to the research questions	156
7.2	Answer to main research question	159
7.3	Contributions	161
7.4	Practical implications	162
7.5	Limitations and future research	164
7.5.1	Short-term future research directions	165
7.5.2	Longer-term future research directions	167
7.6	Final remarks	167
	Bibliography	169
	Glossary	187
	Summary	193
	Samenvatting (Summary in Dutch)	195
	About the author	197
	TRAIL Thesis Series publications	199

Chapter 1

Introduction

This chapter introduces the main challenges motivating the development of new model-based decision support methods for offshore wind farms operations and maintenance (O&M). More specifically, Section 1.1 provides the research background, key definitions, and the main challenges in offshore wind O&M, motivating the problem by discussing why existing decision-support approaches remain limited when uncertainty, accessibility, and logistics must be handled jointly. The overarching research question and sub-questions are formulated in Section 1.2, together with the research approach adopted in this thesis. Section 1.3 summarises the main scientific contributions, and Section 1.4 concludes with the overall thesis structure, supported by a block-diagram representation.

1.1 Background and research motivation

The global energy system is undergoing an accelerated transition driven by climate objectives, energy-security considerations, and the rapidly improving competitiveness of renewable power technologies. In this context, wind energy, and offshore wind in particular, is increasingly positioned as a cornerstone of power-sector decarbonisation, especially for coastal countries with strong wind resources, limited land availability, and growing electricity demand due to electrification of transport, heating, and industry. International roadmaps and policy frameworks have further strengthened this trajectory. For instance, the global commitment to work toward tripling installed renewable capacity by 2030 has sharpened the focus on scalable, mature renewables, notably wind and solar (International Energy Agency, 2024b). In parallel, the International Energy Agency projects a substantial rise in the renewable share of electricity generation toward 2030, with wind and solar accounting for the vast majority of this growth (International Energy Agency, 2024a). Wind power generation has already expanded rapidly, exceeding 2,330 TWh in 2023, and continued acceleration is required to align with net-zero trajectories (International Energy Agency, 2024c). This trend is shown in Figure 1.1.

Offshore wind brings distinct advantages relative to other wind technologies. Compared to onshore wind, offshore sites typically feature higher and more stable wind speeds, lower visual and noise constraints, and access to very large contiguous areas that enable utility-scale developments. These advantages, however, come with increased engineering, installation, and operations complexity. Despite these challenges, evidence indicates that offshore wind is a globally expanding industry.




Figure 1.1: Global electricity generation by technology, 2000-2030 (International Energy Agency, 2024c)

Installed offshore wind capacity has surpassed 80 GW worldwide, with continued growth expected across Europe, Asia, and North America (Global Wind Energy Council, 2025), this can be seen in Figure 1.2.

At the same time, long-term expansion is strongly shaped by a large and evolving project pipeline, policy design, and macroeconomic headwinds (e.g., interest rates, supply-chain constraints), which have led to revisions of near-term deployment expectations in several markets (McCoy et al., 2024).

From a techno-economic standpoint, offshore wind has seen significant cost reductions over the last decade through larger turbines, improved installation practices, better wake-loss management, and more industrialised supply chains. Recent international cost assessments report a




Figure 1.2: New offshore wind installations, global, in MW (Global Wind Energy Council, 2025)

global weighted-average levelised cost of electricity (LCOE¹) for offshore wind of approximately USD 0.075/kWh for projects commissioned in 2023, reflecting a substantial improvement relative to earlier deployment eras (International Renewable Energy Agency, 2024), as shown in Figure 1.3.

At the same time, turbine upscaling also introduces new operational challenges. Figure 1.4 shows that turbine sizes are increasing, with 17 MW units predicted by 2025 (International Energy Agency, 2024b). As turbine ratings and rotor diameters increase, maintenance interventions become more demanding in terms of access, equipment, lifting requirements, and execution risk. Failures in larger turbines can also imply greater production losses per event, while major repairs may require more specialised vessels, tools, and planning effort. As a result, some of the benefits associated with larger machines are accompanied by increased O&M complexity, particularly for heavy maintenance tasks.

Nonetheless, costs remain highly heterogeneous and sensitive to site characteristics (water depth, distance to shore, soil conditions), technology choice (fixed-bottom versus floating), and the maturity of regional supply chains and ports. In Europe, strategic targets reflect both ambition and the scale of the infrastructure challenge: the European Commission has set capacity objectives of at least 60 GW offshore wind by 2030 and 300 GW by 2050 (European Commission, 2020). Similar ambitions exist in other leading markets, such as the United Kingdom's target of up to 50 GW offshore wind by 2030 (HM Government, 2022).




Figure 1.3: LCoE benchmark for energy generation technologies (Global Wind Energy Council, 2025)

¹The Levelized Cost of Energy (LCOE) represents the average cost of producing one unit of electricity over the lifetime of an energy asset. It is commonly expressed as the ratio between the discounted lifetime costs and the discounted lifetime electricity production:

$$\text{LCOE} = \frac{\sum_{t=0}^T \frac{I_t + O_t + M_t + F_t}{(1+r)^t}}{\sum_{t=0}^T \frac{E_t}{(1+r)^t}},$$

where I_t denotes investment expenditures, O_t operating costs, M_t maintenance costs, F_t fuel costs, E_t electricity generated in period t , r is the discount rate, and T is the project lifetime. For offshore wind, $F_t = 0$, so LCOE is mainly driven by capital expenditure, O&M costs, financing assumptions, and energy yield (Darling et al., 2011).




Figure 1.4: Rotor size and power rating increase International Energy Agency (2024b)

As the sector scales, the operational phase becomes increasingly decisive for both profitability and system-level contribution. Offshore wind farms are capital-intensive assets with long lifetimes (typically 25-30 years), and their economic performance is strongly influenced by availability and downtime. Operation and maintenance (O&M) is therefore not a peripheral activity but a central driver of lifecycle cost and energy yield. In multiple studies and industry assessments, O&M is estimated to contribute on the order of 25-30% of the overall LCOE, depending on site conditions and operational strategy (Katsouris & Savenije, 2016). Figure 1.5 provides a high-level illustration of how total expenditure is distributed across the major lifecycle cost categories of a representative fixed-bottom offshore wind farm. The breakdown highlights that, beyond the up-front capital items (wind turbine supply, balance of plant, and installation/commissioning), a substantial share of total spend accrues during the operational phase through operation, maintenance and service (OMS). The values shown in Figure 1.5 follow the cost-category split reported for a 500 MW reference project (6 MW turbines, jacket foundations) and are intended as an indicative benchmark rather than a universal allocation, since the relative shares can shift with site conditions, supply-chain maturity, and O&M strategy (BVG Associates, 2014).

Figure 1.6 further clarifies annual operating expenditure (OpEx) to highlight which cost drivers dominate during the operational phase. The breakdown indicates that maintenance-related spending is the primary contributor, and that a large fraction of maintenance cost is associated with the chartering and use of vessels and specialised equipment, rather than labour or consumable materials alone. This reinforces a core point for offshore O&M planning: logistics is a major cost lever, and decisions that reduce waiting time, improve the utilisation of marine resources, or avoid avoidable mobilisations can have first-order economic impact. The values in Figure 1.6 are based on the representative fixed-bottom offshore wind reference project cost components reported in Stehly et al. (2024).

These costs can increase for harsh environments, far-from-shore sites, and for next-generation installations in deeper waters where floating foundations may become necessary. In these contexts, logistics and accessibility constraints frequently dominate effective downtime.




Figure 1.5: Lifecycle costs for offshore wind, adapted from: Stehly et al. (2024)




Figure 1.6: Main contributors to offshore wind operational cost, adapted from: Stehly et al. (2024)

In offshore wind O&M, maintenance is executed in a spatio-temporally constrained, highly uncertain environment. Key operational decisions, such as when to dispatch crews, which vessels to charter, what spares to pre-position, and which turbines to prioritise, must be made under strong coupling between (i) metocean conditions, (ii) vessel capabilities and port infrastructure, (iii) turbine/component health states, and (iv) the opportunity cost of lost production. Accessibility is often governed by thresholds on wind speed, wave height, and visibility, which shape feasible weather windows and introduce significant stochasticity into execution times. At the same time, the offshore supply chain relies on scarce and expensive resources such as service operation vessels (SOVs), crew transfer vessels (CTVs), heavy lift vessels, helicopters, and specialised crews. These resources

are typically shared across projects and subject to market volatility, contracting constraints, and limited availability. Consequently, maintenance planning becomes a multi-constraint optimization problem in which technical, logistical, and meteorological uncertainty are necessary considerations (McCoy et al., 2024).

Within this setting, the industry has historically relied on combinations of corrective maintenance, time-based preventive maintenance, and condition-based maintenance supported by SCADA systems and condition monitoring. While condition monitoring and predictive maintenance have advanced substantially, the translation of health information into dispatch and scheduling decisions remains non-trivial. Health forecasts (e.g., remaining useful life (RUL)) are uncertain, maintenance actions are heterogeneous (inspection, minor repair, major replacement), and the value of acting early must be weighed against mobilisation costs and the risk of failed access attempts due to weather. Moreover, offshore wind farms are comprised of many turbines and multiple subsystems per turbine, creating opportunities for grouping and opportunistic maintenance, but also creating combinatorial complexity when such opportunities are evaluated jointly with vessel routing and weather uncertainty, as shown in Chapter 2.

These realities motivate the need for model-based decision support across multiple planning levels. At a strategic level, decisions concern long-term contracting, fleet sizing, spare parts policies, and infrastructure investments (ports, warehousing, and digital capabilities). At a tactical level, decision makers must translate seasonal forecasts and expected failure patterns into maintenance campaigns and vessel assignments. At an operational level, planners must update weekly or day-to-day schedules in response to realised weather, failures, and resource availability. The defining difficulty is that offshore wind O&M is simultaneously (a) dynamic, (b) uncertain, and (c) operationally coupled across assets and resources. This has two important implications for research: first, solution methods must handle large state spaces and frequent re-optimization; second, uncertainty must be represented in a way that supports robust decisions without prohibitive computational burden (McCoy et al., 2024).

Optimization-based approaches (including mixed-integer linear programming, heuristics, and decomposition methods) have delivered valuable insights and strong benchmarks for tactical planning under explicit constraints. However, the increasing scale of wind farms, the stochasticity of access and failures, and the desire to leverage high-frequency monitoring data have pushed the community toward sequential decision-making formulations. Data-driven and learning-based approaches, including deep reinforcement learning (DRL), have emerged as promising candidates to support adaptive operational policies that react to evolving system states and uncertainty while learning effective strategies from simulation. At the same time, industry and policy developments continue to expand the scope of relevant constraints: grid integration, curtailment risk, and broader sustainability considerations can all influence the value of production and the prioritisation of interventions (International Energy Agency, 2024a; WindEurope, 2024). Consequently, new directions in offshore wind O&M research increasingly emphasise integrated frameworks that combine metocean-driven accessibility modeling, realistic logistics, cost and downtime modeling, and health-informed decision-making. These directions form the foundation for the methodological choices and comparative analysis developed throughout this thesis.

The motivation of this thesis is thus rooted in a practical gap: offshore wind O&M planning must reconcile three realities simultaneously: (1) uncertain and evolving component health information, (2) strongly weather-constrained accessibility, and (3) logistics constraints and resource coupling across turbines.

1.2 Research questions and approach

The overall research question of this thesis is:

Q: *What are the trade-offs between optimization-based and learning-based decision-making for offshore wind maintenance?*

To address this main research question, the following sub-questions are formulated.

Q1: *What are the state-of-the-art, state-of-practice, and research gaps in offshore wind O&M decision support, particularly regarding the integration of accessibility, logistics, and condition-informed planning?*

As outlined in Section 1.1, offshore wind O&M is governed by a strong coupling between access constraints, scarce marine resources, and uncertain degradation and failure processes. Although the literature includes extensive work on maintenance strategies, logistics, and weather-driven accessibility modeling, these aspects are often treated in partial or disconnected ways. For example, some approaches focus on maintenance timing while simplifying logistics to generic capacity constraints, whereas others model vessels and routing but represent asset condition through static thresholds or aggregated failure rates. In practice, operators must combine these dimensions continuously: condition information affects which assets to prioritise, accessibility determines whether an intervention is feasible, and limited resources impose trade-offs across turbines. A structured synthesis is therefore required to map dominant modeling approaches and assumptions, identify recurring simplifications, and clarify the resulting research gaps that motivate the methodological choices of this thesis.

Q2: *How to model the dynamics of offshore wind farm O&M to support maintenance decision-making?*

A persistent challenge in method development is that results are difficult to compare across studies because modeling assumptions differ substantially. Differences in metocean datasets and access thresholds, vessel and technician availability, transit times, spare-part policies, and cost components can dominate performance metrics such as downtime and total O&M cost. As a result, apparent improvements attributed to a method may instead reflect differences in the case study or environment. To address this, the thesis defines a shared modeling framework comprising: (i) metocean-driven accessibility modeling and weather-window logic, (ii) a representation of offshore logistics resources and operational constraints, (iii) a cost structure that captures both direct maintenance costs and downtime-related losses, and (iv) a reference wind farm case study used consistently throughout the thesis. This shared environment provides a basis for controlled evaluation and aligns with the objective of quantifying the operational and economic value of a maintenance strategy.

Q3: *How can O&M planning be formulated as an optimization model that incorporates condition and prognostic information?*

Optimization-based planning is widely used in offshore wind because it can represent feasibility constraints explicitly and yield interpretable schedules. However, many planning tools still rely on simplified maintenance triggers (e.g., time-based rules) or separate the decision of *what* to maintain from the decision of *when* and *how* to execute maintenance offshore. This thesis therefore develops an optimization-based method for tactical O&M planning that integrates condition/prognostic

information into maintenance selection and timing. The model is designed to capture the coupling between turbine-level decisions and shared logistics resources, represent access limitations and weather-driven feasibility, and quantify trade-offs among preventive interventions, corrective repairs, and downtime losses. The resulting solutions provide insight into the value and limitations of optimization-based decision support.

Q4: *Which modeling requirements should a learning-based offshore wind O&M framework satisfy to address the limitations identified in the optimization-based framework?*

While optimization models can provide high-quality schedules under well-defined assumptions, they are often challenged by the uncertainty and scale that characterise offshore wind. Metocean conditions vary stochastically, failures occur unpredictably, repair durations are uncertain, and condition information is imperfect and continuously updated. Representing these uncertainties explicitly in optimization formulations typically increases problem size substantially (e.g., via scenarios, recourse decisions, or robustness sets) and can make repeated replanning burdensome. This thesis analyses these limitations with respect to scalability (number of turbines, time resolution, resource detail), tractability of uncertainty representation, and the ability to adapt decisions online as new information arrives. The outcome is a clear motivation for investigating methods that can learn adaptive policies for repeated decision-making under uncertainty, while still respecting the operational constraints that dominate offshore O&M.

Q5: *How can offshore wind O&M decision-making be formulated and solved as a learning problem to derive adaptive maintenance policies under offshore constraints?*

Reinforcement learning provides a framework for learning decision policies in stochastic, sequential environments. For offshore wind O&M, the decision maker repeatedly observes an operational situation (asset condition indicators, resource status, and accessibility information) and selects actions (e.g., whether and what to maintain, and which intervention type), with the objective of improving long-run performance in terms of cost and downtime. This thesis develops a learning-based framework aligned with the shared modeling environment defined in Q2, with three requirements: actions must map to operationally meaningful interventions; the environment must reflect weather-driven access constraints and logistics limitations; and the performance signal (reward) must represent the relevant trade-offs between preventive actions, corrective failures, and lost production. Because asset health is inferred from monitoring and prognostics and is therefore uncertain, the learning formulation explicitly accounts for partial observability by using condition and uncertainty features as inputs to the policy.

Optimization-based approaches can represent operational constraints explicitly and provide transparent schedules, but they often require simplifying assumptions to remain solvable at scale and may struggle to adapt decisions online as new information arrives (e.g., updated RUL distributions or revised weather forecasts).

In parallel, Deep Reinforcement Learning (DRL) is increasingly explored for sequential decision-making under uncertainty and high-dimensional state information. In principle, DRL can learn adaptive policies that condition actions on evolving system states, but credible application to offshore wind requires operational feasibility, realistic cost and downtime signals, and careful treatment of partial observability in health and forecasts.

This thesis positions MILP and DRL as complementary paradigms for offshore wind O&M decision support and develops them under a shared modeling environment (metocean, logistics, cost model, and a consistent case study) to enable controlled comparison.

1.3 Contributions

The contributions of this thesis are as follows:

1. A consolidated review of offshore wind O&M decision-support modeling, covering maintenance strategies, logistics constraints, metocean/accessibility impacts, cost drivers, and research gaps. This work has been published as (part of) Borsotti et al. (2026c).
2. A modeling framework (metocean, logistics, and cost) and a reference case study enabling controlled comparison between optimization-based and learning-based approaches.
3. A MILP-based O&M planning model that integrates prognostic information into planning decisions, together with a performance and insights analysis on the reference case study. This work has been published as (part of) Borsotti et al. (2026b).
4. A structured diagnosis of MILP limitations (uncertainty handling, scalability, adaptivity) and a synthesis of DRL methods for maintenance decision-making, including implications of partial observability. This work has been published as (part of) Borsotti et al. (2026a).
5. A DRL framework for offshore O&M planning formulated as a POMDP, incorporating probabilistic RUL features and multi-level maintenance actions, trained and evaluated under the shared case study. This work has been published as (part of) Borsotti et al. (2025).
6. A comparative perspective on the advantages and disadvantages of optimization- and learning-based approaches to maintenance scheduling in offshore wind O&M, grounded in empirical results and practical requirements (feasibility, interpretability, scalability, robustness).

1.4 Thesis outline

This thesis is structured in seven chapters. A graphical representation of the structure and the relationship between the shared modeling framework, the optimization-based workstream, and the learning-based workstream is provided in Figure 1.7. The contents of each chapter are briefly described next:

- **Chapter 2** answers Q1 by reviewing the state of the art in offshore wind farm O&M modeling and decision support, covering maintenance strategies, metocean impacts, logistical constraints, and cost drivers, and identifying key research gaps.
- **Chapter 3** answers Q2, presenting the shared modeling environment used throughout the thesis, including metocean/accessibility assumptions, vessel and logistics modeling, prognostic information and the cost model, together with the reference case study.

- **Chapter 4** addresses Q3 and presents the MILP formulation for tactical O&M planning, solution strategy, and case-study results, including performance analysis and operational insights.
- **Chapter 5** answers Q4 and discusses the limitations of the MILP paradigm for offshore O&M and synthesises the DRL literature relevant to maintenance decision-making, motivating the transition toward POMDP-based DRL methods.
- **Chapter 6** addresses Q5 and presents the DRL environment and POMDP formulation, algorithmic choices, training setup, and results, including comparisons against benchmark maintenance strategies.
- **Chapter 7** concludes the thesis by answering the main research question, consolidating findings across the optimization-based and learning-based contributions, discussing methodological implications and practical deployment considerations, and outlining directions for future research.




Figure 1.7: Thesis structure and methodological flow

Chapter 2

Literature Review

Offshore wind is scaling rapidly, yet its competitiveness remains strongly constrained by the cost and complexity of Operation and Maintenance (O&M). The objective of this chapter is to consolidate the current state of knowledge on how O&M decisions are modelled, optimized, and supported in practice, and to clarify where the literature still falls short. In doing so, the chapter directly addresses Research Question Q1: *What are the state-of-the-art, state-of-practice, and research gaps in offshore wind O&M decision support, particularly regarding the integration of accessibility, logistics, and condition-informed planning?*

To answer Q1, the chapter structures the literature using a strategic-tactical-operational decision-making framework, and maps modeling contributions to the decisions they support, the assumptions they rely on, and the uncertainties they represent.

The chapter is organized as follows. Section 2.1 outlines the main maintenance strategies, and introduces the decision-making framework used throughout the review. Section 2.2 details the review protocol, including database selection, filtering stages, and the classification scheme adopted for synthesis. Section 2.3 surveys *strategic* models, distinguishing between simulation-based strategy evaluation and optimization-based policy design, with particular emphasis on age-threshold approaches. Section 2.4 reviews *tactical* decision models for fleet selection and spare parts management. Section 2.5 addresses *operational* decision-making, focusing on diagnostic and prognostic methods together with routing and scheduling optimization under weather and resource constraints. Section 2.6 synthesizes the main gaps and limitations, with particular attention to prognostic-driven scheduling, cross-layer interconnections, environmentally-inclusive O&M planning, and the persistent scarcity of real operational data. Finally, Section 2.7 concludes the chapter by consolidating the implications of the review for the remainder of the thesis and translating the identified gaps into concrete research directions.

This chapter has been published as M. Borsotti, X. Jiang, R. R. Negenborn: “A review of multi-horizon decision-making for operation and maintenance of fixed-bottom offshore wind farms”, *Renewable and Sustainable Energy Reviews*, vol. 226, article 116450, 2026.

2.1 Introduction

Offshore wind has expanded rapidly in recent years and is expected to play an increasingly important role in the energy transition. However, despite its strong potential, offshore wind remains characterized by high capital and operating costs compared with more mature renewable technologies. In this context, O&M plays a central role, since it directly affects cost, downtime, and energy production. As outlined in Chapter 1, offshore O&M is particularly challenging because of limited accessibility, complex logistics, and uncertainty in component degradation. Therefore, efficient O&M is crucial for reducing offshore wind LCoE. Optimized strategies can cut maintenance frequency and downtime, increasing turbine availability and performance.

2.1.1 Maintenance strategies

In this section, we will provide a concise overview of the main maintenance strategies utilized in the O&M of fixed-bottom OWFs. These strategies include corrective maintenance, preventive maintenance, predictive maintenance, and opportunistic maintenance Ren et al. (2021). These strategies differ in objectives, features, benefits and shortcomings:

- **Corrective maintenance:** repairs or replaces components post-failure. Although reactive and potentially costly due to extended downtime, it effectively addresses unexpected failures Rinaldi et al. (2021).
- **Preventive maintenance:** schedules routine tasks (inspections, lubrication, cleaning, replacements) at fixed intervals to prevent failures and sustain performance, though it may lead to unnecessary maintenance and higher costs Fox et al. (2022).
- **Condition-based maintenance (CBM):** continuously monitors component conditions via sensors to detect deterioration, allowing maintenance to be performed only when needed.
- **Predictive maintenance:** uses sensor data and analytics to forecast failures and optimize scheduling Kou et al. (2022). It reduces unexpected failures by identifying trends and anomalies Le & Andrews (2016); May et al. (2015) but requires a high initial investment in monitoring systems Fox et al. (2022).
- **Prescriptive maintenance:** extends beyond predictive analytics by not only forecasting equipment failures but also prescribing specific maintenance actions to mitigate or completely avoid the issues Fox et al. (2022).
- **Opportunistic maintenance:** uses unplanned opportunities to group tasks, thereby optimising resources, reducing downtime, and enhancing reliability Rinaldi et al. (2021).

The main goal of asset management is not to simply select one maintenance strategy over another, but rather to understand the optimal timing for employing different strategies based on certain age thresholds in the lifetime of components.

Figure 2.1 illustrates that as a component ages, there are specific points or thresholds where it is most effective to switch from one maintenance strategy to another.




Figure 2.1: Maintenance strategies, age thresholds and state of health of a component (adapted from Li et al. (2021))

To determine these age thresholds and make informed decisions, it is crucial to have a high level of confidence in knowing the state of health of the components. This is where health prognostic and diagnostic studies play a vital role. By utilizing condition monitoring data, advanced analytics, and predictive modeling techniques, health prognostic studies enable accurate assessments of the health and RUL of components. This information enhances the effectiveness of maintenance strategies by allowing for timely interventions and proactive measures.

2.1.2 Decision-making framework for OWF O&M

In this review, we adopt a widely accepted decision-making classification that distinguishes between strategic, tactical, and operational levels (Shafiee, 2015). This structure is commonly used in reliability engineering and operations management to differentiate decisions by scope, time scale, and level of detail. Strategic decisions focus on long-term planning (typically 10–25 years), such as selecting the O&M strategy, including activity types, frequency, and budgeting. Tactical decisions cover mid-term planning (months to several years), including spare part logistics, vessel fleet selection, and maintenance resource allocation. Operational decisions address short-term execution (days to weeks), such as scheduling of maintenance tasks, routing of vessels, and response to failures.

The main scope of the three decision-making layers is summarized in Figure 2.2, together with the most common input data and the typical outputs expected for their optimization; the flowchart illustrates the main information flows and dependencies between the strategic, tactical, and operational layers. A more detailed breakdown of the specific inputs, outputs, and constraints associated with each decision-making level is provided later on in Table 2.10 from Section 2.6.2.

Overall, the state of the art in OWFs O&M is constantly evolving, with new technologies and approaches being developed and implemented all the time. In this chapter we review the methodologies and approaches used to optimize the decision-making for O&M over these three layers and will point out the limitations of the existing models and the gaps in the literature. While




Figure 2.2: Decision-making layers, scope, inputs and outputs

prior reviews have provided valuable insights into specific aspects or decision-making layers, our work seeks to bridge these segmented understandings. Moreover, this review stands out by providing a detailed description of the tools and methodologies that are used in different decision-making layers while providing a holistic critique of existing models. Throughout this review, we highlight how various modeling approaches support the development of decision support systems (DSS) for OWF O&M, which play a central role in assisting stakeholders across strategic, tactical, and operational levels.

2.2 Review methodology

This review aims to provide a structured and critical assessment of the literature on OWF O&M, with a particular focus on optimization and decision support methods. The scope includes fixed-bottom OWFs unless stated otherwise. The methodology follows established systematic review principles, adapted to address the specific heterogeneity and multi-horizon nature of O&M models.

The initial literature pool was built through searches in Web of Science, Scopus, ScienceDirect, and Google Scholar. We selected keywords aligned with the main decision layers and modeling approaches relevant to OWF O&M. These include technical, operational, and optimization-focused terms. Table 2.1 presents the keyword set used, always in combination with “offshore” and “offshore wind” to ensure domain specificity. The keyword list was refined iteratively based on preliminary search results and bibliometric keyword co-occurrence analysis.

The initial search yielded 1645 documents. These were filtered in three stages: articles were retained if they contained clear relevance to offshore wind O&M optimization, decision-making, simulation, or planning. Documents were then excluded if they addressed only general wind energy issues without a clear O&M focus or dealt solely with onshore systems. Finally, the remaining set was assessed based on specificity (e.g., modeling depth), recency, citation impact, and diversity of methods or horizons considered.

After this multi-stage filtering, a final curated set of 100 papers was selected to ensure coverage across decision-making layers, methodological diversity, and citation impact. This set formed the basis of the critical review.

To structure the analysis, the selected documents were categorized along two axes: depending on the decision-making layer: strategic, tactical, and operational. Respectively, 57%, 13%, and 30% of articles fell into these categories. Based instead on the topic of focus: this included O&M optimization, O&M simulation, routing and scheduling, fleet selection and spare part management (see Figure 2.3).

To map research trends, a bibliometric analysis was performed using the Bibliometrix package in R Aria & Cuccurullo (2017). Figures 2.4 to 2.6 summarize the most cited sources, and documents, along with the annual scientific production in this domain.

Table 2.1: Search themes and keywords used in the literature search

Search theme	Keywords
Maintenance strategy and performance	Operation and maintenance; opportunistic maintenance; reliability; availability; operation and maintenance simulation
Maintenance approaches	Condition-based monitoring; operation and maintenance optimization; life cycle assessment
Logistics and resources	Spare parts management; vessel routing; fleet selection
Maintenance execution	Maintenance scheduling




Figure 2.3: Clustering by focus topic.




Figure 2.4: Most relevant sources according to the bibliometric analysis.




Figure 2.5: Most relevant documents according to citation count.




Figure 2.6: Annual scientific production over time related to offshore wind O&M.

Although this review does not claim to be exhaustive, we have aimed for breadth and representativeness by combining quantitative relevance (e.g., citations, coverage) with qualitative selection (e.g., methodological depth, diversity, and recency). The next sections are structured around the three decision-making layers and provide a critical analysis of the most relevant modeling contributions in each.

2.3 Strategic decision-making

Strategic decision-making in OWF O&M involves high-level planning activities that shape the long-term performance and cost-effectiveness of assets. These decisions are typically made prior to or early in the operational phase and span the entire lifecycle of the wind farm. These decisions help define overarching maintenance strategies, guide long-term budgeting, and determine the key factors that will drive the initiation of maintenance cycles over time.

Given the scope and complexity of strategic O&M planning, this review distinguishes between two complementary modeling approaches: simulation models and optimization models. Simulation models are primarily used to evaluate the performance of predefined strategies under uncertainty and help explore the impact of logistical constraints, metocean conditions, and failure dynamics over time. Optimization models, in contrast, aim to determine the most cost-effective strategy by solving a mathematical formulation of the decision problem, often under simplifying assumptions. We separate these two families of models in our review to better reflect their distinct purposes,

data requirements, and methodological approaches. The following subsections present simulation models used for strategy evaluation and optimization models used for long-term policy design, respectively.

2.3.1 Strategy evaluation

This section reviews O&M models by organizing the literature around key strategic planning objectives. We begin with scenario-based simulation tools used to evaluate cost, logistics, and reliability over the project lifecycle. Next, we cover analytical models designed for early-stage planning with reduced computational burden. We then review state-based formalisms that model degradation and maintenance dynamics. This is followed by a discussion on how financial impacts and uncertainty are handled. Finally, we examine how metocean and accessibility modeling supports maintainability-aware site planning.

At the strategic level, it is important to assess the cost-effectiveness and technical feasibility of OWFs over their full lifecycle. As offshore turbines move further from shore, with increasingly complex configurations and harsh metocean conditions, accurate O&M modeling becomes central to reducing uncertainties in financial planning and strategic design decisions. A growing body of research has responded to this need with increasingly sophisticated simulation frameworks, capable of capturing not only turbine behavior and failure dynamics but also logistical constraints, weather-induced access limitations, and component-specific maintenance policies. This section reviews the most prominent O&M simulation models in the literature and highlights methodological innovations, data requirements, and the treatment of uncertainty.

A lot of effort has been put towards the development of such models; Rinaldi et al. (2021) provide a comprehensive review of current O&M modeling practices for offshore wind, including traditional and emerging maintenance paradigms.

By simulating different operational scenarios, simulation tools can be used to predict the expected maintenance needs of a wind farm over the long term, and to identify potential bottlenecks and risks in the maintenance process. This approach involves modeling the behavior of the wind turbines, the impact of different types of maintenance activities, and the logistics of personnel and equipment required for maintenance.

In recent research on O&M simulation tools for OWFs, several distinct approaches have been proposed. For the purpose of this review, we have focused on presenting the primary and extensively referenced models. These tools aim to evaluate the costs, risks, and operational strategies associated with maintenance activities. The four main ones are here summarized and have been extensively compared and analyzed in Dinwoodie et al. (2015):

UiS (Endrerud et al., 2014): adopts a multi-method simulation approach, combining agent-based and discrete event paradigms. The agent-based layer is used to model the individual attributes and behaviors of wind turbines, technicians, vessels, and maintenance managers, enabling decision-making at the agent level. In parallel, the discrete event paradigm models the time-driven execution of tasks within each agent. Specifically, it structures activities like travel, transfer, and repair as time-stamped events in a queue, processed sequentially based on resource availability and operational logic. This hybrid structure enhances the model's ability to capture the complexity and variability inherent in offshore O&M marine logistics.

ECUME (Douard et al., 2012): is a comprehensive O&M simulation tool developed by EDF that assesses the average cost of operation for OWF projects. The tool is based on event modeling through Monte Carlo simulation and uses Hidden Markov Models for inaccessibility risk assessment. It considers both deterministic and probabilistic cash flows. Deterministic cash flows being capital costs and operational costs, while probabilistic cash flows account for corrective maintenance costs due to failures and condition-based maintenance costs.

NOWIcob (Hofmann & Sperstad, 2013, 2014): uses a time-sequential Monte Carlo technique that incorporates weather uncertainties, different maintenance tasks and their associated costs. A key advantage of the model is the ability to represent various vessel concepts for wind farm access, as well as different maintenance regimes.

StrathOW (Dinwoodie et al., 2013; Dalgic et al., 2014): combines a Multivariate Auto Regressive climate model with a Markov-Chain Monte Carlo approach. The strength of this tool lies in capturing the impact of climate, wind farm specifications, and operational strategies on O&M costs while maintaining computational simplicity. The outputs from this model, such as vessel costs and lost revenue, are used as input for decision support analysis through Bayesian Belief Networks and decision trees.

These tools collectively contribute to evaluating O&M strategies and improving the efficiency and profitability of OWF operations and are used in other works such as Judge et al. (2019) where a comprehensive financial analysis model for OWFs is introduced, featuring stochastic time-series simulation modules that can be used to analyse technologies, strategies, and procedures throughout the wind farm's lifecycle. The model supports both fixed and floating turbines, multiple policies, and sites, while its O&M module calculates OPEX, energy production, and other critical performance indicators using NOWIcob.

In order to simulate the O&M activities of a wind farm over the long term, a discrete-event simulation approach is generally adopted. Figure 2.7 presents a generalized step-by-step process for the O&M simulation of wind farms over the long term.

A key challenge of simulating the O&M of OWFs is the complexity of the models and the amount of data required to accurately simulate the behavior of an OWF. To address this issue, recent research has introduced more computationally efficient and analytically tractable models for long-term O&M decision-making. For instance, Centeno-Telleria et al. (2023) propose a Markov chain-based analytical model that captures the key operational states of offshore renewable systems, enabling rapid long-term planning and extensive sensitivity analyses. Compared to conventional




Figure 2.7: Flowchart of common O&M simulation models

simulation models, their approach reduces computational burden while retaining accuracy in capturing downtime patterns. This method is further applied by the same authors (Centeno-Telleria et al., 2024b) to compare floating wind farms across the North Sea and Iberian Peninsula, linking maintenance-induced downtime with energy production to support strategic site-level planning. In a subsequent study (Centeno-Telleria et al., 2025), they assess heavy maintenance strategies, onsite versus towing, for floating wind turbines using a hybrid Markov framework, offering valuable insights for early-stage logistical and reliability planning. These analytical approaches demonstrate that Markov-based models can deliver strategic insights comparable to Monte Carlo simulations, but with far lower computational requirements, making them well-suited for use in DSSs at the design or early operational stage.

In addition, several other notable contributions further extend the study and development of O&M simulation tools for OWFs. Three state-based formalisms are commonly used to model reliability and its financial consequences: Markov chains, Semi-Markov or Markov-reward processes and Petri Nets. Markov chains represent memory-less transitions between states with time-independent or time-dependent rates. They are often used to simulate turbine degradation and availability, as in Stock-Williams & Swamy (2019) and Pattison et al. (2016). Semi-Markov or Markov-reward processes combine probabilistic transitions with cost or reward structures assigned to each state. This is applied for example in Stock-Williams & Swamy (2019). Finally, Petri Nets, such as those employed by Elusakin et al. (2021), provide a graphical framework for modeling concurrent activities, resource contention, and maintenance workflows, including crew and vessel constraints. For example, a reliability model for OWTs that includes degradation, inspection, and maintenance processes is introduced in Le & Andrews (2016). This model employs a Petri net approach to capture the stochastic dynamics of these processes, integrating both degradation and maintenance operations, and yielding key results such as state probabilities, expected task counts, and the downtime impact of various maintenance strategies. A stochastic Petri network (SPN) model for O&M scheduling of floating turbines (FOWTs) and their supporting structures is presented in Elusakin et al. (2021). This model accounts for various factors, including component deterioration and renewal processes, and utilizes data from literature and wind energy industry databases. It was tested on a 5 MW wind turbine on a spar buoy platform.

The reviewed articles offer insights into asset modeling, integrated maintenance policies, financial analysis, reliability and maintenance management. Table 2.2 presents a typology of decision-making, failure, and metocean modeling approaches applied in O&M simulation models. This overview is intended for researchers building simulation frameworks and looking to identify which model combinations are commonly used or underexplored. It helps in selecting compatible methods for new modeling studies or comparative benchmarking.

Table 2.3 compares O&M simulation models based on how they address uncertainty and the types of maintenance strategies implemented. It is intended for researchers and practitioners seeking to identify existing models suitable for simulating different O&M strategies under uncertainty. The table allows a quick assessment of which studies consider logistics, repair types, and advanced strategies like CBM or opportunistic maintenance.

Table 2.4 identifies the turbine components most frequently modeled in O&M simulations. This can support practitioners deciding which components to prioritize for predictive modeling, and researchers aiming to focus on underrepresented areas (e.g., yaw systems, cooling).

Table 2.2: Methods, failure models, and metocean models used in O&M simulation models

	Simulation method							Failure model							Metocean model							
	Monte Carlo Simulation	Petri Net	Stochastic Petri Net	Fuzzy Logic	Numerical simulation	Bayesian Belief Network	Discrete event simulation	Exponential probability distribution	4-states Weibull distribution	3-state Weibull distribution	2-States Weibull distributions	Homogeneous Poisson Process	Non-homogeneous Poisson Process	stochastic Petri Net	Two-state Markov process	Time-dependent failure rates	Constant failure rates	Average "Wait for good weather"	Markov Chain Model	stochastic Petri Net	Autoregressive moving-average model	Masked Autoregressive
Douard et al. (2012)	✓										✓						✓					
Hofmann & Sperstad (2013)	✓										✓								✓			
Dinwoodie et al. (2013)						✓				✓							✓					✓
Endrerud et al. (2014)							✓						✓			✓						
Le & Andrews (2016)	✓	✓							✓								✓		✓			
Zhu et al. (2016)					✓					✓						✓		✓				
Judge et al. (2019)	✓											✓					✓		✓			
Dao et al. (2020)				✓											✓		✓		✓		✓	
Elusakin et al. (2021)			✓						✓					✓			✓			✓		
Taboada et al. (2021)					✓			✓								✓		✓				

Table 2.3: Types of uncertainty, strategies, and repair types handled by O&M simulation models

	Uncertainties							Strategies				Repair types										
	Costs	Failure rates	Prediction signals	Logistic time	Travel time	Repair time	Deterministic	Corrective	Scheduled	CBM	Opportunistic	Replacement	Major repair	Medium repair	Minor repair	Manual reset	Type 5 ⁵	Type 4 ⁴	Type 3 ³	Type 2 ²	Type 1 ¹	Not specified
Douard et al. (2012)	✓	✓						✓	✓								✓		✓	✓	✓	
Hofmann & Sperstad (2013)			✓	✓		✓		✓	✓	✓												✓
Dinwoodie et al. (2013)	✓							✓	✓													✓
Endrerud et al. (2014)							✓	✓	✓													✓
Le & Andrews (2016)							✓	✓	✓	✓	✓						✓	✓	✓	✓	✓	
Zhu et al. (2016)			✓					✓	✓	✓	✓		✓		✓							
Judge et al. (2019)							✓	✓	✓	✓	✓	✓	✓		✓	✓						
Dao et al. (2020)	✓	✓						✓														✓
Elusakin et al. (2021)							✓	✓	✓	✓							✓	✓	✓	✓		
Taboada et al. (2021)							✓	✓	✓	✓				✓		✓						

¹Heavy component, requires external crane, ²Heavy component, requires internal/external crane (> 800–1000kg),

³Small parts, requires internal crane (< 800–1000kg), ⁴Small parts, inside nacelle, ⁵Small parts, outside nacelle (Le & Andrews, 2016)

In the reviewed literature, two complementary approaches are used to represent the financial impact of O&M decisions. *Deterministic cash-flow models* assume a single expected value for certain cost or revenue terms. *Probabilistic cash-flow models*, by contrast, incorporate uncertainty, e.g., failure rates, electricity prices, or weather-induced downtime, as random variables. The models that integrate this yield distributions for cost metrics, typically estimated through Monte Carlo simulation, as applied in Pattison et al. (2016) and Stock-Williams & Swamy (2019), enabling scenario-based analysis of lifetime costs. To mention another example, an analysis of the reliability and O&M management of OWTs is provided in Taboada et al. (2021), where various maintenance

Table 2.4: Components considered in O&M simulation models

	Blades	Hub	Pitch system	Generator	Electrical system	Control system	Hydraulic system	Mechanical brakes	Yaw system	Converter	Transformer	Tower	Shaft	Bearings	Gearbox	Lubrication system	Cooling system	Mooring lines	Floating foundation
Douard et al. (2012)	✓		✓	✓	✓	✓		✓	✓				✓	✓	✓				
Hofmann & Sperstad (2013)	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓		✓	✓		
Dinwoodie et al. (2013)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓		
Endrerud et al. (2014)	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓			✓	✓		
Le & Andrews (2016)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
Zhu et al. (2016)	✓	✓	✓	✓						✓		✓	✓						
Judge et al. (2019)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓				
Dao et al. (2020)	✓		✓	✓	✓					✓			✓	✓	✓				
Elusakin et al. (2021)	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓			✓	✓
Taboada et al. (2021)	✓		✓	✓	✓	✓		✓	✓				✓	✓	✓				

strategies are compared and optimized while evaluating their cost-benefit profiles using stochastic and probabilistic life cycle cost (LCC) models.

Uncertainty remains a key challenge in long-term O&M modeling, particularly in the estimation of failure probabilities and major component replacement rates for next-generation offshore turbines, for which field data are scarce. In fact, the results of O&M simulation models are highly dependent on the assumptions and parameters used in the model, which can be difficult to validate and adjust, given the scarcity of real-life data, for this reason many papers such as Scheu et al. (2017), Dao et al. (2020) and Zhu et al. (2016) focus on uncertainties and on evaluating their impact on the simulation results. Jenkins et al. (2022) applied structured expert elicitation to estimate major replacement rates for future offshore wind technologies. Their approach helped quantify not only expected values but also credible bounds, which are essential for robust decision-making in early planning stages. More recently, Centeno-Telleria et al. (2024a) extended this work by performing a detailed sensitivity analysis on replacement cost and frequency assumptions for floating wind, showing that even small shifts in these parameters can significantly alter lifecycle cost projections and optimal maintenance strategies. These studies emphasize the importance of capturing uncertainty explicitly and demonstrate how expert-based methods can complement data-driven approaches when empirical evidence is limited. Furthermore, various sensitivity analyses have been carried out, demonstrating the use of simulation tools to investigate the effects of various inputs and uncertainties on O&M costs. For example, Martin et al. (2016) identified 14 important inputs for calculating O&M costs, including failure rates, component costs, repair duration, and shift length. For availability, they considered factors like failure rates, duration of the shifts, and resource availability. The sensitivity analysis aimed to understand the impact of these inputs on O&M costs and availability. The impact of the chance of failure on asset performance evaluation is highlighted in Scheu et al. (2017). They investigated the implications of different failure models and frequencies on O&M models and simulation tools. By employing various failure models and frequencies, they compared the resulting availability of assets and drew conclusions about the effects on O&M strategies. A maintenance policy for wind turbines using a multi-component system approach is proposed in Zhu et al. (2016) where the authors analytically derive the system's lifecycle costs under the proposed policies using numerical simulations. Additionally, the study factors in decision-making uncertainties by exploring how variations in prediction signal accuracy affect both the selection of decision criteria and the overall performance of the maintenance strategies. Finally, the effects of uncertainties related to

weather, repair time, and the operational range of vessels on maintenance costs are examined in Ambühl & Sørensen (2017). Using simulation tools, they analyzed the impact of these uncertainties on maintenance costs. The study aimed to quantify the influence of these factors and provide insights into optimising maintenance strategies.

An essential aspect in long-term O&M modeling is the treatment of metocean conditions, which strongly influence site accessibility and thereby turbine availability. While traditional simulation models incorporate metocean effects through stochastic weather windows or time-sequential data (e.g., NOWIcob, StrathOW), recent analytical approaches have also begun to integrate these factors explicitly. Centeno-Telleria et al. (2024b) extend their Markov-based framework by incorporating high-resolution metocean datasets, such as significant wave height and wind speed, to compute site-specific accessibility profiles. This enables a more realistic estimation of downtime and its impact on energy yield. Rowell et al. (2022) further demonstrate how floating wind sites differ markedly from fixed-bottom locations in terms of access windows, which has major implications for maintenance vessel strategies and service intervals. To support such evaluations at early planning stages, Konuk et al. (2023) propose a technology-informed accessibility metric that combines metocean statistics with logistic feasibility, offering a standardized way to rank offshore renewable sites by maintainability. These studies highlight that modeling accessibility is not only critical for accurate O&M cost and performance prediction, but also increasingly feasible thanks to the growing availability of open metocean datasets.

Overall, O&M simulation models have evolved from basic cost estimators into multi-layered decision support tools that integrate failure dynamics, weather accessibility, vessel logistics, and financial forecasting. The continued development of probabilistic models, expert-informed uncertainty quantification, and hybrid analytical frameworks is expanding their applicability across design, planning, and operational stages. While simulation remains computationally demanding and data intensive, emerging approaches such as Markov-based models and structured expert elicitation offer scalable alternatives that retain critical realism. As offshore wind enters a new phase of deployment with larger, floating turbines in remote locations, the role of simulation in supporting robust, cost-effective O&M strategies will only become more pivotal.

2.3.2 Policy design

This subsection reviews optimization models developed to design cost-effective O&M policies at the strategic level. The primary objective of these models is to define the optimal timing to perform maintenance on the wind turbine components based on their age with respect to their expected lifetime. These models are increasingly integrated into DSSs, which combine data, modeling tools, and interfaces to assist planners in determining long-term strategies under uncertainty. The significant impact of age groups and age thresholds on the total cost of maintenance is highlighted in Sarker & Ibn Faiz (2016), where the authors propose a cost model for OWT maintenance based on opportunistic preventive maintenance.

Optimal age replacement policies for both parallel and series systems featuring dependent components are examined in Safaei et al. (2020). The study employs a copula framework to model the interdependencies and to formulate the corresponding optimal decision problems. The results offer valuable insights to reduce costs or increase system availability, and the paper concludes with a discussion on applying these findings to OWT systems.

A maintenance strategy for OWFs that integrates degradation failures, incidents, and age-based opportunities is proposed in Xie et al. (2019) and Xie et al. (2020). This strategy aims to reduce O&M costs, and its comparative analysis shows significant cost reductions with respect to traditional opportunistic maintenance strategies. Accessibility constraints are taken into account by modeling weather series using the Markov chain method, and the optimal preventive maintenance age along with the opportunistic maintenance age are evaluated through numerical optimization to minimize the total expected cost of maintenance.

The opportunistic maintenance strategy presented in Li et al. (2020) optimize subsystem intervals by grouping tasks using a state transition degradation model. The results demonstrate the effectiveness of the proposed model in reducing the total maintenance cost compared to conventional preventive maintenance strategies.

A framework for optimising maintenance strategies with multiple objectives is introduced in Li et al. (2022b). The study identifies three key uncertainties impacting the model: the random nature of failure times, discrepancies between actual and predicted failure occurrences, and the unpredictability of maintenance outcomes. To resolve conflicts among objectives, Monte Carlo simulation is employed along with the Non-dominated Sorting Genetic Algorithm (NSGA-II) to determine optimal maintenance decisions.

A single-component maintenance optimization model designed for OWT upkeep under time-varying cost conditions is presented in Schouten et al. (2022). Building on existing policies, the study demonstrates that the optimal strategy in such a context is a time-dependent age replacement policy.

Table 2.5 summarizes key studies on age threshold optimization for O&M planning. It helps modelers compare how age-related decision variables, constraints (like accessibility), and optimization techniques are handled across different works. This table supports the understanding and practical design of age-based preventive maintenance policies.

In summary, the reviewed literature shows that strategic decision-making in OWF O&M relies on both simulation and optimization models to evaluate and define long-term maintenance strategies.

Table 2.5: Optimization objectives, design variables, constraints, assumptions, and methods in studies using age-threshold-based maintenance policies

	Objective Function	Design Variables	Constraints	Assumptions	Methodology
	Yearly O&M Cost Total O&M Cost Availability Downtime	Age Thresholds Age Groups Maintenance Intervals	Age groups of same length Accessibility Component Dependencies Failure Replacement Maintenance Interval Safety	Same components WT Weibull Failure Distribution Fixed Maint. Time Negligible Maint. time Perfect Maintenance	Exhaustive Search Markov Method Simulation Analysis Copula Framework Monte Carlo NSGA-II LP/MILP
Sarker & Ibn Faiz (2016)	✓	✓ ✓	✓	✓ ✓ ✓	✓
Xie et al. (2019)	✓	✓	✓	✓ ✓	✓ ✓
Xie et al. (2020)	✓	✓	✓	✓ ✓	✓
Safaei et al. (2020)	✓ ✓	✓	✓	✓	✓
Li et al. (2020)	✓	✓	✓	✓ ✓	✓
Li et al. (2022b)	✓ ✓	✓	✓	✓ ✓	✓ ✓
Schouten et al. (2022)	✓	✓	✓	✓	✓

Simulation models support scenario-based assessment of predefined strategies under uncertainty, helping quantify cost, reliability, logistics, and accessibility trade-offs over the wind farm lifecycle. Optimization models, by contrast, are used to prescribe maintenance policies, for example by identifying cost-effective age thresholds, grouping rules, or risk-based intervention criteria under simplified cost-reliability formulations. The reviewed studies therefore highlight the importance of determining appropriate age thresholds, accounting for risk and time-varying costs, and exploiting multiple maintenance opportunities when defining long-term strategies. Together, these tools form the backbone of strategic decision-support systems for early-stage investment, design, and planning decisions.

2.4 Tactical decision-making

Transitioning from the strategic overview, our focus shifts to the tactical echelon. In this section, we investigate the tactical perspective of O&M optimization for OWFs. We analyze the main factors and objectives in fleet selection and spare part management for OWF O&M.

2.4.1 Fleet selection

The selection of the fleet for OWF O&M involves determining the optimal combination of vessels and resources needed to effectively and efficiently carry out maintenance activities. One of the primary goals is to minimize the overall O&M costs associated with the wind farm, considering factors such as vessel capabilities, seakeeping abilities, chartering costs, crew transfer costs, maintenance costs, and revenue loss due to downtime Sperstad et al. (2017); Dalgic et al. (2014). Additionally, different maintenance strategies have to be considered in fleet selection. The optimal fleet should be capable of supporting these strategies to ensure the reliability and availability of the wind turbines Dinwoodie et al. (2012).

This literature review provides an overview of the current research on fleet selection for OWFs' O&M, highlighting significant findings and differences in approaches.

In Sperstad et al. (2017), the authors compare six decision support tools for selecting optimal vessel fleets for O&M operations, including MAINTSYS, Shoreline Endrerud et al. (2014, 2015), the MARINTEK fleet optimization model Halvorsen-Weare et al. (2013), and ECN's O&M Tool Obdam et al. (2011). Notably, the comparison also includes ECUME and StrathOW, which are originally developed for strategic-level analysis (as reviewed in Section 2.3.1), but are here evaluated based on their ability to support medium-term decisions such as fleet composition. This reflects the cross-horizon applicability of certain tools when adapted to different planning contexts. The study assesses the robustness of these tools in minimizing O&M costs under varying weather conditions, vessel specifications, and failure scenarios.

Logistics planning for OWF operations is addressed in Dalgic et al. (2015b), where a Monte Carlo simulation approach is used to evaluate cost-effective allocation of O&M resources, considering different vessel types and weather conditions. The study emphasizes sustaining continuous power production by efficiently operating O&M fleets. A mathematical model to optimize fleet selection, taking into account turbine availability, O&M costs, and revenue loss, is introduced in Dalgic et al. (2015d). Factors such as weather conditions, vessel capabilities, and maintenance

strategies are considered. Uncertainty in weather conditions, travel times, and repair times is addressed in Halvorsen-Weare et al. (2013) through a mathematical model aimed at minimizing the expected total cost of the vessel fleet over the wind farm's lifetime while satisfying turbine availability requirements. A branch-and-cut algorithm is employed to solve this model. A time-domain Monte Carlo approach for investigating the optimal chartering strategy for jack-up vessels is proposed in Dalgic et al. (2015e). Additionally, the study highlights regional collaborations between different parties as an alternative solution for cost reduction, integrating climate parameters, failure characteristics, and vessel specifications to determine the best overall strategy. The potential benefits of resource sharing among service providers for OWF maintenance are examined in uit het Broek et al. (2019). Instead of each provider leasing its own Jack-Up vessel, the study proposes two alternatives: one where the vessel is purchased and shared, and another that combines vessel and harbor sharing. Simulation modeling is used to evaluate these scenarios, revealing significant cost benefits that vary with the number and distance of service providers involved. Considering future major maintenance costs for OWFs is also a quite important challenge, as highlighted in Dinwoodie & McMillan (2014), where a modeling framework assesses different operating strategies under various scenarios, incorporating offshore heavy lift vessel strategies. Four operational strategies are proposed: fix on fail, batch repair, annual charter and purchase based on different life failure distributions. A mixed-integer linear programming (MILP) model for vessel fleet selection and maintenance scheduling is presented in Gutierrez-Alcoba et al. (2019). The MILP model considers a priori information to optimize fleet composition and scheduling. A heuristic approach is introduced to assess the effect of a priori information on the MILP model's performance. A two-stage stochastic programming model for fleet optimization in OWF maintenance is proposed in Stålhane et al. (2019). The model integrates preventive and corrective maintenance tasks, optimising vessel fleet composition, and scheduling. The study considers uncertainty in maintenance needs and weather conditions. Genetic algorithms are employed to optimize O&M asset configurations in Rinaldi et al. (2020), with the approach aiming to reduce operating costs while enhancing reliability and availability. Finally, Stålhane et al. (2016) present a DSS designed to aid in fleet selection for OWF maintenance. They use a stochastic optimization model to determine the optimal configuration of vessels.

The selected literature describes simulation models, mathematical optimizations, and heuristic methods, to address the challenges associated with fleet selection. Table 2.6 outlines the main modeling approaches for offshore maintenance fleet optimization. It is intended for researchers designing fleet selection models and for decision-makers comparing strategies such as purchasing vs chartering. The table can also guide tool developers in choosing appropriate constraints and objectives for real-world fleet planning tools.

2.4.2 Spare part management

Spare part management in the O&M of OWFs involves effectively managing the availability, procurement, storage, and utilization of spare parts required for the maintenance and repair of wind turbines. It encompasses various activities and considerations to ensure uninterrupted operation, minimize downtime, and optimize costs. Efficient spare part management is crucial for the O&M of OWFs, impacting maintenance costs, downtime, and overall profitability. SPM aims to ensure that the right spare parts are available at the right time and place, minimizing the inventory holding costs and the downtime costs due to stockouts. However, determining the optimal inventory levels

Table 2.6: Fleet optimization approaches based on objective functions, constraints, and methods

Halvorsen-Weare et al. (2013) Dalgic et al. (2015d) Dalgic et al. (2015e) Stålhamre et al. (2019) uit het Broek et al. (2019) Rinaldi et al. (2020)	Objective Function	
	↖	Fixed Costs of vessels and base
	↖	Variable Cost of chartering
	↖	Downtime Cost
	↖	Total O&M Cost
	↖	Availability
	↖	Reliability
	Design Variables	
	↖	Vessels to Purchase
	↖	Vessels Chartering
	Constraints	
	↖	Budget
	↖	Seakeeping Ability / Capacity
	↖	Vessel Availability
	Assumptions	
	↖	Tasks are Pre-Determined
	↖	Deterministic
	↖	Fixed Inputs
	Methodology	
	↖	MIP
	↖	Monte Carlo Simulation
	↖	Stochastic Prog.
	↖	Dantzig-Wolfe Decomposition
	↖	Genetic Algorithms
	↖	Reliability Analysis

for spare parts is a complex and challenging task, as it involves multiple factors such as the failure characteristics of the components, the lead time and cost of procurement, the storage and handling costs, the maintenance policies and strategies, and the uncertainty and variability of demand.

Regardless of the specific methods applied, the optimization of inventory level generally consists in the following steps:

1. **System Analysis:** Analyze historical failure data to identify patterns, frequencies, and severity. Categorize critical components by risk using methods like FMEA Ferdinand et al. (2018) or Markov processes Zhang et al. (2018) to determine spare parts needs.
2. **Demand Forecasting:** Estimate future failure probabilities and timings, considering demand uncertainty and lead times, to predict spare parts consumption and replenishment requirements.
3. **Cost Modeling:** Develop a cost model that includes procurement, handling, stockout, and excess inventory costs.
4. **Inventory optimization:** Balance holding costs against benefits to determine optimal order quantity, reorder point, and safety stock based on criteria like total cost minimization, service level, or risk reduction.

The above steps, summarized in Figure 2.8, provide a general framework for SPM optimization; however, different methods and techniques may be applied within each step depending on the specific characteristics and requirements of each problem. In this chapter, we review some of the most common approaches for SPM optimization, specifically in the context of OWFs.

A critical review Tusar & Sarker (2022b) compares various spare parts control strategies for OWFs. They systematically analyze prominent models, highlighting significant contributions and claiming cost reductions up to 51% compared to traditional scheduling methods, while integrated approaches have demonstrated 27% savings. This review provides valuable guidelines for effective and efficient spare parts management throughout the wind turbine's life cycle.




Figure 2.8: Generic flowchart for Spare Part Management optimization, adapted by Tusar & Sarker (2022b)

A model that integrates wind turbine accessibility and restrictive factors to minimize maintenance costs is proposed in Tracht et al. (2013). Their findings reveal that restricted accessibility affects preventive maintenance scheduling and spare parts demand, suggesting the potential to reduce costs and enhance profitability by utilizing resources and technicians more efficiently.

Optimization schemes and component updating techniques have been explored to optimize spare parts quantity and reduce costs. An optimization scheme based on component updating and a Markov process is introduced in Zhang et al. (2018). By considering availability as the optimization objective, they effectively reduce spare parts management costs. Simulation analysis demonstrates substantial cost savings compared to traditional methods. This approach presents a promising solution to spare parts quantity optimization.

A problem similar to SPM optimization has been tackled by Neves-Moreira et al. (2021), where a mixed-integer programming model for determining and validating repair kits under varying weather conditions is proposed. They emphasize the importance of adapting repair kits based on weather forecasts and demonstrate considerable downtime reductions through emergency resupplies. Their methodology comprises three phases, offering insights into the repair kit problem and aiding practitioners in defining repair kits under different conditions.

Table 2.7 synthesizes key characteristics of spare parts management (SPM) optimization methods in the context of offshore wind farms.

2.5 Operational decision-making

As we presented the medium-term considerations of fleet and inventory management, our attention turns to the short-term horizon. In this section, we explore the short-term perspective of O&M optimization for OWFs. This horizon is dedicated to the short-term scheduling of maintenance tasks, on a daily or weekly basis, where real-time factors like weather, component condition, and logistical constraints play a crucial role. Here, we present some of the latest developments in prognostic and diagnostic methods, which are used to inform decisions made at this horizon. We also present an overview of different techniques for optimising routing and scheduling decisions.

2.5.1 Diagnostics and prognostics

This section reviews the importance of diagnostics and prognostics for OWTs and some of the methods and applications in this field. The ability to assess and predict component failures has a

Table 2.7: Characteristics of main SPM methods proposed in the literature

	Objective Function		Design Variables			Constraints			Methodology	
	Downtime Cost	O&M Cost	Spare Parts Inventory	Detection Time	Detection Frequency	Availability	Budget	Detection Cycle	FMEA	Markov Process
Ferdinand et al. (2018)	✓		✓			✓	✓		✓	
Zhang et al. (2018)		✓		✓	✓	✓		✓		✓

direct impact on optimising the maintenance strategies employed over the wind farm's lifetime. By exploring CBM technologies in this section, we aim to underline their significance in influencing the broader decision-making process and highlight the synergy between technology and strategic planning.

Diagnosis means evaluating the current health or performance of a system based on its current state and past behaviour, while prognosis means predicting the future health or performance of a system based on its current state and past behaviour. Diagnosis and prognosis are essential for CBM because they provide information about the condition and the RUL of the wind turbine components, which can be used to optimize the maintenance decisions (Pires Leite et al., 2018). For example, diagnosis can help identify the type, location and severity of a fault, while prognosis can estimate the time to failure or the probability of failure within a given time horizon.

There are several methods and techniques for the diagnostics and prognostics of offshore wind turbines (OWTs), which can be classified into data-driven, model-based, or hybrid approaches (Kang et al., 2019):

- Data-driven methods use historical or real-time data from sensors or SCADA systems to extract features and patterns that indicate faults or degradation. One prominent area of research focuses on SCADA data analysis for CBM (Tautz-Weinert & Watson, 2017). The importance of interpreting SCADA data to enable effective CBM is highlighted in Yang et al. (2013). These approaches enable researchers to extract valuable insights from operational data that are readily available to turbine operators without the need for potentially expensive sensors. Operational curve analysis, such as power curve, rotor speed curve, and blade pitch angle curve, is used to detect deviations from normal behavior, as shown in Pandit & Infield (2018), aiding in early fault detection and proactive maintenance. Deep learning techniques also play a significant role in this domain. For instance, Rahimilarki et al. (2022) present a time-series fault detection method using convolutional neural networks (CNNs), capable of identifying anomalies from turbine operational data. A recent contribution to data-driven methods is the Gaussian mixture autoencoder developed by Fernandez-Navamuel et al. (2025), which enables uncertainty-aware damage identification in floating wind turbines by modeling complex, nonlinear data distributions. In addition, Papatzimos et al. (2019) introduces a probabilistic model to evaluate alarm counts for individual subassemblies and the entire wind farm. The model achieved 99.3% accuracy for long-term forecasts and over 83% accuracy for short-term predictions, using two years of met mast data. Colone et al. (2019) combines statistical classification of SCADA data with event tree analysis to predict abnormal component behavior and assess predictive maintenance benefits, comparing Naïve Bayes and Neural Network classifiers.
- Model-based methods rely on physical or mathematical representations of wind turbine components to simulate their behavior and compare predicted and observed responses. In Ghane et al. (2018), a high-fidelity gearbox model is used to estimate and detect wear in the main shaft bearing via statistical fault diagnostics. Similarly, Romero et al. (2018) present a CBM system analyzing power-dependent vibrations in the drivetrain to identify specific fault-related patterns. For floating offshore wind turbines, tailored architectures have been proposed. A mixed model and signal-based fault diagnosis architecture to detect and isolate faults specific to floating turbine dynamics is introduced in Liu et al. (2021).

- Hybrid approaches combine data-driven and model-based techniques to capitalize on their respective strengths. Bayesian methods are particularly well suited to this purpose. For example, Asgarpour & Sørensen (2018) present a generic diagnostic model that integrates outputs from multiple fault detection systems such as vibration, temperature, and oil particle sensors, using Bayesian updating to reduce uncertainty based on inspection outcomes. Digital Twin (DT) technologies can also be interpreted as comprehensive hybrid solutions, since they integrate physical models, operational data, and updating mechanisms into a unified asset representation. A review of DT-based frameworks for OWFs (Xia & Zou, 2023) highlights their roles in failure monitoring, RUL prediction, safety and ecological management, O&M decision-making, and design optimization.

Table 2.8 provides a comparative summary of prominent health monitoring methods used for condition monitoring and prognostics in offshore wind turbines. It is intended for the reader to evaluate the relative strengths and scopes of different fault detection and prediction models. The table distinguishes between diagnostic and prognostic models, between SCADA-based and sensor-rich CBM methods, and lists typical analytical approaches. This overview supports informed decisions about the integration of health-monitoring technologies into predictive maintenance pipelines.

By using diagnostic and prognostic information, maintenance activities can be planned and scheduled according to the current and future condition of the wind turbine components. This can avoid unnecessary or premature maintenance actions, as well as prevent catastrophic failures or downtime due to undetected faults or degradation.

2.5.2 Routing and scheduling

The routing and scheduling problem in the O&M of OWFs involves determining the optimal routes and schedules for vessels, technicians and resources involved in maintenance activities. It aims to

Table 2.8: Characteristics of main health diagnosis and prognosis methods proposed in the literature

	Methodology									SCADA vs CBM		Prognostic vs Diagnostic		
	Gaussian process	Normal operational conditions	Convolutional Neural Network	FDAE	FIE	Statistical fault diagnosis	Naïve Bayes classifier	Neural Network	Bayesian update	Wind speed statistics	SCADA	CBM	Prognostic	Diagnostic
Yang et al. (2013)		✓									✓			✓
Romero et al. (2018)		✓										✓		✓
Pandit & Infield (2018)	✓										✓			✓
Ghane et al. (2018)						✓						✓		✓
Asgarpour & Sørensen (2018)									✓			✓		✓
Colone et al. (2019)							✓	✓			✓		✓	
Papatzimos et al. (2019)										✓	✓			✓
Liu et al. (2021)				✓	✓						✓	✓		✓
Rahimilarki et al. (2022)			✓								✓			✓

efficiently allocate resources, minimize travel time and costs, optimize task sequencing, and ensure timely completion of maintenance tasks. This section reorganizes the reviewed literature into three modeling perspectives: fleet logistics, task sequencing and resource allocation, finally providing information on how weather, operational constraints and uncertainties have been incorporated into operational DSSs.

One of the primary goals of vessel routing is to optimize the paths taken by vessels to access and service the OWTs. This involves considering factors such as wind and wave conditions, vessel capabilities, and the location of the wind farm Irawan et al. (2017). The objective is to minimize the cost by minimizing travel time and distance, reduce fuel consumption, and ensure safe and efficient transportation to and from the wind farm Ade Irawan et al. (2023). A mathematical model by Irawan et al. (2021) considers vessels capable of remaining offshore for several shifts while managing large repairs, the paper also presents a precise calculation of downtime costs. A simulation is employed to assess the model's performance, using a rolling horizon heuristic method to solve it. Moreover, the study demonstrates how coupling a mathematical model with simulation can help evaluate strategic decisions related to vessel fleet composition. Another mathematical model for planning and executing maintenance tasks is presented in Raknes et al. (2017). The study also introduces a heuristic solution method based on a rolling-horizon approach and includes computational experiments to support strategic decisions on fleet composition. Docking operations have been studied in Wu (2014) which focuses on evaluating operational limits for different access systems. Their work introduced methodologies for numerical analysis of accessing techniques, such as active motion-compensated access devices and fenders. These methods can be utilized to assess docking operability, contributing to the safe and efficient transportation of technicians and materials. To improve the daily route planning of maintenance vessel operations, Lazakis & Khan (2021) developed an optimization framework known as OptiRoute. This framework employs heuristic and clustering techniques while taking into account weather, vessel characteristics and failure data. The experimental results showed that the operational windows were significantly extended, particularly when Service Operation Vessels (SOVs) and Crew Transfer Vessels (CTVs) were used in tandem. Additionally, Sinha et al. (2013) introduced a software tool named GESTIONE to optimize OWT maintenance. The package uses techniques such as Bayesian Networks, Artificial Neural Networks, and FMECA, utilizing condition monitoring data to predict component failures, thereby reducing O&M costs and providing insights into wind turbine operations. Furthermore, Dawid et al. (2018) introduced a decision support tool for OWF vessel routing under uncertainty. Their methodology incorporated uncertain maintenance action times into vessel routing, considering realistic OWF problems. The tool highlighted the importance of considering uncertainties in vessel routing decisions.

Efficient task sequencing is crucial to minimize downtime and maximize the utilization of resources. This includes determining the order in which maintenance tasks are performed, considering their dependencies, and assigning suitable time slots to each task. The aim is to minimize waiting times, avoid conflicts in resource allocation, and optimize the overall schedule for maintenance activities Fan et al. (2019); Dalgic et al. (2015b). Contributing to the optimization of the scheduling of maintenance activities, risk-based planning methods for O&M of wind turbines are discussed in Nielsen & Sørensen (2014) with the goal of evaluating the probability of failure based on inspections and finally minimizing the expected costs. They present various approaches, including decision rules based on crude Monte Carlo simulations, limited memory influence diagrams, Bayesian updating and Markov decision process. The study concludes that decision rules based on the probability of failure yield lower costs compared to the methods based on Monte Carlo simulations. The

integration of advanced mathematical techniques in maintenance management systems has been explored to improve the efficiency and cost-effectiveness of offshore wind power generation. A novel architecture and system for Reliability Centred Maintenance (RCM) was presented in Pattison et al. (2016). Their integrated system demonstrated fault detection, reliability and maintenance modeling, and optimized maintenance scheduling, leading to maximized availability and revenue generation of turbines. A model that maximizes the profit of an OWF was developed in Mazidi et al. (2017) by solving the Single-Level (SLP) and Bi-Level (BLP) optimization problem, considering various network and wind farm constraints. The model incorporates information from condition monitoring systems and provides a scalable solution for planning maintenance schedules in a deregulated power system. Lastly, Liu et al. (2019) proposed an O&M scheduling strategy model for OWFs that aimed to minimize maintenance scheduling costs. By optimising maintenance division, the number of model optimizations can be simplified, providing practical guidance for scheduling O&M tasks in OWFs. Alternative modeling techniques such as constraint programming also show promise. Froger et al. (2018) develop a short-term scheduling model for offshore wind farms that simultaneously handles vessel routing, technician assignment, and turbine-specific fault types. Their framework flexibly incorporates spatial constraints (e.g., distance, location), temporal availability, technician skills, and weather-induced access windows. While not yet commonly integrated into prognostic-driven frameworks, such constraint-based models offer powerful tools for enhancing the realism and feasibility of short-term maintenance planning.

The routing and scheduling problem also involves allocating resources such as vessels, crew members, and equipment to different maintenance tasks Dai et al. (2015). It requires considering the availability and capabilities of resources, task requirements, and constraints. The objective is to optimize resource allocation to ensure efficient task execution and minimize idle time. An optimization model that considers multiple vessels, time periods, O&M bases, and wind farms was developed in Irawan et al. (2017). The study aimed to minimize overall maintenance costs, including travel, technician, and penalty-related expenses. The results showed that using multiple vessels and O&M bases can significantly reduce maintenance costs, leading to improved cost-effectiveness. The model was later extended to include a service operation vessel (SOV) and a safe transfer boat (STB), enabling the transfer of technicians and equipment to turbines, as described in Ade Irawan et al. (2023). A mathematical formulation of this problem is presented in Dai et al. (2015), where the authors focus on determining the cheapest assignments of turbines and routes to vessels. Resource-related decision-making in corrective maintenance has also been addressed through a mathematical model proposed in Nachimuthu et al. (2019), which identifies cost-effective resource combinations and offers considerable cost savings. The model accounts for uncertainties in turbine failure information and supports stakeholders in making critical short-term decisions.

Decision support systems and optimization models play an important role in solving the routing and scheduling problem. These tools help evaluate different scenarios, incorporate real-time data, consider uncertainties in weather conditions and task duration, and optimize routing and scheduling decisions based on specific objectives, constraints, and performance metrics. Weather and operational constraints that affect vessel O&M activities have to be taken into account. This includes considering the effect of weather conditions on vessel operability, such as maximum wave height for safe access to turbines. Constraints related to vessel maintenance, crew shift schedules, and regulatory requirements may also be considered in the optimization process Dalgic et al. (2015a). Historic wind farm data have also been used to evaluate maintenance planning methods and estimate their respective return on investment (ROI), as shown in Stock-Williams & Swamy (2019). Automated daily maintenance planning was shown to bring significant financial

benefits to the O&M of OWFs. Uncertainties in maintenance activities pose significant challenges in OWF maintenance. To address this, efficient solution methods for both deterministic and uncertain maintenance routing problems have been proposed in Irawan et al. (2021). Their simulation-based optimization algorithm considered uncertain factors such as travel time, required maintenance time, transfer time for technicians and equipment, and unpredictable broken-down turbines. These methods provide valuable insights into handling uncertainties and optimising maintenance strategies.

Table 2.9 presents an overview of key optimization and simulation methods used for routing and scheduling tasks in offshore wind farm maintenance. This table is intended to provide guidance in evaluating suitable methods for task sequencing, vessel routing, and scheduling under uncertainty. It allows comparison across heuristic, exact, and machine learning-based approaches, providing a foundation for selecting a method based on computational requirements, uncertainty modeling, or solution precision.

In conclusion, the selected literature highlights representative optimization models, simulation-based algorithms, mathematical formulations, and decision support tools employed in the routing and scheduling of O&M activities in OWFs. The studies emphasize the importance of minimizing maintenance costs, considering uncertainties, and integrating advanced techniques for improved reliability, cost-effectiveness, and revenue generation.

2.6 Gaps and limitations

This section explores both the limitations of widely adopted modeling approaches and the underrepresented opportunities in the current offshore wind O&M literature. While our review has focused primarily on mainstream methods, such as stochastic modeling, optimization-based planning, and rule-based simulations, several promising alternatives remain relatively unexplored or fragmented across decision-making layers. We also highlight overlooked dimensions such as the limited inte-

Table 2.9: Methodologies used in the main Routing & Scheduling methods proposed in the literature

	MILP	ILP	Branch and Bound	GA	Random Forests	Dynamic Bayesian Network	Memetic Algorithm	MIP	Rolling-Horizon Heuristics	Large Neighbourhood Search	Analytical model	Heuristic optimization	SLP/BLP	Bayesian Network	Artificial Neural Network	FMECA	Monte Carlo Simulation	Particle Swarm optimization
Sinha et al. (2013)														✓	✓	✓		
Dai et al. (2015)			✓															
Pattison et al. (2016)					✓	✓	✓					✓						
Raknes et al. (2017)								✓	✓			✓						
Irawan et al. (2017)	✓	✓						✓	✓	✓		✓						
Mazidi et al. (2017)													✓					
Dawid et al. (2018)																	✓	
Liu et al. (2019)												✓						✓
Stock-Williams & Swamy (2019)				✓								✓						
Nachimuthu et al. (2019)											✓							
Lazakis & Khan (2021)												✓						
Ade Irawan et al. (2023)										✓		✓						

gration of biodiversity considerations into decision frameworks. Addressing these gaps is essential for advancing more holistic, cost-effective, and environmentally conscious O&M solutions.

2.6.1 Prognostic-driven scheduling

The advent of advanced data analytics and predictive modeling has revolutionized short-term scheduling approaches. By integrating real-time data from sensors with prognostic models, it is now possible to predict component health and maintenance needs with greater accuracy Le & Andrews (2016). This technological progression has facilitated the development of prognostic-driven approaches, which are known to be highly effective in optimising short-term maintenance schedules Van Horenbeek & Pintelon (2012).

A recent contribution by Frederiksen et al. (2024) proposes a methodology to quantify the efficiency and economic benefit of predictive maintenance compared to preventive strategies. Their approach introduces separate indicators for predictive accuracy and realized maintenance efficiency, validated through a realistic OWF case study. The study finds that predictive maintenance can yield cost savings, but only if the model achieves high accuracy and the system responds efficiently. Otherwise, the economic advantage may be lost.

The selected literature outlines different strategies for optimising both vessel routing and the timing of maintenance operations in OWFs. These approaches focus on reducing travel times, streamlining resource allocation, and ensuring transportation is conducted safely and efficiently to cut costs Irawan et al. (2017); Ade Irawan et al. (2023). However, many of these methods have traditionally relied on fixed age thresholds for scheduling maintenance. Such thresholds are determined by models that base maintenance timing solely on a component's age rather than its real-time condition Sarker & Ibn Faiz (2016).

Age threshold optimization techniques have been extensively explored in Section 2.3.2. These techniques use stochastic models to determine the most cost-effective timing for maintenance actions Safaei et al. (2020). Although effective, these traditional methods do not fully make use of the potential benefits of prognostic information, which could further enhance scheduling efficiency.

Recent advancements highlight a shift towards prognostic-driven scheduling models that utilize probabilistic RUL predictions to inform maintenance decisions Li et al. (2020). Unlike traditional fixed age threshold models, these approaches assess the actual condition and anticipated future health of components, enabling a more adaptive and responsive maintenance strategy that may reduce costs and enhance system reliability Van Horenbeek & Pintelon (2012).

The integration of prognostics for scheduling O&M marks a significant shift from reactive to proactive maintenance strategies. Advanced monitoring and prognostic strategies have shown potential for substantial cost savings and improved operational efficiency. For example, studies have indicated that employing prognostic-driven models can lead to a reduction in O&M costs of up to 8%, and a reduction in lost production by up to 11%, primarily through early intervention to prevent failures and major component replacements Turnbull & Carroll (2021).

Recently, Borsotti et al. (2024) presented a review of prognostic-driven scheduling models used in other fields where this approach is already commonly used, such as aircraft maintenance, identifying potential applications for OWFs. The model they propose incorporates a Model Predictive Control (MPC) framework, which optimizes maintenance schedules based on probabilistic RUL

predictions. This framework has shown promising results in dynamically adjusting maintenance plans to improve cost-effectiveness.

The adoption of predictive models must also consider the trade-off between their predictive performance and the practical cost-benefit of implementing them, as emphasized in recent techno-economic evaluations Canet et al. (2023). High-performance models often require substantial data acquisition and integration efforts, including real-time condition monitoring systems and tailored prognostic algorithms. These requirements can drive up operational costs and introduce complexity that may not yield proportional benefits unless carefully matched to the asset's criticality, degradation profile, and accessibility constraints.

Moreover, several studies highlight that the high upfront capital cost of condition monitoring sensors remains a barrier to large-scale deployment. Shafiee (2015) argues that operators must weigh these sensor investments against expected gains in availability and maintenance efficiency.

In conclusion, while traditional age-threshold optimization has provided a useful foundation for medium- to long-term planning, such as selecting the strategy or initiating criteria for maintenance interventions, it is inherently limited by its reliance on static assumptions about failure rates and component aging. Prognostic-driven models, by contrast, enable more dynamic and responsive decision-making by leveraging condition-based information and probabilistic RUL estimates. These models allow maintenance schedules to adapt in real time to asset health, environmental conditions, and operational constraints, aligning short-term actions with evolving system risks. However, their implementation requires careful evaluation of predictive performance relative to the costs and logistical burden of data acquisition and system integration.

2.6.2 Considering the interconnections

The challenges related to strategic, tactical and operational decision-making in the O&M of OWFs have been addressed separately in the literature. However, it is important to note that many of the design variables, constraints, and optimization objectives that are used in these separate approaches are interconnected and affect each other across the different decision-making layers.

For example, strategic decisions can have a significant impact on the availability and readiness of resources for maintenance in the medium- and short-term horizons. Similarly, operational decisions, such as the routing and scheduling of maintenance tasks, can impact the overall maintenance strategy and budget allocation in the long-term.

Considering these dependencies holistically can help resolve limitations faced when models operate in silos. For example, simulation tools like NOWIcob and StrathOW provide detailed short- and medium-term maintenance cost and accessibility estimates, but their outputs are rarely integrated into long-term strategic planning, potentially missing cumulative effects of climate or logistic constraints Dinwoodie et al. (2013); Hofmann & Sperstad (2013). Similarly, analytical models such as those proposed by Centeno-Telleria et al. Centeno-Telleria et al. (2023, 2024b) offer fast long-term insights, but may oversimplify short-term scheduling constraints unless coupled with detailed operational models. Furthermore, insights from short-term condition monitoring, like those derived from SCADA-based fault diagnostics or CNN-based anomaly detection Tautz-Weinert & Watson (2017); Rahimilarki et al. (2022) are often underutilized when updating maintenance plans or long-term maintenance strategies, which normally make use of fixed failure rates to

model degradation. These examples highlight how cross-horizon information flow could improve responsiveness, resource planning, and overall lifecycle cost-efficiency.

Therefore, it is crucial to consider the interdependencies and interactions between the different decision-making layers when addressing the O&M optimization problem.

Table 2.10 presents an interconnection matrix illustrating how core variables such as maintenance plans, personnel availability, or vessel capacity, interact with different offshore wind O&M sub-problems. The matrix highlights whether each variable acts as an input, constraint, or output within models for maintenance strategy, spare parts management, fleet selection, task scheduling, and vessel routing. This synthesis can be useful for researchers building O&M models, as it helps identify variable dependencies and coordination requirements across hierarchical decision layers.

Therefore, there is a pressing need to develop integrative frameworks that link decisions across horizons, facilitating feedback loops, information flow, and dynamic adaptation. This not only enables a system-level optimization of O&M but also helps align short-term operations with long-term goals. Decisions related to O&M are not isolated events but rather dynamic processes that evolve over the lifespan of a wind farm. The decisions made in one decision-making layer can significantly impact the effectiveness and outcomes of decisions made in other layers. The conditions, performance, and demands of OWFs are subject to continual changes and unforeseen circumstances Shafiee (2015).

To address these challenges, it is of utmost importance to develop a mechanism or framework that can integrate and fuse different algorithms and strategies for more effective decision-making Ren et al. (2021). Such a mechanism would enable the seamless flow of information, insights, and data across different echelons, allowing for a holistic view of the O&M landscape. By integrating and fusing the various decision-making processes, operators can make more informed and comprehensive decisions that account for the interdependencies and interconnections between different aspects of O&M.

Developing such a mechanism or framework holds several key benefits. Firstly, it promotes synergy and coherence in decision-making, avoiding conflicts or suboptimal outcomes caused by isolated or disconnected strategies. Secondly, it facilitates knowledge sharing and cross-learning between different stakeholders and disciplines involved in O&M. This collaboration enhances the

Table 2.10: Interconnection matrix.

Problem	Maintenance plan	Age thresholds	Personnel availability	Personnel skill set	Availability of spare parts	Critical components	Maintenance tasks	Task completion time	Availability of vessels	Fleet mix	Vessel speed	Vessel capacity
Maintenance strategy	O	O	I	I	C	O	O	I	I	C	I	I
Spare parts management	I	C			O	I	I					
Fleet selection	I	C	O	O			I	C	O	O	O	O
Task scheduling	I	I	C	C	I	I	I	O	I	C	I	C
Vessel routing	I		C	C	I		I	C	I	O	C	C

Note: I = Input; O = Output; C = Constraint. Each letter indicates whether a variable serves as an input, output, or constraint in the corresponding optimization sub-problem.

collective understanding and expertise in managing OWF O&M. Lastly, it enables the identification of potential trade-offs, synergies, and opportunities that may arise from integrating different algorithms and strategies.

To achieve this, future research efforts should focus on developing advanced decision support systems that can integrate and fuse different algorithms, models, and strategies. These systems should be adaptable and flexible to accommodate the dynamic nature of O&M requirements. Additionally, they should leverage emerging technologies such as artificial intelligence, machine learning, and big data analytics to maintain horizontal coherence (within one horizon) and vertical consistency (across horizons).

In conclusion, rather than viewing each decision horizon as isolated, we advocate for a holistic modeling approach. By accounting for cross-layer feedbacks, constraints, and dependencies, such integrated systems can significantly enhance the robustness, agility, and economic performance of offshore wind O&M strategies.

2.6.3 Environmentally-inclusive decision-making

In the realm of OWF O&M, there exists a growing interest in integrating environmental concerns into decision-making frameworks. However, different schools of thought must be distinguished. *Nature-Inclusive Design (NID)* refers to approaches that enhance positive environmental impacts, e.g., through habitat creation or biodiversity-friendly structures. In contrast, eco-conscious or techno-environmental frameworks aim to quantify and minimize environmental harm, often through metrics like lifecycle emissions or ecological cost functions.

Several studies have started incorporating environmental aspects into O&M-related decision support. For instance, Garcia-Teruel et al. (2022) evaluate the Global Warming Potential (GWP) of floating OWFs and shows that the O&M phase can account for up to 49% of lifecycle GHG emissions, depending on fleet composition and maintenance strategy. This demonstrates that O&M choices significantly influence environmental performance. Similarly, Canet et al. (2023) propose a techno-environmental framework where wind turbine design and operation are optimized beyond purely economic metrics, incorporating carbon footprint explicitly into the objective function.

Despite this progress, these approaches still fall short of fully integrating ecological considerations or biodiversity objectives into maintenance planning. For example, vessel routing decisions or repair schedules rarely account for impacts on marine life or sensitive habitats. This is where the principles of Nature-Inclusive Design remain underexplored in the O&M domain.

A few studies hint at this potential. The work in Hernandez et al. (2021) highlight how vessel type influences emissions, suggesting environmental trade-offs in logistic planning. Similarly, the authors of Hall et al. (2022) argue that Environmental Impact Assessment (EIA) should not be treated as a static regulatory requirement, but rather as a dynamic decision-support tool employed throughout the project lifecycle, including the O&M phase.

To bridge this gap, future efforts should develop hybrid frameworks that combine quantitative environmental assessment with ecological sensitivity mapping, integrating tools such as lifecycle assessment (LCA) and real-time monitoring of ecological indicators. These could help optimize O&M not only for cost and reliability but also for biodiversity and environmental preservation.

Such integration would yield multiple benefits, including compliance with increasingly stringent environmental regulations, minimized disruption to marine ecosystems and habitats, and improved public perception and stakeholder acceptance, particularly in ecologically sensitive regions.

Ultimately, aligning O&M decision-making with both sustainability and biodiversity goals can significantly enhance the long-term viability and social license of offshore wind developments.

2.6.4 Lack of real-life data

One of the challenges in the research of the O&M of OWFs is the availability of actual real-life data. This is due to the reluctance of turbine manufacturers and wind farm owners to distribute such information. This obstacle has already been pointed out in articles such as Sinha & Steel (2015) in the context of failure prognostics, Gintautas & Sørensen (2017) in their research regarding weather window prediction and many others. The reluctance to share data may stem from competitive reasons, as the performance and reliability of a wind farm can be a key factor in attracting new customers and investors. Sharing this information may also expose potential vulnerabilities and weaknesses that could be exploited by competitors. Moreover, the data collected by the turbine manufacturers and wind farm owners is often considered proprietary information and can be subject to intellectual property laws and contractual obligations. As such, the sharing of this data may require legal agreements and negotiations, which can be time-consuming and costly. This lack of real-life data can limit the accuracy and validity of O&M simulations and modeling. Without access to actual data, researchers may have to rely on assumptions and estimations, which can introduce inaccuracies and bias into the analysis.

The limitation is not necessarily that the relevant operational data do not exist, but rather that they are fragmented across different stakeholders and are rarely available to academic researchers in an integrated form. Wind farm operators may have access to substantial parts of these datasets through their asset-management and maintenance-management systems, while other information may remain under the control of turbine manufacturers, O&M contractors, vessel operators, or condition-monitoring providers. Consequently, access to the combination of datasets required to calibrate and validate an integrated O&M decision-support model remains particularly difficult.

For condition and degradation modelling, the most valuable information includes SCADA measurements and alarm logs, dedicated condition-monitoring signals such as vibration and temperature measurements, inspection records, component age and failure histories, observed failure modes, and maintenance histories. Together, these data would enable the calibration and validation of both degradation models and the diagnostic or prognostic models used to infer component health. For the decision-making formulations considered later in this thesis, such data are particularly relevant for estimating how the hidden condition of a component evolves over time and how reliably this condition can be inferred from the information available to the maintenance planner.

A second group of data is required to represent maintenance execution and offshore logistics. Relevant sources include work orders and maintenance-management records containing intervention type, duration, technicians involved, spare parts consumed, and restoration outcome; vessel dispatch and transit logs; observed waiting and transfer times; measured and forecast metocean conditions; and records of unsuccessful or delayed offshore access. Spare-part inventories and procurement lead times would allow material availability to be represented dynamically, while component repair and replacement costs, vessel charter and mobilisation costs, and turbine production data would support

calibration of the economic objective. Combining these health, maintenance, logistics, accessibility, and cost datasets would therefore enable substantially stronger validation of O&M decision-support models. Synthetic datasets remain useful for controlled methodological comparisons, as adopted in this thesis, but cannot substitute for project-specific calibration and out-of-sample validation using real operational data.

2.7 Conclusions

This chapter addressed **Research Question Q1**, namely: *What are the state-of-the-art, state-of-practice, and research gaps in offshore wind O&M decision support, particularly regarding the integration of accessibility, logistics, and condition-informed planning?* By reviewing the literature through a strategic-tactical-operational decision-making framework, the chapter provided a structured synthesis of how offshore wind O&M decisions are currently modelled, what assumptions dominate existing methods, and where important limitations remain.

Across the three decision layers, the review showed that the literature has produced a broad range of valuable modeling approaches. At the strategic level, simulation and optimization models support long-term evaluation of cost, reliability, and maintenance policy design. At the tactical level, the literature addresses questions such as fleet composition, spare-parts policies, and resource allocation. At the operational level, routing, scheduling, diagnostics, and prognostics have received increasing attention, particularly in relation to short-term intervention planning under weather and resource constraints. Taken together, these contributions establish a strong basis for model-based decision support in offshore wind O&M.

At the same time, the review identified four recurring gaps that are particularly relevant for this thesis. First, prognostic and condition-informed maintenance is increasingly discussed in the literature, but its integration into actual planning and scheduling models remains limited. Second, many models treat accessibility, logistics, and asset health information only partially or in isolation, rather than as tightly coupled elements of the same decision problem. Third, the connection across strategic, tactical, and operational decision levels remains weak, with few approaches providing a consistent cross-horizon decision-support structure. Fourth, the development and validation of advanced O&M decision models is still constrained by the scarcity of real operational and prognostic data.

These findings define the methodological role of the remainder of the thesis. Chapter 3, building on the findings of this literature review, establishes a shared modeling framework, including metocean-driven accessibility, logistics assumptions, and prognostic information. Sharing this common framework, Chapter 4 and Chapter 6 investigate how offshore wind O&M planning can be formulated as an optimization and a learning problem, respectively, explicitly incorporating accessibility, logistics resources, and condition-informed maintenance decisions.

Chapter 3

OWF Characteristics & Dynamics

This chapter addresses Research Question Q2: *How to model the dynamics of offshore wind farm O&M to support maintenance decision-making?* To answer this question, the chapter defines which aspects and which components should be taken into account when modeling the dynamics of an OWF. Its purpose is to establish a common modeling environment in which optimization-based and learning-based approaches can be evaluated on a like-for-like basis. In this way, differences observed in later chapters can be attributed primarily to the decision-making paradigm itself, rather than to inconsistencies in case-study assumptions, data inputs, or performance metrics.

Figure 3.1 provides a schematic overview of the offshore wind farm O&M system represented in this chapter. The figure highlights the main interacting elements that define the modeling environment: the wind farm layout, the onshore O&M base, the access assets used to reach the farm, the metocean conditions that determine accessibility, the critical turbine components whose degradation generates maintenance demand, the uncertain RUL information available to the decision models, and the maintenance actions that affect both component condition and production losses. These interactions motivate the modeling requirements introduced below and provide the basis for the mathematical definitions used in the remainder of the chapter.

This chapter is organized as follows. Section 3.1 introduces the rationale and design requirements of the modeling environment shown in Figure 3.1. Section 3.2 defines the fixed-bottom offshore wind farm layout, turbine power-production model, and spatial structure of the offshore site. Section 3.3 introduces the critical turbine components whose degradation generates maintenance demand, while Section 3.4 and Section 3.5 define the transportation assets, metocean conditions, and accessibility constraints governing offshore operability. Section 3.6 presents the synthetic RUL prediction generator used to provide uncertain prognostic information to the decision models. Section 3.7 defines the maintenance-action space, including the restorative effects, resource requirements, and state-update logic associated with minor repairs, major repairs, and replacements. Together, these sections define the offshore wind O&M environment in which maintenance decisions are coordinated and applied. The decisions themselves may follow predefined maintenance strategies or be selected dynamically through the optimization- and learning-based approaches developed in Chapters 4 and 6. Finally, Section 3.8 summarizes the resulting dataset structure, costs, KPIs, and modeling interfaces used consistently throughout the remainder of the thesis.

The modelling environment presented in this chapter builds on the framework introduced in Marco Borsotti, Rudy R. Negenborn, and Xiaoli Jiang: “Model Predictive Control Framework




Figure 3.1: Schematic overview of the offshore wind O&M environment.

for Optimizing Offshore Wind O&M”, in *Advances in Maritime Technology and Engineering*, CRC Press, pp. 533–546, 2024, presented at the 7th International Conference on Maritime Technology and Engineering (MARTECH 2024), Lisbon, Portugal, 14–16 May 2024. The present chapter consolidates the modelling assumptions, data-generation procedures, accessibility logic, maintenance-action taxonomy, and benchmark outputs used by the optimization- and learning-based decision-support models developed in the following chapters.

3.1 Requirements of the modeling environment

Chapter 2 showed that offshore wind O&M decision-support models often treat asset degradation, weather-driven accessibility, marine logistics, and maintenance decision-making in partial or disconnected ways. Some approaches focus on maintenance timing while simplifying logistics to generic capacity constraints, whereas others represent vessels and routing in greater detail but rely on simplified asset-health information or static failure assumptions. This fragmentation makes it difficult to compare methods fairly, since the modeling environment itself can dominate results such as downtime, total O&M cost, vessel utilisation, or the number of failures. A controlled and standardised environment is therefore needed before alternative decision-support methods can be meaningfully compared.

In this sense, the chapter provides the benchmark layer of the thesis. It specifies the physical, operational, environmental, and prognostic assumptions that remain common across the optimization-based and learning-based workstreams. The environment consolidates the wind farm configuration, the set of critical components, the available transportation and access assets, the generation of synthetic metocean-driven accessibility data, the generation of synthetic RUL predictions, and the maintenance-action framework used to update component condition over time. Together, these elements define a standardised environment through which all subsequent models interact with the same underlying offshore wind O&M system.

Based on the findings in Chapter 2, the environment is designed to satisfy five requirements:

- R1.** Represent the dynamics of a fixed-bottom offshore wind farm at sufficient detail for maintenance planning. The environment must capture key elements such as the wind farm layout, turbine characteristics, and spatial relations that determine travel distances, production losses, and the basic structure of the maintenance planning problem. A core requirement is modularity: the benchmark must support a variable number of turbines while preserving a consistent and reproducible spatial structure, so that optimization- and learning-based methods can be tested across different problem scales.
- R2.** Include critical components whose degradation drives maintenance demand. The model must explicitly represent the turbine subsystems that generate maintenance needs, so that interventions are linked to meaningful degradation and failure processes.
- R3.** Model weather-driven accessibility and logistics feasibility. The environment must represent offshore access assets, their operability limits, and the metocean conditions that determine whether maintenance actions are feasible in practice.
- R4.** Provide uncertain prognostic information. Since the thesis evaluates methods that use condition and prognostic information, the environment must generate RUL estimates and associated uncertainty in a way that is realistic enough to support meaningful decision making, while remaining controlled enough to enable repeatable experiments.
- R5.** Define inputs, outputs, costs, and KPIs. The environment must define a common information and evaluation interface so that the optimization model and the learning framework are assessed under the same assumptions, state variables, and performance indicators.

Figure 3.2 summarises these requirements and shows how each one motivates a corresponding modeling block in the remainder of the chapter.

The remainder of the chapter follows this logic. Section 3.2 defines the fixed-bottom offshore wind farm configuration and spatial layout, thereby addressing Requirement **R1**. Section 3.3 introduces the critical component set based on a Failure Modes, Effects and Criticality Analysis, addressing Requirement **R2**. Section 3.4 describes the transportation and access assets together with their operability limits, while Section 3.5 presents the historical weather source and the synthetic metocean operability generator; together, these sections address Requirement **R3**. Section 3.6 defines the synthetic RUL prediction generator used to emulate prognostic information under uncertainty, addressing Requirement **R4**. Section 3.7 then specifies the maintenance action set, restorative assumptions, resource requirements, and state-update logic. Finally, Section 3.8 summarises the




Figure 3.2: Design requirements of the offshore wind O&M environment

resulting data structure and outputs, showing how the environment provides a common interface for the MILP and DRL models and thereby addresses Requirement **R5**.

3.2 Wind farm configuration and spatial layout

This section defines the wind farm representation used in the O&M environment. The formulation is inspired by reference-case approaches for offshore wind O&M simulation, where simplified but well-documented case studies are used to compare models under controlled assumptions (Dinwoodie et al., 2015). The objective is therefore not to reproduce the exact geometry of a specific commercial wind farm, but to define a transparent, modular, and reproducible spatial representation that can be used consistently by the optimization- and learning-based models developed in later chapters.

3.2.1 Power production model

An OWF is represented as a set of N_{WT} wind turbines, each with rated power P_r [MW]. The total installed capacity of the wind farm is therefore

$$P_{\text{farm}} = N_{\text{WT}} P_r. \quad (3.1)$$

This aggregation is commonly used in wind-farm modeling to describe the nominal capacity of a project. The wind farm is assumed to be served by a single onshore Operations & Maintenance (O&M) base in the Netherlands, located at a one-way nautical distance d [km] from the offshore site. This abstraction is consistent with offshore wind O&M models in which travel between an onshore base, offshore turbines, and/or offshore service locations is a central driver of maintenance time and cost (Dinwoodie et al., 2015; Irawan et al., 2022).

The turbine layout is defined as a regular grid, i.e., a rectangular array, to avoid confounding benchmarking results with site-specific cabling design, wake-optimized micro-siting, or irregular concession-boundary effects. Regularised layouts are a standard simplification in benchmark and simulation studies because they provide a reproducible spatial structure while retaining the essential information needed for logistics calculations, namely turbine positions and inter-turbine distances. Let turbine i have planar coordinates (x_i, y_i) , with uniform inter-turbine spacing Δx and Δy . Distances between turbines are computed in the two-dimensional plane and subsequently translated into travel times using vessel-specific speeds, as commonly done in maintenance routing and location-routing formulations for offshore wind farms (Irawan et al., 2022).

Each turbine is modelled as a utility-scale unit characterised by five main power-performance parameters: cut-in wind speed v_{in} [m/s], rated wind speed v_r [m/s], cut-out wind speed v_{out} [m/s], rotor diameter D [m], and hub height H [m]. The cut-in wind speed v_{in} is the minimum wind speed at which the turbine starts producing electrical power. Below this threshold, the aerodynamic power available in the wind is insufficient to operate the turbine efficiently, and the power output is assumed to be zero. The rated wind speed v_r is the wind speed at which the turbine first reaches its rated power P_r ; above this value, the turbine control system limits the power output to avoid overloading the generator and drivetrain. The cut-out wind speed v_{out} is the wind speed above which

the turbine is shut down for safety, so that structural loads and component stresses remain within acceptable limits.

The rotor diameter D determines the swept area of the rotor and therefore the amount of wind energy that can be captured by the turbine. Larger rotors intercept a larger air mass and are generally associated with higher energy yield, but they also increase the scale and complexity of maintenance interventions, especially for blade-related tasks. The hub height H defines the elevation at which the wind speed is evaluated. This is important because wind speed generally varies with height, and the wind speed at hub height is the relevant input for the turbine power curve. Together, these variables determine the relationship between the metocean time series and the available power production of the turbine.

In the O&M environment, these parameters are important for two reasons. First, they determine the available production that is lost when a turbine is unavailable due to failure, waiting time, or maintenance. Second, they provide a consistent way to translate wind-speed conditions into economic consequences through the downtime cost model. This representation follows the conventional use of wind turbine power curves in power-performance and annual-energy-production calculations (Commission, 2022; Wan et al., 2010; Carrillo et al., 2013).

Hourly energy production is first computed at the turbine level from the single-turbine power curve $P(v)$ and the wind-speed time series. The function $P(v)$ maps a hub-height wind speed v to the corresponding power output of one turbine. For hour h , let v_h [m/s] denote the wind speed at hub height and let $P(v_h)$ denote the corresponding single-turbine power output. The available energy production of one turbine is

$$E_{i,h}^{\text{avail}} = P(v_h) \Delta h, \quad (3.2)$$

where $E_{i,h}^{\text{avail}}$ [MWh] is the available energy production of turbine i during hour h , $P(v_h)$ [MW] is the corresponding single-turbine power output, and $\Delta h = 1$ h is the temporal resolution of the metocean and production time series. In this benchmark, all turbines are exposed to the same farm-level wind-speed time series, so $E_{i,h}^{\text{avail}}$ is identical for all turbines before availability losses are applied.

Farm-level available production is then obtained by summing the available production over all turbines that are operational during hour h :

$$E_h^{\text{farm}} = \sum_{i=1}^{N_{\text{WT}}} a_{i,h} E_{i,h}^{\text{avail}} = \sum_{i=1}^{N_{\text{WT}}} a_{i,h} P(v_h) \Delta h, \quad (3.3)$$

where E_h^{farm} [MWh] is the realised farm-level energy production during hour h , and $a_{i,h} \in \{0, 1\}$ is the availability indicator of turbine i . If turbine i is operational, $a_{i,h} = 1$; if it is unavailable due to failure, waiting time, or maintenance, $a_{i,h} = 0$. The corresponding lost production is therefore

$$E_h^{\text{lost}} = \sum_{i=1}^{N_{\text{WT}}} (1 - a_{i,h}) E_{i,h}^{\text{avail}}. \quad (3.4)$$

The O&M decision model therefore evaluates farm-level consequences by aggregating turbine-level unavailability over time.

3.2.2 Modular layout generation

A core design requirement of the reference case study is *modularity*: the benchmark must support a variable number of turbines while preserving a consistent and reproducible spatial structure. This enables systematic scalability studies, for example $N_{WT} \in \{10, 25, 50, 100\}$, without changing the downstream model interface. Given an input dataset specifying the number of turbines N_{WT} , the layout is generated procedurally according to a deterministic grid-based rule:

1. Read N_{WT} from the dataset metadata.
2. Compute the grid dimension as

$$n = \lceil \sqrt{N_{WT}} \rceil, \quad (3.5)$$

where n denotes the number of columns and rows of the enclosing square grid. Turbines are placed row-wise on an $n \times n$ grid. If $n^2 > N_{WT}$, only the first N_{WT} grid points are populated, so that the remaining cells are left empty and the last row may be only partially filled. Equation (3.5) is a benchmark-design rule introduced in this thesis to obtain a compact, reproducible layout for any selected value of N_{WT} .

3. Apply a user-defined spacing s to convert grid indices into Cartesian coordinates. The coordinate of turbine i placed at row r_i and column c_i is

$$(x_i, y_i) = (c_i s, r_i s). \quad (3.6)$$

The point $(x_0, y_0) = (0, 0)$ is defined as the wind-farm entry point, i.e., the reference origin used for marine logistics calculations.

Equations (3.5) and (3.6) are not meant to prescribe an optimal wind farm layout. Instead, they define a reproducible synthetic benchmark. This distinction is important because wind farm layout optimization is a separate design problem, typically involving wake losses, bathymetry, electrical layout, exclusion zones, and concession constraints. Here, these effects are deliberately excluded so that the benchmark isolates the O&M decision-making problem rather than the wind-farm design problem.

The grid-based structure is intentionally simple: it offers a transparent mapping between N_{WT} and spatial extent, and avoids introducing wind-farm-specific geometric bias into the benchmarking layer. Figure 3.3 illustrates three example layouts generated according to the same procedural rule for $N_{WT} \in \{10, 25, 50\}$, shown in Fig. 3.3a–3.3c, respectively.

The layout generator computes and stores a pairwise distance matrix:

$$D_{ij} = \left\| \begin{bmatrix} x_i \\ y_i \end{bmatrix} - \begin{bmatrix} x_j \\ y_j \end{bmatrix} \right\|_2 = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}, \quad (3.7)$$

for all turbine pairs (i, j) . The use of spatial distances between turbines is standard in offshore wind maintenance routing formulations, where turbine locations determine feasible routes, travel times, and vessel utilisation costs (Irawan et al., 2022). The distance matrix is used as a lookup table by routing heuristics and optimization formulations, avoiding repeated geometric calculations at runtime. The same precomputed matrix $\mathbf{D}$ is shared across all models, ensuring that differences




Figure 3.3: Example regular grid layouts for different wind farm sizes.

in routing and travel costs arise from decision logic rather than from differences in numerical implementation.

When travel time is required, distances are converted using vessel-specific speeds:

$$\tau_{ij}^v = \frac{D_{ij}}{u_v}, \quad (3.8)$$

where τ_{ij}^v is the travel time between turbines i and j for vessel type v , and u_v is the corresponding vessel speed. This distance–speed conversion is a standard representation in vessel routing and maintenance logistics models, where sailing times are determined by the distance between service locations and the speed of the selected vessel (Irawan et al., 2022).

This module is designed to be directly replaceable with real wind-farm coordinates when available. In that case, the grid-generation and spacing rules in Equations (3.5)–(3.6) are replaced by the actual turbine coordinates of the specific OWF, while the distance matrix $\mathbf{D}$ is computed identically using Equation (3.7). This preserves the downstream interface to all models and ensures that the benchmarking framework remains compatible with both synthetic and real case studies.

3.3 Critical components set

The previous section defined the wind farm as a set of N_{WT} turbines, indexed by i , and introduced the turbine-level availability indicator $a_{i,h}$, which determines whether turbine i contributes to farm-level production during hour h . The present section defines what causes this availability indicator to change from an operational to a non-operational state. In particular, each turbine is decomposed into a set of critical components, indexed by c , whose degradation and failure generate maintenance demand. The component set therefore provides the link between the turbine-level production model in Eqs. (3.2)–(3.4) and the component-level RUL, failure, and maintenance models introduced in the following sections.

A failure of component c on turbine i implies that the corresponding turbine is unavailable until a restorative maintenance action is completed. In terms of the production model, this means

that the turbine availability indicator becomes $a_{i,h} = 0$ during the failure, waiting, and repair period, causing the available energy $E_{i,h}^{\text{avail}}$ to contribute to lost production rather than realised farm output. Defining the critical component set is therefore necessary before introducing the RUL generator, maintenance actions, and cost model, because it specifies the physical subsystems to which degradation trajectories, failure probabilities, repair actions, vessel requirements, and downtime consequences are attached.

A wind turbine is represented here as a series system of critical subsystems, meaning that failure of any one of the selected critical components is assumed to render the turbine unavailable until a restorative maintenance action is completed. This abstraction is widely adopted in offshore wind O&M modeling because it provides a tractable way to link component-level degradation and failure events to turbine-level downtime and production loss. Within this benchmark, the selected critical components are chosen to capture both high-frequency maintenance demand and high-impact failures with strong logistical implications. Figure 3.4 provides a simplified schematic of the wind turbine and nacelle, highlighting the main structural elements and the internal drivetrain configuration used to position the critical components considered in this thesis.

To define a tractable yet representative component set for benchmarking, a Failure Modes, Effects and Criticality Analysis (FMECA) is conducted following standard FMEA practice (Commission, 2018). Starting from a broad turbine bill of materials, candidate subsystems are screened using the four criteria summarized in Table 3.1: severity, occurrence, detectability, and logistics criticality. Each criterion is rated on a five-level scale, where increasing scores indicate a higher contribution to overall maintenance criticality; for detectability, a higher score denotes greater




Figure 3.4: Wind turbine schematic showing the main structural and drivetrain components. The numbered elements are: (1) tower, (2) nacelle, (3) blades and rotor, (4) generator, (5) gearbox, (6) main bearing, (7) main shaft / drivetrain shaft, and (8) hub and pitch system assembly.

Table 3.1: FMECA scoring rubric used to screen turbine subsystems for inclusion in the benchmark component set.

Criterion	1	2	3	4	5
Severity	Negligible production effect; no safety exposure; no secondary damage	Limited production loss; local effect only	Noticeable production loss; moderate local damage	Major production loss; possible safety exposure or secondary damage	Turbine-level outage or severe safety / cascading damage consequences
Occurrence	Rare failure mode; low empirical hazard	Occasional failure mode	Moderate recurring failure frequency	Frequent failure mode	Very frequent / high-hazard failure mode
Detectability	Easy early detection with standard monitoring	Generally detectable with routine CM signals	Detectable only through trend analysis or combined indicators	Weak / intermittent signatures; limited early warning	Difficult to detect before failure; poor CM observability
Logistics criticality	Short intervention; minimal tools; no strong vessel dependence	Simple planned intervention; low logistic burden	Moderate repair duration; planned crew/vessel coordination needed	Long repair; large spare parts; narrow access window	Heavy-lift or highly vessel-dependent repair; long downtime exposure

difficulty of early fault identification. After assigning subsystem scores, a composite criticality index can be computed through a weighted additive form,

$$CI_i = w_S S_i + w_O O_i + w_D D_i + w_L L_i, \quad (3.9)$$

Here, CI_i is the criticality index of subsystem i , while S_i , O_i , D_i , and L_i denote severity, occurrence, detectability, and logistical impact. Severity measures the consequence of failure, occurrence its expected frequency, detectability the likelihood that degradation can be detected before functional failure, and logistical impact the offshore-specific difficulty of executing the required intervention. In this thesis, the additional logistical-impact term is included because components are not only critical when they fail often or severely, but also when their maintenance requires scarce vessels, long access windows, specialised equipment, or extended downtime. The additive form is used because it provides a transparent screening metric in which each dimension contributes explicitly to the final score. The weights w_S , w_O , w_D , and w_L allow the index to be adapted to the purpose of the analysis. In this thesis, equal weights are adopted because no stakeholder-specific preference was available that would justify prioritising one criticality dimension over another. The resulting index therefore gives equal importance to failure consequences, failure frequency, detectability, and offshore logistical burden.

The candidate set is based on the offshore wind turbine subsystem categories analysed by Carroll et al. (2016). Their empirical study covers approximately 350 offshore wind turbines and provides subsystem-level failure frequencies, repair categories, material costs, active repair durations, and

technician requirements. The occurrence score is mapped directly from the total annual subsystem failure rate, λ_i , using

$$O_i = \begin{cases} 1, & \lambda_i < 0.15, \\ 2, & 0.15 \leq \lambda_i < 0.25, \\ 3, & 0.25 \leq \lambda_i < 0.40, \\ 4, & 0.40 \leq \lambda_i < 0.80, \\ 5, & \lambda_i \geq 0.80, \end{cases} \quad (3.10)$$

where λ_i is expressed in failures per turbine per year. Severity and logistics-criticality scores are assigned by matching the reported repair category, material cost, active repair duration, technician requirement, and heavy-lift implications to the rubric in Table 3.1. Detectability is assessed from the relative availability and maturity of condition-monitoring methods reported in wind turbine reliability and FMECA studies (Pérez et al., 2013; Scheu et al., 2019).

The numerical ratings in Table 3.2 are therefore not values reported verbatim by Carroll et al. (2016); rather, they constitute a transparent mapping of the published empirical evidence to the common five-level rubric adopted in this thesis. The total failure rates are obtained by summing the minor-repair, major-repair, major-replacement, and unclassified-cost failure rates reported for each subsystem.

Table 3.2: Literature-derived FMECA scores used to select the critical turbine subsystems. Failure rates and repair characteristics are derived from Carroll et al. (2016).

Subsystem	λ_i [turbine ⁻¹ year ⁻¹]	S_i	O_i	D_i	L_i	CI_i	Retained
Generator	0.999	5	5	2	5	4.25	Yes
Blades	0.520	5	4	3	5	4.25	Yes
Gearbox	0.633	5	4	2	5	4.00	Yes
Main bearing	0.227	5	2	3	5	3.75	Yes
Pitch/hydraulic system	1.076	4	5	3	3	3.75	Yes
Contactors/circuit breaker/relay	0.430	4	4	2	4	3.50	No
Hub	0.235	5	2	2	5	3.50	No
Other components ^a	1.005	2	5	4	2	3.25	No
Yaw system	0.189	4	2	3	4	3.25	No
Tower/foundation	0.185	5	2	3	3	3.25	No
Electrical components	0.435	3	4	2	3	3.00	No
Controls	0.428	3	4	2	3	3.00	No
Power supply/converter	0.180	4	2	2	4	3.00	No
Sensors	0.346	2	3	4	2	2.75	No
Pumps/motors	0.346	3	3	3	2	2.75	No
Transformer	0.065	5	1	2	3	2.75	No
Grease/oil/cooling liquids	0.471	2	4	2	2	2.50	No
Safety system	0.392	2	3	2	2	2.25	No
Heaters/coolers	0.213	2	2	2	2	2.00	No
Service items	0.125	1	1	1	1	1.00	No

^aThe “other components” category in Carroll et al. (2016) aggregates auxiliary items such as lifts, ladders, hatches, door seals, and nacelle seals. Although these events occur frequently, they are predominantly minor interventions and do not constitute a single physically homogeneous degradable subsystem to which a component-specific lifetime and RUL model can be assigned.

The five highest-ranking physically distinct subsystems are retained, resulting in an effective inclusion cut-off of

$$CI_i \geq 3.75.$$

The resulting component set is

$$C = \{\text{blades, main bearing, gearbox, generator, pitch system}\}.$$

The number of retained components is limited to five to balance physical representativeness against the dimensionality of the subsequent optimization and learning problems.

The ranking captures two complementary sources of O&M criticality. First, the pitch/hydraulic system and generator generate frequent maintenance demand. Carroll et al. (2016) identify the pitch/hydraulic system as the largest individual contributor to the observed offshore failure rate, while the generator is also a major contributor to both minor and major interventions. Second, blades, the gearbox, the generator, and the main bearing represent lower-frequency but high-consequence interventions associated with long repair durations, large crews, expensive replacement parts, specialised vessels, and prolonged downtime. In particular, generator and gearbox failures account for the vast majority of the major replacements observed by Carroll et al. (2016).

The hub also exhibits high repair severity and logistical burden, but its lower occurrence and relatively greater detectability result in a score below the inclusion cut-off. The aggregated “other components” category has a high occurrence rate but consists predominantly of minor auxiliary failures. It is therefore less suitable for the component-specific degradation and prognostic models developed in this thesis. The selected set consequently captures both frequent faults that drive technician and vessel activity and high-impact failures that drive heavy-vessel mobilisation and prolonged production loss.

3.4 Transportation and access assets

The previous section defined the critical turbine subsystems that generate maintenance demand in the O&M environment. Once a component requires inspection, repair, or replacement, the feasibility and cost of the intervention depend not only on the component itself, but also on the offshore resources needed to reach the turbine and execute the task. Transportation and access assets therefore form the link between component-level maintenance needs and system-level O&M performance.

This section defines the access assets considered in the benchmark and explains how their operational capabilities constrain maintenance execution. Each transport mode is characterised by its role in offshore O&M, its typical use case, and its operability limits with respect to wind speed and significant wave height. These limits are subsequently used in Section 3.5 to translate metocean conditions into binary accessibility states. In this way, the transportation module provides the logistical interface between maintenance actions, weather windows, travel feasibility, downtime, and cost.

Marine logistics is a primary driver of offshore wind O&M performance because it determines (i) whether technicians and equipment can physically reach the asset, and (ii) how long corrective and preventive tasks remain pending under adverse metocean conditions. For benchmarking the optimization- and learning-based decision-support models developed in this thesis, the reference

case study includes a modular *access fleet* layer: each transportation mode is represented by (a) operability limits (based on weather conditions and seakeeping abilities of the transportation mode selected) and (b) capacity constraints (crew/load).

Common means of transportation for accessing the plant and carrying out maintenance activities are crew transfer vessels (Figure 3.5a), service operation vessels (Figure 3.5b), jack-up vessels (Figure 3.5c) and helicopters (Figure 3.5d), which are described in the following subsections.

3.4.1 Crew transfer vessel

CTVs are the standard nearshore solution for daily technician transfers and small-component logistics. They typically depart from an onshore base or a mother vessel, execute one or multiple turbine visits, and return to port within the same day. Transfer is performed via the turbine boat landing/ladder, which creates a relatively strict operability threshold in significant wave height. A commonly adopted limit in O&M models is $H_s \leq 1.5$ m (Hu, 2019; Donnelly et al., 2024). For sites located significantly far from shore or for longer operations, an offshore accommodation strategy such as a service operation vessel is usually more efficient (Hu, 2019; Neves-Moreira et al., 2021).



(a) Crew Transfer Vessel.



(b) Service Operation Vessel.



(c) Jack-up Vessel.



(d) Helicopter transfer for nacelle access.

Figure 3.5: Transportation and access assets.

3.4.2 Service operation vessel

SOVs constitute an offshore-based access strategy in which technicians are accommodated onboard and turbine access is normally achieved using a motion-compensated walk-to-work gangway and, where applicable, daughter craft. This reduces repeated port transit and can increase accessibility for far-from-shore wind farms (Hu, 2019; Donnelly et al., 2024). The limiting significant wave height is system-specific rather than universal, because it depends on vessel response, gangway design and location, dynamic-positioning capability, transfer procedures, and operator restrictions. Hu (2019) reports limiting values of approximately $H_s = 2.5\text{--}4.5\text{ m}$ across different motion-compensated access systems, while values between 2.5 m and 3.5 m are commonly applied to SOV-based access in O&M studies. Donnelly et al. (2024) use $H_s = 2.5\text{ m}$ as a baseline SOV operability limit and $H_s = 3.5\text{ m}$ as a higher-accessibility case. These values represent the operability limit of the access system.

3.4.3 Jack-up vessel

Jack-up vessels are required for heavy-lift maintenance tasks such as major component exchange. Operability is driven by safe positioning, jacking, and hoisting constraints. Typical modeling values are $H_s \leq 2.0\text{ m}$ during positioning and $v_{\text{wind}} \leq 10\text{ m/s}$ during crane operations (Donnelly et al., 2024). Mobilisation and demobilisation costs dominate the economics of this vessel type.

3.4.4 Helicopter

Helicopters are used for small teams or urgent inspections, transferring personnel to nacelle platforms or offshore substations. They bypass wave restrictions entirely but are limited by wind and visibility. The typical wind-speed limit for safe transfer is around 20 m/s, while nacelle work often requires lower limits (e.g. 12 m/s) (Hu, 2019). Load restrictions confine this mode to light-maintenance and inspection tasks.

Using the operability limits defined in Table 3.3, each hourly record of significant wave height and wind speed can be converted into an accessibility indicator for each transportation mode.

Table 3.3: Representative operability limits for transportation and intervention assets used to map hourly metocean conditions into accessibility windows. Values compiled from Hu (2019); Donnelly et al. (2024).

Asset	Primary role	H_s limit (m)	Wind limit (m/s)	Notes
CTV	Technician transfer (port-based, daily)	1.5	12	Boat-landing transfer; limited payload and crew capacity.
SOV	Technician transfer + offshore accommodation	3.5	17	Offshore-based; gangway limit depends on design.
Jack-up vessel	Major component exchange (crane)	2.0	10	H_s for positioning; wind for hoisting operations.
Helicopter	Technician transfer (nacelle/helideck)	–	20	Not constrained by H_s ; wind and visibility dependent.

These indicators define whether an operation can be initiated or continue at a given time step and are therefore a key exogenous input to both the MILP optimization and the DRL environment. This mapping ensures that both models are benchmarked under identical environmental constraints.

3.4.5 Spare parts management

The access assets described above determine whether technicians, tools, and equipment can be transported to the offshore site. A further logistical requirement concerns the availability of the specific spare parts needed to complete the selected maintenance action. Beyond the availability and operability of access assets, the execution of a maintenance action also depends on the availability of the required spare parts. Spare parts management is therefore an important element of offshore wind O&M logistics, since missing components or consumables can delay otherwise feasible interventions and increase turbine downtime. In the present modelling environment, however, spare parts management is not treated as an inventory optimisation problem. Instead, spare-part availability is represented as an operational feasibility constraint: a maintenance action can be undertaken only if the spare part required for the selected component and intervention type is available at the time of execution. In this sense, lack of spare parts is modelled analogously to a limitation in accessibility, since it can prevent a maintenance action from being performed even when the metocean conditions and access assets would otherwise allow it.

3.5 Weather data and synthetic metocean dataset generation

The previous sections defined the physical wind farm layout, the critical component set, and the transportation assets available for offshore interventions. However, whether these assets can actually execute maintenance at a given time depends on the metocean conditions at the site. Wind speed and significant wave height determine the accessibility of different transport modes and therefore shape the waiting times, downtime, and feasibility of maintenance actions.

This section defines the weather and accessibility module used in the O&M environment; the objective is to generate realistic long-horizon accessibility signals for each transport mode. To this end, historical wind and wave data are first converted into vessel-specific binary operability states. These empirical state sequences are then used to estimate month-dependent Markov transition matrices, which are finally sampled to create synthetic hourly accessibility datasets. The resulting data provide a common metocean-driven feasibility input for the MILP and DRL models developed in later chapters.

3.5.1 Historical metocean data source

Long-term weather measurements are essential to evaluate maintenance strategies under realistic variability in accessibility and waiting times. However, obtaining long time series for currently operational Dutch offshore wind farms can be challenging due to limited public release. This thesis uses the long-term hindcast database compiled by Li et al. (2015) for five European offshore sites. Their dataset is based on simultaneous hourly wind and wave hindcast data from 2001 to 2010 and provides joint information on mean wind speed at 10 m height and significant wave height, among

other parameters. This multi-year coverage enables robust characterization of seasonal patterns and inter-annual variability, which is critical for assessing O&M decision models. In this work, only two parameters are retained as exogenous drivers of marine operability: significant wave height (H_s) and wind speed at 10 m (U_{10}), consistent with the key variables required to determine access feasibility and (where relevant) to approximate production conditions. In order to obtain many long time series of realistic metocean-driven accessibility conditions for benchmarking purposes, we implemented a synthetic operability generator that preserves the empirical temporal structure of accessibility through a month-dependent Markov process. Rather than generating continuous wind and wave time series, the generator produces vessel-specific hourly operability states (accessible vs. non-accessible) derived from historical wind speed and significant wave height records. This directly provides the binary accessibility indicators used by both the MILP and DRL models.

3.5.2 Vessel-specific operability states

Hourly historical weather observations are read and the following information is extracted:

- `datetime` $\rightarrow h \in \{1, \dots, T\}$ at hourly resolution,
- `month(h)` $\in \{1, \dots, 12\}$ extracted from the timestamp,
- wind speed $U(h)$ at hour h ,
- significant wave height $H_s(h)$ at hour h .

Considering the following vessel set:

$$\mathcal{V} = \{\text{CTV, SOV, Jack-up, Helicopter}\}.$$

Each vessel's operational limits are read from Table 3.3 containing: maximum admissible wind speed per vessel type ($\bar{U}_v$), and maximum admissible significant wave height per vessel type ($\bar{H}_{s,v}$).

For each hour h and vessel type v , an operability indicator is computed by thresholding wind speed and wave height:

$$O_v(h) = \mathbb{I}[U(h) \leq \bar{U}_v \wedge H_s(h) \leq \bar{H}_{s,v}], \quad (3.11)$$

where $O_v(h) = 1$ denotes *operational*, and $O_v(h) = 0$ denotes *non-operational*.

This mapping converts continuous metocean variability into discrete operational feasibility states consistent with how access constraints enter both optimization and sequential decision formulations.

3.5.3 Month-dependent Markov-chain estimation

To preserve seasonal differences in persistence and switching behaviour, a separate two-state Markov chain is estimated for each month and vessel type. Let $S_v(h) \in \{0, 1\}$ denote the discrete operability state for vessel v at hour h , and let $m(h)$ denote the month associated with hour h . For each vessel v and month m , we estimate a transition matrix:

$$\mathbf{P}_v^{(m)} = \begin{bmatrix} p_{v,00}^{(m)} & p_{v,01}^{(m)} \\ p_{v,10}^{(m)} & p_{v,11}^{(m)} \end{bmatrix}, \quad p_{v,ij}^{(m)} = \mathbb{P}(S_v(h+1) = j \mid S_v(h) = i, m(h) = m). \quad (3.12)$$

The transition probabilities are obtained by maximum likelihood estimation (MLE) from empirical transition counts within each month:

$$p_{v,ij}^{(m)} = \frac{N_{v,ij}^{(m)}}{\sum_{j'} N_{v,ij'}^{(m)}}, \quad (3.13)$$

where $N_{v,ij}^{(m)}$ is the number of observed $i \rightarrow j$ transitions in month m for vessel v .

3.5.4 Synthetic operability series generation

Using the month-dependent transition matrices, we can thus generate synthetic hourly operability databases spanning any Y years.

$$T_{\text{syn}} = Y \cdot 365 \cdot 24 \text{ hours.}$$

For each vessel type v , the synthetic series is generated sequentially:

1. Initialise the state $S_v(1)$ by sampling from the empirical state distribution of the corresponding calendar month (or by fixing an initial state).
2. For each subsequent hour $h = 1, \dots, T_{\text{syn}} - 1$:

$$S_v(h+1) \sim \text{Categorical} \left(\mathbf{P}_v^{(m(h))} [S_v(h), :] \right), \quad (3.14)$$

where $\mathbf{P}_v^{(m(h))} [S_v(h), :]$ denotes the row of transition probabilities conditioned on the current state and month.

This mechanism reproduces month-specific persistence of access/no-access regimes while remaining computationally lightweight and fully reproducible under fixed random seeds.

The output is stored as a time-indexed dataframe per vessel type:

$$df_v = \{datetime(h), S_v(h)\}_{h=1}^{T_{\text{syn}}},$$

which can be directly consumed by the maintenance simulation/optimization layers.

Figure 3.6 summarizes the baseline workflow adopted to generate the synthetic operability dataset from historical metocean observations. Starting from hourly wind speed and significant wave height records, the data are first time-indexed and combined with vessel-specific operability thresholds to derive historical binary accessibility states. These state sequences are then used to characterize seasonal accessibility dynamics and to estimate month-dependent Markov transition matrices for each transport mode. Finally, the estimated transition structure is sampled sequentially to generate long-horizon synthetic hourly operability series, which preserve the empirical persistence and switching behaviour of accessible and non-accessible weather windows while remaining computationally lightweight.




Figure 3.6: Workflow for generating synthetic hourly operability datasets from historical metocean observations.

3.5.5 Validation of the synthetic operability dataset

To verify that the synthetic generator preserves the targeted monthly switching structure, we recompute transition matrices from the synthetic series and compare them to the original MLE matrices. For vessel v and month m , let $\hat{\mathbf{P}}_{v,\text{syn}}^{(m)}$ denote the transition matrix estimated from the synthetic series. We quantify the discrepancy using mean squared error (MSE) averaged over months and matrix entries:

$$\text{MSE}_v = \frac{1}{12} \sum_{m=1}^{12} \frac{1}{4} \sum_{i \in \{0,1\}} \sum_{j \in \{0,1\}} \left(p_{v,ij}^{(m)} - \hat{p}_{v,ij,\text{syn}}^{(m)} \right)^2. \quad (3.15)$$

Low values of MSE_v indicate that the synthetic generator reproduces the intended month-conditioned operability dynamics. This numerical agreement is shown compactly in Figure 3.7 and is further supported by the visual comparisons in Figures 3.8–3.9, which confirm consistency in seasonal accessibility levels, switching behaviour, and persistence of operational and non-operational weather windows.

A compact numerical summary of this validation is reported in Figure 3.7, which shows the vessel-wise mean squared error between the historical and synthetic monthly transition matrices. The uniformly low values confirm that the synthetic generator closely reproduces the targeted month-conditioned operability dynamics across all transport modes considered.




Figure 3.7: Mean squared error between the historical and synthetic operability transition matrices for CTV, SOV, JuV, and Helicopter.

Beyond the scalar MSE metric in Eq. (3.15), the similarity between historical and synthetic operability dynamics is also assessed visually for the four transport modes most relevant to the benchmark applications: CTV, SOV, Jack-up Vessel, and Helicopter. Figure 3.8 compares the month-specific transition probabilities governing persistence and switching between operational and non-operational hourly states. In particular, the off-diagonal terms quantify the probability of entering or leaving an accessible weather window at the next hour, while the diagonal terms reflect persistence of favourable or unfavourable conditions.

Figure 3.9 summarizes the seasonal accessibility structure through month-by-vessel heatmaps of the implied percentage of operational hours. The close agreement between historical and synthetic patterns across vessels and seasons supports the use of the synthetic operability database as a realistic long-horizon benchmark input for the MILP and DRL maintenance models.

Finally, because offshore maintenance performance depends not only on the fraction of accessible hours but also on the duration of consecutive accessible and inaccessible windows, Figure 3.10 compares the expected lengths of operational and non-operational hourly spells implied by the historical and synthetic Markov chains. This highlights whether the synthetic generator reproduces the persistence of calm-weather access windows and prolonged storm-driven inaccessibility periods.

The adopted approach is intentionally decision-centric: it generates the binary accessibility signal required by the O&M decision models, rather than aiming to reconstruct the full joint distribution of wind and wave processes.

3.6 Synthetic RUL prediction generator

The RUL module introduces the time-dependent health dynamics of the selected critical components. While the previous sections define the wind farm layout, component set, transport assets, and metocean-driven accessibility states, the RUL generator determines when each component would fail in the absence of maintenance. For each turbine-component pair (i, c) , where $i \in \{1, \dots, N_{WT}\}$




Figure 3.8: Comparison of historical and synthetic month-dependent hourly transition probabilities for CTV, SOV, JuV, and Helicopter.




Figure 3.9: Month-by-vessel heatmaps of the implied percentage of operational hours for the historical and synthetic datasets.

and c belongs to the critical component set defined in Section 3.3, a latent true RUL is generated and updated over time. This variable creates the maintenance demand that later propagates to the turbine availability indicator $a_{i,h}$ in Eq. (3.3), and therefore to farm-level production and lost production in Eqs. (3.3)–(3.4).

Given the current scarcity of real-life condition-monitoring and prognostic datasets for offshore wind components, the benchmarking framework relies on a synthetic RUL prediction generator. The objective is to create a prognostic dataset that is not only realistic, but also flexible enough




Figure 3.10: Expected lengths of consecutive operational and non-operational hourly weather windows for CTV, SOV, JuV, and Helicopter.

to represent different levels of monitoring quality and prediction uncertainty. For this reason, the generator is designed as a modular and interchangeable layer of the framework, so that synthetic predictions can be seamlessly replaced by real operational data as soon as such data become available. This section defines the synthetic RUL prediction generator adopted in this thesis.

The generator is designed to provide a point estimate and an uncertainty measure, while remaining fully reproducible, scalable, and consistent with maintenance-induced resets and partial rejuvenation effects.

The generation flow is summarised as:

1. sample a latent *true* RUL from a component-specific Weibull lifetime model;
2. compute a time-dependent uncertainty level that emulates improving prognostics over time;
3. generate a point prediction by perturbing the true RUL with Gaussian noise;
4. define a predictive distribution;
5. compute cumulative failure probabilities over a given time horizon from the predictive distribution.

3.6.1 Latent true RUL generation

The latent lifetime of each critical component is represented using a two-parameter Weibull distribution. The Weibull model is adopted because its shape parameter can represent decreasing, constant, or increasing failure hazards and because it is widely used to model the reliability of mechanical and electromechanical wind turbine components (Le & Andrews, 2016; Li et al., 2021). More specifically, a shape parameter $\alpha_c < 1$ represents a decreasing hazard rate, $\alpha_c = 1$ corresponds to a constant hazard rate, and $\alpha_c > 1$ represents an increasing hazard rate associated with ageing and wear-out.

The component-specific Weibull parameters are adopted from the generic offshore wind case study presented by Li et al. (2021), whose reliability inputs were assembled primarily from published wind turbine failure and maintenance studies. The parameters should therefore be interpreted as literature-derived benchmark values rather than as site-specific estimates calibrated to a particular turbine design or offshore wind farm.

For each component c , the latent lifetime is defined as

$$L_c \sim \text{Weibull}(\alpha_c, \lambda_c), \quad (3.16)$$

where α_c is the dimensionless shape parameter and λ_c is the scale parameter expressed in days. A newly initialised component, or a component installed following full replacement, receives one realisation from this distribution. The sampled value represents the hidden physical lifetime used by the simulation environment and is not directly observed by the maintenance planner.

The corresponding component-specific probability density function is

$$f_c(d; \lambda_c, \alpha_c) = \frac{\alpha_c}{\lambda_c} \left(\frac{d}{\lambda_c} \right)^{\alpha_c - 1} \exp \left[- \left(\frac{d}{\lambda_c} \right)^{\alpha_c} \right], \quad (3.17)$$

where $d \geq 0$ denotes the component lifetime in days.

Sampling is performed using inverse-transform sampling:

$$L_c = \lambda_c [-\ln(U)]^{1/\alpha_c}, \quad U \sim \text{Unif}(0, 1). \quad (3.18)$$

Sampling is performed because the benchmark requires a latent fixed lifetime for each component installed on each turbine. The Weibull distribution defines the statistical variability of component lifetimes across the wind farm, but the simulation must assign one concrete lifetime realisation to every turbine–component pair (i, c) . This sampled lifetime represents the hidden true failure time of that component in the absence of maintenance.

At observation time k , the true remaining useful life is obtained from the sampled latent lifetime and the component virtual age $A_c(k)$:

$$\text{RUL}_{\text{true},c}(k) = \max \{L_c - A_c(k), 0\}. \quad (3.19)$$

To avoid unrealistically short initial lifetimes in the benchmark dataset, the implementation includes a minimum-life rejection rule:

$$\text{RUL}_{\text{true},c} \leftarrow \text{resample until } \text{RUL}_{\text{true},c} > \text{RUL}_{\text{min}}, \quad (3.20)$$

with $RUL_{\min} = 0.3 \cdot 365$ days.

The minimum-life rejection rule avoids conflating the benchmark with early defect-liability or warranty cases. Very early failures may be treated differently in practice, since OEM warranties, defect-warranty clauses, service agreements, or insurance products can shift repair responsibility away from ordinary owner/operator O&M budgets (Hörmann, 2016; Munich Re, 2023; SGS, 2026). The benchmark instead targets the stable operational phase, where maintenance timing, access constraints, downtime losses, and logistics decisions are the primary objects of evaluation.

Figure 3.11 illustrates the Weibull density and the induced truncated density.

Component-specific Weibull probability density functions (PDF) and cumulative distribution functions (CDF) are shown in Figures 3.12a–b. These figures show the component-specific Weibull lifetime models used to initialise the latent true RUL values. Panel (a) reports the probability density functions, indicating the most likely lifetime ranges for each component, while panel (b) reports the corresponding cumulative distribution functions, indicating the probability that a component has failed before a given RUL value. The different curve shapes reflect the component-specific Weibull parameters: components with lower scale parameters or steeper CDFs are more likely to fail earlier, whereas broader density curves represent larger variability in possible lifetimes. The vertical dashed line indicates the imposed minimum initial RUL, $RUL_{\min} = 0.3 \cdot 365$ days, below which sampled lifetimes are rejected and resampled.

3.6.2 Time-varying uncertainty model

At each observation time k , the generator provides an uncertainty measure $\sigma_c(k)$ (in days) for component c , intended to emulate how prognostic accuracy may improve as the component approaches end-of-life, as more condition information becomes available.

Let t_k denote the elapsed time input to the predictor, and let $RUL_{\text{true},c}(k)$ denote the true remaining life (days) at time k . The implementation defines:

$$\sigma_{\max}(k) = \gamma_{\max} RUL_{\text{true},c}(k), \quad \sigma_{\min}(k) = \gamma_{\min} RUL_{\text{true},c}(k), \quad (3.21)$$




Figure 3.11: True RUL draws (days) under Weibull sampling with minimum-life rejection.




Figure 3.12: Weibull lifetime models per component: (a) PDFs; (b) CDFs.

where $\gamma_{\max}$ and $\gamma_{\min}$ are scaling parameters that define the upper and lower bounds of the prediction uncertainty envelope as proportions of the true RUL. In this way, the synthetic predictor is allowed to exhibit larger uncertainty when failure is still far away and progressively lower uncertainty as the component approaches end of life.

The standard deviation is then updated using a logistic schedule:

$$\sigma_c(k) = \sigma_{\max}(k) - \frac{\sigma_{\max}(k) - \sigma_{\min}(k)}{1 + \exp\left(-\frac{t_k - a RUL_{\text{true},c}(k)}{b RUL_{\text{true},c}(k)}\right)}, \quad (3.22)$$

where:

- a controls the *timing* of the transition,
- b controls the *steepness* (abrupt vs gradual tightening).

Higher values of a shift the tightening of the prediction uncertainty to later stages of component life, whereas lower values cause the predictor to become confident earlier. Higher values of b produce a smoother and more progressive reduction in uncertainty, whereas lower values generate a steeper transition between the high- and low-uncertainty regimes.

The parameters $\gamma_{\max}$, $\gamma_{\min}$, a , and b were not calibrated using operational prognostic data. Instead, they define controlled synthetic scenarios representing different initial levels of prediction uncertainty and different rates at which prognostic confidence improves as component failure approaches. Their purpose is to enable reproducible testing of the decision-making methods under explicitly defined uncertainty conditions. In an operational application, the synthetic uncertainty schedule can be replaced by prediction-error distributions estimated from historical prognostic outputs and observed component lifetimes.

Figure 3.13 illustrates the resulting evolution of the prediction standard deviation as the component approaches its end of life. Figures 3.14 and 3.15 further show how varying $(\gamma_{\max}, \gamma_{\min})$ and (a, b) , respectively, changes the uncertainty trajectory.




Figure 3.13: Prediction uncertainty scheduling model.




Figure 3.14: Sensitivity of σ/RUL_{true} to the uncertainty envelope $(\gamma_{max}, \gamma_{min})$.




Figure 3.15: Effect of (a, b) on uncertainty timing and steepness. (a) Varying a shifts the onset of improvement; (b) varying b changes steepness.

3.6.3 Prediction model

Given the true remaining life at time k , the point prediction is generated as:

$$RUL_{\text{pred},c}(k) = \max(RUL_{\text{true},c}(k) + \epsilon_c(k), 0), \quad \epsilon_c(k) \sim \mathcal{N}(0, \sigma_c^2(k)), \quad (3.23)$$

where $\sigma_c(k)$ is given by Eq. (3.22). The truncation at zero ensures physical consistency (negative remaining life is not permitted). A component is flagged as failed when its true remaining life reaches zero; in such cases the predictor returns null RUL values.

To support risk-aware decision-making (e.g., computing failure probabilities over a planning horizon), the generator converts the point prediction into a predictive distribution. Specifically, for each component c at observation k , a truncated normal distribution is defined:

$$\mathcal{R}_c(k) \sim \text{TruncNormal}\left(RUL_{\text{pred},c}(k), \sigma_c(k); [0, \infty)\right). \quad (3.24)$$

Figure 3.16 shows representative predictive densities at different elapsed times, illustrating how the distribution narrows as $\sigma_c(k)$ decreases.

For a given horizon d (days), the generator computes the cumulative probability that the component fails within d days using the predictive CDF:

$$\mathbb{P}(\mathcal{R}_c(k) \leq d) = F_{\mathcal{R}_c(k)}(d), \quad (3.25)$$

Moreover, discrete day-by-day probabilities over a horizon H , can be computed as:

$$p_c^d(k) = \mathbb{P}(d-1 < \mathcal{R}_c(k) \leq d) = F_{\mathcal{R}_c(k)}(d) - F_{\mathcal{R}_c(k)}(d-1), \quad d = 1, \dots, H. \quad (3.26)$$

To illustrate the end-to-end behaviour of the generator under a steadily decreasing true RUL, Figure 3.17a first compares the relative uncertainty schedule over the last 500 days of component life for two representative cases: an optimistic baseline case, in which uncertainty reduces progressively and high-confidence predictions are reached relatively early, and a delayed-tightening case, in which




Figure 3.16: Predictive densities at different elapsed times (same RUL_{true}), highlighting uncertainty tightening.

the reduction in uncertainty is postponed until close to failure. This comparison provides the basis for interpreting the practical behaviour shown in Figures 3.17b and 3.17c.

Figure 3.17b illustrates the RUL prediction under the optimistic baseline setting, together with an uncertainty band derived from $\sigma_c(k)$. In this case, the prediction band contracts earlier along the degradation trajectory, representing monitoring conditions under which prognostic confidence improves well before end of life.

Figure 3.17c shows the corresponding behaviour under delayed tightening. Here, the uncertainty band remains wider for most of the remaining lifetime and narrows only in the final phase, representing situations in which high-confidence predictions are available only close to failure.

The generator is fully reproducible under fixed random seeds and the sensitivity figures provided in this section offer a transparent basis for selecting parameterisations that match the intended benchmarking use-case (e.g., early confidence vs late confidence availability).

The overall logic of the synthetic RUL prediction generator is illustrated in Figure 3.18. Starting from a sampled latent true RUL, the procedure enforces a minimum-life threshold to avoid unrealistically short initial lifetimes. For accepted samples, a time-dependent uncertainty level is computed, after which the generator produces both a point prediction and a predictive distribution. These outputs are then used to derive cumulative failure probabilities over the decision horizon.




Figure 3.17: Illustration of uncertainty behaviours in the synthetic RUL prediction generator. (a) Relative uncertainty schedules for optimistic baseline and delayed-tightening case. (b) Prediction path and associated uncertainty band under the optimistic baseline setting. (c) Prediction path and associated uncertainty band under delayed tightening.




Figure 3.18: Workflow of the synthetic RUL prediction generator

3.7 Maintenance actions

The previous sections define the state variables that describe the offshore wind O&M environment before a decision is made: the wind farm contains N_{WT} turbines indexed by i , each turbine contains the selected critical components indexed by c , each component has a latent true RUL and a predicted RUL, and each transport asset has time-dependent accessibility determined by metocean conditions. The maintenance-action module defines how decisions modify these variables. In particular, a maintenance action changes the component health state, updates the latent and predicted RUL trajectories, affects the turbine availability indicator $a_{i,h}$ introduced in Eq. (3.3), and therefore propagates to farm-level production and lost production in Eqs. (3.3)–(3.4). At the same time, each action is associated with specific vessel requirements, execution times, and costs, linking the component-health model to the transport and accessibility variables defined in Sections 3.4 and 3.5.

As discussed in Chapter 1 and Chapter 2, offshore wind O&M is not limited to a binary *maintain / do not maintain* choice. In practice, operators combine corrective, preventive, condition-based,

predictive, and opportunistic strategies, and interventions may range from relatively light repairs to full component replacement depending on the observed or predicted severity of degradation, the expected RUL, and the offshore logistical opportunity available at the time of intervention.

This thesis therefore adopts a *multi-level maintenance action set* for each critical component:

$$\mathcal{A} = \{\text{minor repair, major repair, replacement}\}. \quad (3.27)$$

This choice is consistent with the offshore wind O&M literature reviewed in Chapter 2, where maintenance models frequently distinguish among replacement, major repair, and minor repair because these actions differ substantially in vessel requirements, execution time, cost, and restorative effect.

The use of multiple maintenance levels serves two purposes. First, it improves realism by recognising that offshore interventions differ not only in direct cost but also in restorative strength and logistical burden. Second, it makes the O&M models more informative, because a decision-support method must learn not only *when* to intervene, but also *which type of intervention* is most appropriate given the current health state, remaining life, and resource availability.

The three maintenance levels considered in the benchmark are interpreted as follows:

- **Minor repair:** A minor repair represents a limited intervention aimed at restoring function or slowing degradation without fully renewing the component. Typical examples include localized blade repair, replacement of a small subassembly, lubrication, sealing, sensor replacement, or calibration and adjustment activities. Such tasks are generally associated with shorter offshore interventions, lower direct cost, and lighter logistical requirements. In many cases they can be executed using port-based resources such as crew transfer vessels (CTVs). In reliability terms, a minor repair is treated as an *imperfect maintenance* action: it improves the condition of the asset, but does not return it to an as-good-as-new state.
- **Major repair:** A major repair corresponds to a more extensive restorative action than a minor repair, but still stops short of full component replacement. Examples include repair or exchange of larger internal assemblies, more substantial structural restoration, or interventions requiring longer work duration, more technicians, and more capable access assets. In offshore wind practice, these tasks generally imply stricter weather and vessel constraints than minor actions, and they are typically justified when degradation is advanced but full replacement is either unnecessary or economically unattractive. As with minor repair, major repair is modelled here as *imperfect maintenance*, but with a larger restorative effect.
- **Replacement:** Replacement is the strongest intervention and corresponds to complete renewal of the failed or highly degraded component. This action is appropriate when: (i) functional failure has already occurred, (ii) the predicted degradation state is too severe for repair to be cost-effective, or (iii) the planner chooses full renewal because an offshore campaign already mobilises the required heavy logistics. Replacement is treated as an *as-good-as-new* action and therefore resets the degradation process by drawing a new lifetime from the component-specific Weibull distribution defined in Section 3.6.

Figure 3.19 summarises the reliability interpretation adopted for the three maintenance levels. Minor repair and major repair are both treated as imperfect maintenance actions, but with different




Figure 3.19: Effects of minor, major repairs and replacements.

restorative effects: major repair yields a larger condition improvement than minor repair, while neither returns the component to an as-good-as-new state. Replacement, by contrast, is treated as a perfect maintenance action and therefore resets the component condition by initiating a new lifetime cycle.

3.7.1 Virtual-age formulation of maintenance effects

To model the physical effect of maintenance on component life, this thesis adopts a *virtual-age* formulation.

Let L_c denote the component-specific latent lifetime sampled according to Eq. (3.16), and let $RUL_{\text{true},c}^-(t)$ denote the corresponding true RUL:

$$RUL_{\text{true},c}^-(t) = \max(L_c - A_c^-(t), 0). \quad (3.28)$$

The benchmark assumes that imperfect maintenance reduces the accumulated virtual age by a fixed fraction. For action $a \in \{\text{minor}, \text{major}\}$,

$$A_c^+(t) = (1 - \rho_a) A_c^-(t), \quad (3.29)$$

where ρ_a is the maintenance effectiveness factor. Following the imperfect-maintenance formulation adopted in Li et al. (2022c), different intervention levels are represented through action-specific restoration factors. In this thesis, the following values are adopted:

$$\rho_{\text{minor}} = 0.30, \quad \rho_{\text{major}} = 0.50. \quad (3.30)$$

Hence, a minor repair reduces virtual age by 30%, whereas a major repair reduces virtual age by 50%.

The corresponding post-maintenance true RUL becomes

$$RUL_{\text{true},c}^+(t) = L_c - A_c^+(t) = L_c - (1 - \rho_a)A_c^-(t) = RUL_{\text{true},c}^-(t) + \rho_a A_c^-(t). \quad (3.31)$$

Equation (3.31) clarifies an important modeling point: the life extension generated by imperfect maintenance is proportional to the *accumulated virtual age* rather than directly to the current RUL. This is consistent with the notion of imperfect repair in maintenance theory, where the intervention partially rejuvenates the asset without erasing its degradation history.

For full replacement, the component is assumed to return to an as-good-as-new state:

$$A_c^+(t) = 0, \quad L_c^{\text{new}} \sim \text{Weibull}(\alpha_c, \lambda_c), \quad RUL_{\text{true},c}^+(t) = L_c^{\text{new}}. \quad (3.32)$$

Thus, replacement does not merely extend the existing lifetime; it starts a new stochastic life cycle consistent with the Weibull-based synthetic lifetime generator used throughout the thesis.

In terms of RUL_{true} , if the component has current virtual age $A_c^-(t)$ and current true remaining life $RUL_{\text{true},c}^-(t)$, then:

$$RUL_{\text{true},c}^+(t) = \begin{cases} RUL_{\text{true},c}^-(t) + 0.30A_c^-(t), & \text{minor repair,} \\ RUL_{\text{true},c}^-(t) + 0.50A_c^-(t), & \text{major repair,} \\ L_c^{\text{new}}, \quad L_c^{\text{new}} \sim \text{Weibull}(\alpha_c, \lambda_c), & \text{replacement.} \end{cases} \quad (3.33)$$

If, for implementation convenience, virtual age is approximated by elapsed age since the last full renewal, Eq. (3.33) can be applied directly using the tracked age state of the component. In this way, minor and major actions are explicitly assigned lower and higher restorative impact respectively, while replacement samples a new Weibull life.

Figure 3.20 illustrates the lifetime-update logic adopted in the benchmark. The descending lines represent degradation over time, while the vertical jumps correspond to maintenance-induced restoration. The “no action” trajectory shows continued depletion of RUL until failure. A major repair produces a larger upward jump because it applies a 50% virtual-age decrease, whereas a minor repair produces a smaller jump corresponding to a 30% virtual-age decrease. Replacement is conceptually different from both: it does not simply shift the current degradation path upward, but starts a new stochastic life trajectory drawn from the Weibull model.

3.7.2 Effect on the predictive distribution

Maintenance actions must also update the prognostic layer, not only the latent degradation state. After any action at time t , the predictive distribution introduced in Section 3.6 is re-centred around the updated true RUL. Using the notation already adopted in the thesis, the post-maintenance predictive distribution is:

$$RUL_c^+(t) \sim \text{TruncNormal}(\mu_c^+(t), \sigma_c^+(t); [0, \infty)), \quad (3.34)$$

with

$$\mu_c^+(t) = RUL_{\text{pred},c}^+(t), \quad (3.35)$$




Figure 3.20: Update of component RUL under no action, major repair, and minor repair.

$$RUL_{\text{pred},c}^+(t) = RUL_{\text{true},c}^+(t) + \epsilon_c(t), \quad (3.36)$$

$$\epsilon_c(t) \sim \mathcal{N}\left(0, (\sigma_c^+(t))^2\right). \quad (3.37)$$

In other words, the maintenance action first updates the hidden degradation state through Eqs. (3.29)–(3.32), and only then the new RUL prediction and its uncertainty are generated. This preserves coherence between the physical lifetime model and the prognostic information presented to the maintenance planner.

3.7.3 Action-specific resource requirements

Besides modifying the latent degradation state, each maintenance action is associated with specific execution requirements in terms of vessel class, intervention duration, and technician demand. This is essential for the benchmark because the economic and operational consequences of an action depend not only on its restorative effect, but also on the offshore resources that must be mobilised to carry it out.

In offshore wind O&M, maintenance actions on different components can require substantially different logistics. For example, full blade replacement typically requires heavy-lift capabilities and long intervention times, whereas a minor corrective action on the same component may be executable using a CTV and a small technician team. Likewise, major repairs generally require more capable access assets and longer offshore occupation than minor repairs, even when they do not escalate to full replacement.

To represent this explicitly, the benchmark assigns to each component–action pair:

- a required vessel class;

- an active intervention duration in hours;
- a required number of technicians.

The parameterisation combines component reliability data, empirical offshore repair records, and established vessel-task compatibility principles. The Weibull shape and scale parameters are adopted from the offshore wind benchmark of Li et al. (2021). Intervention durations are adapted from the component- and repair-category data reported by Carroll et al. (2016), while the association of minor, major, and replacement actions with progressively more capable offshore assets follows the logistics hierarchy described by Dalgic et al. (2015c). Technician requirements are informed by the empirical crew-demand data of Carroll et al. (2016) and simplified to a consistent benchmark hierarchy of two, four, and eight technicians for minor repair, major repair, and replacement, respectively.

These parameters are used jointly with the accessibility model of Section 3.4 and the cost model in subsequent chapters. In this way, the choice between minor repair, major repair, and replacement affects not only the post-maintenance RUL, but also the feasibility, waiting time, and direct execution cost of the intervention.

The values in Table 3.4 reflect the increasing logistical burden associated with stronger interventions. Minor repairs are assumed to be executable with a CTV and a two-technician crew, major repairs require an SOV-based intervention and four technicians, while full replacement requires a jack-up vessel and mobilises eight technicians. This hierarchy is consistent with the intended benchmark logic: stronger maintenance actions provide greater lifetime restoration, but also require scarcer and more expensive offshore resources and longer occupation times. As a result, the action selected by the planner affects not only the post-intervention degradation trajectory, but also the operational feasibility and cost of execution.

Generally speaking, the action set is intended to represent the following operational logic: Replacement is mandatory after failure and may also be selected preventively when the component is critically degraded, when a major offshore campaign is already mobilised, or when the remaining life extension obtainable from repair is insufficient to justify the intervention. Major repair is selected when degradation is advanced but the component is still recoverable, and when a stronger rejuvenation effect is preferred to avoid near-term repeat interventions. Minor repair is selected when the component is degraded but not yet critical, especially when an opportunistic offshore access window exists and a relatively cheap intervention can meaningfully postpone failure.

This hierarchy is consistent with the maintenance strategies discussed earlier in the thesis. Corrective strategies restore failed components, preventive and predictive strategies intervene before failure, and opportunistic strategies exploit already-mobilised resources to reduce marginal cost and downtime.

Table 3.4: Component-specific maintenance action requirements used in the benchmark

Component	Shape	Scale [days]	Replacement			Major repair			Minor repair		
			Time [h]	Vessel	Techs	Time [h]	Vessel	Techs	Time [h]	Vessel	Techs
Blades	3	1847	292	JuV	8	21	SOV	4	11	CTV	2
Bearing	2	1811	41	JuV	8	15	SOV	4	8	CTV	2
Gearbox	3	1477	236	JuV	8	19	SOV	4	10	CTV	2
Generator	2	1594	86	JuV	8	17	SOV	4	9	CTV	2
Pitch system	3	1144	41	JuV	8	15	SOV	4	8	CTV	2

3.8 Benchmark implementation

The previous sections defined the main modeling blocks of the offshore wind O&M environment shown in Figure 3.1, including the wind farm layout, transportation assets, metocean accessibility model, critical component set, synthetic RUL prediction generator, maintenance actions, and production and cost consequences. In this section, these elements are integrated into a common benchmark simulation environment that provides the executable interface used by the optimization- and learning-based models developed in the following chapters.

Within this benchmark environment, each turbine–component pair (i, c) is represented through a time-indexed state trajectory describing the evolution of component health, prognostic information, maintenance actions, accessibility conditions, and resulting production and cost impacts over time. The benchmark therefore acts as the common operational layer through which maintenance decisions are evaluated under identical assumptions regarding degradation, logistics, weather accessibility, and economic consequences.

For each turbine–component pair (i, c) , the benchmark stores the following information:

- True state: $\text{age} * i, c(t)$, $\text{RUL}^{\text{true}} * i, c(t)$, and a failure indicator. These variables define the latent physical condition of the component and determine when failures occur, when turbines become unavailable, and how downtime propagates to production losses and costs.
- Prognostic information: $\mathcal{R}_{i,c}(t)$ and uncertainty bounds. These variables represent the imperfect RUL information available to the decision-support models. The optimization and DRL approaches therefore interact with uncertain prognostic information rather than directly observing the latent true RUL.
- Maintenance history: action type, start time, duration, selected access asset, and post-maintenance state update. These variables record how maintenance interventions modify component condition, resource usage, accessibility requirements, downtime, and economic outcomes over time.

Together, these variables provide the common simulation interface used throughout the remainder of the thesis. The benchmark environment can therefore be operated using predefined maintenance strategies or through the optimization- and learning-based decision-making approaches introduced in Chapters 4 and 6.

Together, these three groups of variables ensure that the dataset contains both the hidden ground-truth degradation process and the observable information used for decision-making. This distinction is essential for comparing optimization- and learning-based O&M strategies under the same component dynamics, prognostic uncertainty, maintenance effects, and cost consequences.

This structure enables identical information interfaces across the MILP and DRL models and supports transparent sensitivity analysis with respect to prognostics quality (e.g., varying $\sigma(h)$ or update frequency).

Figure 3.21 summarizes the environment developed in this chapter using an Extended Design Structure Matrix (XDSM). XDSM is a graphical formalism for representing complex design, analysis, and simulation processes by combining the main process sequence on the diagonal with the associated data exchanges shown off-diagonal Lambe & Martins (2012). Here, it is used to

show how the simulation environment is built, the main modeling steps, the key inputs required by each step, and the principal information flows between modules. The diagonal blocks correspond to the main process steps, the upper boxes summarize the main inputs, the off-diagonal grey boxes denote exchanged information, and the left-hand boxes reference the governing equations and tables associated with each block.

This subsection completed the definition of the benchmark outputs by specifying how true degradation states, prognostic information, and maintenance histories are stored and passed to subsequent models. This closes the modeling chain introduced in this chapter, from the physical representation of the wind farm and its critical components, to the synthetic generation of accessibility and health information, and finally to the standardized data structures that support maintenance decision-making.

3.8.1 Modeling scope and abstractions

The modeling environment is designed to balance representativeness and tractability. Its purpose is to expose the optimization-based and learning-based methods developed in Chapters 4 and 6 to the same offshore wind O&M system, rather than to reproduce all operational details of a specific wind farm. The framework should therefore be interpreted as a controlled benchmark environment: it captures the main couplings among component degradation, prognostic information, offshore accessibility, maintenance actions, and cost consequences, while abstracting from several site-specific and execution-level details.

The wind farm is represented through a modular fixed-bottom reference configuration. This makes it possible to vary the number of turbines while preserving a reproducible spatial structure and consistent travel-distance calculations. At the same time, the layout does not reproduce the bathymetry, cable routing, foundation-specific constraints, port-interface characteristics, or operational procedures of a particular project. Similarly, the metocean module translates weather conditions into vessel-specific accessibility states. This captures the decision-relevant effect of weather on whether maintenance can be performed, but it abstracts from detailed marine-operation dynamics such as interrupted transfers, visibility limitations, partial task completion, sea-state-dependent execution times, and task continuation across multiple weather windows.

The degradation and prognostic layers are also defined synthetically. True component lifetimes are sampled from Weibull distributions, while the prognostic module provides uncertain RUL estimates with time-varying uncertainty. This structure enables controlled and repeatable evaluation of decision-making methods under uncertain health information, but it does not claim to reproduce the full statistical structure of real SCADA, inspection, or condition-monitoring data. The same logic applies to the maintenance model. Minor repairs, major repairs, and replacements are represented through a virtual-age formulation with action-specific restorative effects, costs, durations, and resource requirements. This provides a consistent action taxonomy for both optimization and learning, while simplifying the physical mechanisms through which actual repairs influence future degradation.

The logistics representation is kept at planning level. The framework includes access assets, action-specific resource requirements, travel assumptions, and cost components, but does not model detailed vessel routing, technician rostering, spare-part inventory dynamics, port congestion, contractual vessel availability, or multi-farm resource sharing. These abstractions are appropriate for




Figure 3.2.1: XDSM representation of the O&M environment defined in Chapter 3.

comparing decision-making logic under common assumptions, but they limit the direct operational interpretation of the resulting schedules.

These assumptions are therefore benchmark-design choices rather than claims of full operational fidelity. They make it possible to isolate the effect of the decision-making paradigm under common information, cost, degradation, and accessibility assumptions. At the same time, they delimit the scope of the conclusions drawn in Chapters 4 and 6.

3.9 Conclusions

This chapter addressed **Research Question Q2**: *How to model the dynamics of offshore wind farm O&M to support maintenance decision-making?* To answer this question, the chapter established a standardized reference framework for offshore wind farm O&M planning, designed to support consistent verification, validation, and comparison of alternative decision-support methods.

The chapter defined a representative fixed-bottom offshore wind farm together with a transparent spatial layout, a tractable but representative set of critical components, a fleet of transportation and access assets, and a consistent taxonomy of maintenance actions. In addition, it introduced two synthetic but statistically grounded data-generation layers: a month-dependent Markov-chain model for metocean-driven accessibility and a synthetic RUL prediction generator combining Weibull-based component lifetimes with time-varying prognostic uncertainty. Coupled with the virtual-age formulation of maintenance effects, these elements provide a coherent representation of failures, accessibility windows, intervention feasibility, condition evolution, and post-maintenance state updates.

The main contribution of this chapter is therefore not only the specification of a numerical case study, but the construction of a common modeling environment. By exposing all subsequent methods to the same wind farm configuration, logistics assumptions, weather-driven accessibility structure, maintenance action set, and prognostic information interface, the framework enables controlled and like-for-like comparison between decision-support paradigms.

This common environment provides the basis for the methodological developments in the remainder of the thesis. Chapter 4 uses the benchmark defined here to formulate tactical offshore wind O&M planning as a mixed-integer linear programming problem under accessibility, resource, and prognostic constraints. Chapter 5 then analyses the limitations of this optimization-based paradigm and positions those limitations in relation to the broader literature on sequential and adaptive maintenance decision-making. Finally, Chapter 6 builds on the same benchmark environment to develop a deep reinforcement learning framework, allowing the two decision-support paradigms investigated in this thesis to be assessed on a directly comparable basis.

Chapter 4

Optimization-Based O&M Planning

This chapter addresses **Research Question Q3**: *How can O&M planning be formulated as an optimization model that incorporates condition and prognostic information?* In answering Q3, the chapter develops the first methodological decision-support paradigm of this thesis: a mixed-integer linear programming (MILP) framework for tactical offshore wind farm maintenance planning. Building on the research gaps identified in Chapter 2 and the modeling environment established in Chapter 3, the objective here is to show how prognostic information can be translated into explicit, tractable planning decisions under offshore operational constraints.

More specifically, this chapter formulates tactical O&M planning as a constrained optimization problem in which maintenance timing, intervention type, and campaign initiation are determined jointly while accounting for weather-driven accessibility, shared logistics resources, and uncertain component health. In contrast to traditional age-threshold policies, the proposed approach integrates probabilistic RUL information into a rolling-horizon stochastic optimization framework, allowing maintenance decisions to respond to evolving prognostic estimates rather than relying solely on static maintenance triggers. In this way, the chapter operationalises the core thesis premise that offshore wind O&M planning should combine condition information, access feasibility, and marine logistics within a single decision-support model.

The role of this chapter within the thesis is therefore twofold. First, it provides a concrete optimization-based answer to Q3 by developing and testing a prognostic-driven tactical planner on the common benchmark introduced in Chapter 3. Second, it establishes the optimization baseline against which the later learning-based framework will be interpreted. The results show the value of explicitly incorporating prognostic information into maintenance planning, while also exposing the structural assumptions and innate shortcomings that become central in Section 4.5 of this chapter, where the limitations of the MILP paradigm are analysed in detail and used later in Chapter 5 to motivate the transition toward DRL-based decision-making.

The chapter is organized as follows. Section 4.1 introduces the context of this chapter. Section 4.2 positions the proposed model with respect to prognostic-driven maintenance planning. Section 4.3 presents the optimization methodology, including the rolling-horizon logic, prognostic representation, maintenance and cost model, and MILP formulation. Section 4.4 summarizes the benchmark case-study setup used for evaluation. Section 4.5 reports the numerical results and discusses the operational insights obtained from the model. Finally, Section 4.6 concludes

the chapter and highlights the implications of the MILP results for the methodological transition developed in the next chapters.

This chapter has been published as Marco Borsotti, Rudy R. Negenborn, and Xiaoli Jiang: “Stochastic Optimization of Predictive Maintenance Scheduling for Offshore Wind Farms”, *Ocean Engineering*, vol. 357, article 125424, 2026, doi: 10.1016/j.oceaneng.2026.125424.

4.1 Introduction

Offshore wind energy is rapidly expanding, but operations and maintenance (O&M) remains a major challenge and cost driver. O&M activities for offshore turbines are significantly more expensive and complex than those onshore; for example, an average fixed-bottom offshore wind project incurs about 30.3 EUR/MWh in O&M costs, compared to roughly 12.1 EUR/MWh for onshore wind (Fox et al., 2022). In fact, O&M expenditures can comprise around one-third of an offshore wind farm’s total lifetime costs. These high costs are driven by the harsh marine environment, accessibility constraints, and the scale of offshore equipment, making reliability and efficient maintenance crucial for reducing the levelized cost of energy. Improving maintenance strategies to curb unplanned downtime and optimize interventions is therefore essential to enhance the economic viability of offshore wind power.

Maintenance strategies in the wind industry range from reactive (run-to-failure) approaches to proactive schemes (Fox et al., 2022). Historically, many offshore wind operators relied on reactive or time-based periodic maintenance, either fixing components only after a breakdown or replacing them at fixed intervals, but these traditional approaches have clear drawbacks. Reactive maintenance can maximize component life utilization, yet it often leads to catastrophic failures with lengthy downtime and high repair costs, especially in offshore conditions where repairs are slow and weather-dependent. Strict periodic maintenance, on the other hand, may replace parts too early (wasting useful life) or too late (risking unexpected failure) if actual degradation deviates from assumed schedules. Opportunistic maintenance has emerged as an intermediate strategy: it seeks to perform preventive tasks whenever a convenient opportunity arises (e.g., grouping repairs during a scheduled vessel visit or favorable weather window) to reduce overall costs (Fox et al., 2022). While opportunistic scheduling mitigates some inefficiencies, it still does not fully account for the actual condition of each component.

With advances in turbine monitoring and prognostics, the industry is moving toward more condition-based and predictive maintenance approaches (Fox et al., 2022). In condition-based maintenance (CBM), real-time sensor data is continuously analyzed to assess equipment health; maintenance is triggered when measured parameters exceed certain thresholds or anomaly criteria. Predictive maintenance extends this by using operational data and analytics to forecast failures before they occur, enabling just-in-time interventions. The most recent evolution is prescriptive maintenance, which not only predicts impending failures but also recommends optimized O&M actions. These proactive strategies are particularly compelling for offshore wind, where unplanned failures are extremely costly and access is limited by weather and vessel availability.

Recent research efforts have started to integrate prognostic information into maintenance decision-making to capture the dynamic nature of component degradation. Traditionally, the tasks of predicting failures and planning maintenance were handled separately: many studies focused

on RUL prognostics in isolation, while others optimized maintenance schedules assuming static failure distributions (Zhuang et al., 2023). To address this shortcoming, newer approaches aim to tightly couple prognostics with O&M planning so that maintenance can be dynamically adjusted as component health evolves (Vianna & Yoneyama, 2018; Lee & Mitici, 2023a; Jain et al., 2021). This allows decisions to be made online using the latest available health data.

Several studies have demonstrated the potential benefits of such integration, confirming that maintenance planning is increasingly moving toward tighter integration with asset condition information and dynamic operational factors. For example, Ying et al. (2026b) highlight the value of coupling condition assessment with opportunistic maintenance planning while accounting for offshore operational drivers such as environmental conditions, electricity prices, resource constraints, and wind-farm layout. In parallel, Ying et al. (2025) explicitly frame opportunistic maintenance in terms of subsystem fault-chain relationships, highlighting that dependencies among subsystems can materially affect maintenance timing and grouping decisions. To address the dynamic update of maintenance timing and action selection as new prognostic information becomes available, Borsotti et al. (2024) present a predictive maintenance model that integrates probabilistic RUL forecasts into a rolling optimization framework, reducing O&M costs by 8.7% compared to age-based methods. Uncertainty in prognostic information is considered by He et al. (2025), addressing biased and incomplete prognostics using a Bayesian calibration strategy for offshore wind. Finally, Xie et al. (2020) propose a predictive-opportunistic maintenance approach that incorporates lead-time and task dependency, achieving cost savings of up to 39%. These examples underscore that integrating RUL information into maintenance planning can significantly reduce downtime and costs in offshore wind operations. This strengthens the case for moving beyond static or purely age-based maintenance logic.

However, several limitations remain in current work. Most existing models either assume simplified binary decisions (maintain or not) as in Zhuang et al. (2023); Lee & Mitici (2023a), do not fully integrate with vessel logistics as in Raknes et al. (2017), or fail to update decisions dynamically as degradation evolves. Others focus solely on maintenance routing assuming fixed task lists, for example Lazakis & Khan (2021). A holistic framework that links RUL-based prognosis, dynamic optimization, and marine logistics is still lacking.

In this thesis, we address this gap by proposing a tactical, closed-loop maintenance planning framework that integrates probabilistic RUL estimates with stochastic optimization to schedule opportunistic maintenance in offshore wind farms. We embed a probabilistic representation of the RUL within a receding-horizon MPC loop in which new prognostic measurements update the belief state and its uncertainty contracts over time, and the tactical plan is repeatedly re-optimized under these updated beliefs. Finally, we link offshore feasibility inputs to tactical decisions by incorporating weather-driven accessibility, enabling the planner to jointly select maintenance timing and action types under uncertainty. Overall, these elements establish a closed-loop prognostics-informed tactical planner that reduces downtime and total O&M cost under offshore operational constraints.

4.2 Literature review

In this section, we first describe prognostic models, what kind of information they provide, their limitations, and how prognosis has been used to improve maintenance scheduling. To do that, we

explore methodologies from various industries and highlight the gap in the integration of prognostic information into O&M models for offshore wind.

At the moment, much research is in progress to create models and develop tools to assess the current and future state of health of components. For example, sensor-based monitoring and the use of data-analytic techniques have improved failure prognosis (Bhuria et al., 2024) by collecting and analyzing real-time information from sensors and making predictions on the future state of health of the system, allowing for timely maintenance actions (Kou et al., 2022; Le & Andrews, 2016).

Incorporating Prognostics and Health Management (PHM) information into O&M planning models offers significant potential for improving maintenance strategies. Ideally, predictive models like CBM and data-driven prognostics should be integrated into scheduling decisions to enable timely and cost-effective interventions. However, their practical implementation faces substantial challenges. Across works that treat predictive maintenance, a clear theme is the coupling of prognostic outputs with maintenance decisions, using prognostics to trigger or optimize maintenance timing rather than following fixed schedules. For instance, prognostic alerts or alarms are used by de Pater et al. (2022) and da Costa et al. (2023) to signal when an asset likely needs attention. Others feed continuous-valued RUL estimates into optimization algorithms (de Pater & Mitici, 2021; Zhuang et al., 2023) or into learned policies (Lee & Mitici, 2023a).

The cost/benefit trade-off of acting on prognostics is a recurring consideration: Chen et al. (2021) explicitly compute it via an online cost evaluation, Jain et al. (2021) incorporate it in their objective function, and Vianna & Yoneyama (2018) include operational penalty costs in their optimization. All these studies find that using prognostic info yields significant improvements such as lower maintenance costs, fewer failures, or higher availability, validating the core premise of predictive maintenance.

In general, the information provided by prognostic models consists of an estimated time to failure and an associated confidence limit (Sikorska et al., 2011), although there are instances where only the point prediction is used to set a threshold for initiating a maintenance cycle as in He et al. (2023a), where the net expected revenue is maximized for multi-component systems by optimizing the RUL percentage that triggers the replacement of the component.

One of the issues with integrating prognostic information in O&M models is the scarce availability of real-life prognostics data, which has been addressed in works such as Borsotti et al. (2024) and He et al. (2023b), by constructing a synthetic dataset of RUL prognostic data with a pre-generated prediction error, simulating an increasing confidence in the predictions as more Condition Monitoring (CM) data becomes available in time.

The presence of incomplete and uncertain prognostic information also poses a significant challenge, He et al. (2025) propose a probabilistic framework that compensates for it by calculating the expected cost of different maintenance actions based on the predicted RUL of components considering the associated uncertainty.

Many authors acknowledge that prognostics are not perfect and have built-in ways to handle uncertainty. A common approach is to use probabilistic RUL predictions (distributions or confidence bounds) instead of single-point estimates. Lee & Mitici (2023a) use Convolutional Neural Networks (CNN), Zhuang et al. (2023) use Bayesian Deep Learning (DL), and Jain et al. (2021) adopt Weibull Recurring Neural Networks (RNN) to output an entire RUL distribution. This allows decisions to be made with risk awareness, e.g. scheduling earlier if the risk of failure is high. Threshold-based policies with safety margins are another technique: de Pater et al. (2022) trigger an alarm only when

the predicted RUL falls below a conservative threshold, explicitly padding for error. Camci et al. (2019) and Vianna & Yoneyama (2018) embed reliability constraints (like *failure probability* > $X\%$) to ensure a conservative buffer against prognostic errors.

An insight from Jain et al. (2021) is that the optimal maintenance timing itself should adapt as prognostic uncertainty evolves. Early on, when predictions are uncertain, a policy might schedule maintenance a bit early (to be safe), but as more data improves the prediction, the policy can afford to push maintenance closer to failure.

Optimization models such as de Pater et al. (2022); de Pater & Mitici (2021); Zhuang et al. (2023); Camci et al. (2019); Vianna & Yoneyama (2018) typically use PHM inputs as fixed (sometimes stochastic) parameters, and solve over the long-term horizon for the maintenance plan that minimizes cost or maximizes availability. They guarantee optimality under the assumptions made, and can naturally handle constraints (spare inventory, maintenance crew availability, etc.) by inclusion in the optimization formulation.

In order to take into account new information about the failure dynamics of the critical components, O&M planning models should integrate new operational data to update the maintenance strategy, as seen in Li (2023). Approaches such as Model Predictive Control (MPC), for example, can enable dynamic scheduling by continuously updating maintenance plans based on real-time component degradation data, as in Borsotti et al. (2024), where probabilistic RUL predictions are used to evaluate the expected cost of performing different maintenance tasks. Similarly, in Kong & Liu (2025), a closed-loop MPC framework that accounts for probabilistic fatigue degradation and inspection outcomes is presented.

In recent offshore wind O&M optimization works, offshore operational mechanisms can be modeled explicitly, for example, Si et al. (2025b) propose a resource-centered maintenance strategy that leverages wind-wave information to estimate vessel accessibility and quantify maintenance opportunities linked to low-wind periods, while also accounting for key offshore resource drivers such as wave conditions, vessel fleet characteristics, and port-related limitations. In a different paper (Si et al., 2025a) the authors further develop a holistic opportunistic maintenance framework that distinguishes internal and external maintenance opportunities and explicitly links weather (wind speed and wave height) to access feasibility; they also frame offshore O&M decision-making as the combination of tactical decision support and operational maintenance strategy. Along the same line, recent work has also emphasized that environmental uncertainty itself deserves explicit treatment within offshore maintenance models. In particular, Ying et al. (2026a) show the relevance of representing uncertain environmental conditions as a dedicated input to maintenance planning. These contributions motivate representing (a) weather-driven accessibility and (b) seasonality in expected production loss when developing tactical, rolling-horizon maintenance planners for offshore wind farms. In the present study, this perspective is adopted through a decision-oriented representation of weather uncertainty, where offshore conditions are translated into accessibility information that can be embedded directly within a prognostics-informed rolling-horizon optimizer.

In reviewing the current state of predictive maintenance, it is evident that significant advancements have been made, although most research was developed within other industrial contexts and do not translate directly to the offshore wind sector, for example, the predictive scheduling approaches reviewed, such as de Pater et al. (2022); de Pater & Mitici (2021), only consider binary maintenance decisions of whether to maintain or not, neglecting diverse actions: minor, major repairs, and replacements.

On the other hand, current maintenance scheduling approaches for offshore wind farms remain fragmented, often separating strategical and tactical planning from short-term operational execution, relying on fixed thresholds or time intervals to trigger maintenance actions. Most existing models do not incorporate prognostic information and do not adapt maintenance plans in response to changes in the state of the system.

Therefore, there remains a gap in providing a dynamic, closed-loop framework that integrates real-time prognostic data to plan maintenance tasks for offshore wind farms, which this work addresses.

To tackle these challenges, we propose a closed-loop stochastic optimization model that integrates probabilistic RUL prognosis into tactical and operational maintenance planning. Using an MPC approach, our framework continuously updates component state observations, dynamically adapting maintenance decisions in response to evolving system conditions and uncertainties. Furthermore, our model explicitly incorporates offshore-specific logistical and weather constraints, diverse maintenance action options (minor repairs, major repairs, replacements), and is flexible enough to integrate prognostic information from generic models, synthetic datasets, or actual operational data.

In the following section we explain in detail the methodology implemented in the proposed model.

4.3 Methodology

In this section, we present a stochastic optimization framework for predictive maintenance scheduling in offshore wind farms, integrating probabilistic RUL prognostics with a closed-loop MPC strategy, dynamically updating maintenance schedules based on real-time prognostic data and operational constraints.

First, we describe our MPC-based rolling horizon approach, detailing how probabilistic RUL predictions inform both medium-term tactical decisions and short-term operational scheduling adjustments. Next, we specify how the shared probabilistic RUL information defined in Chapter 3 is mapped into the optimization model through a daily discretisation of failure probabilities. We also introduce an opportunistic maintenance strategy, demonstrating how the model effectively groups maintenance tasks to optimize resource utilization, particularly vessel selection. Then, we detail the maintenance model used, explaining how the state of each component is updated following different maintenance actions, the calculation of the expected lifetime of the components and the expected maintenance costs, considering material expenses, labor, downtime, and transportation. Finally, we present the mathematical formulation of our optimization model, incorporating constraints related to logistics, resource availability, and real-time prognostic updates, thus bridging tactical planning and short-term execution.

4.3.1 RUL prognostics input to the MILP

The prognostic information used by the MILP is not generated within this chapter, but inherited directly from the environment defined in Chapter 3. More specifically, the optimization model receives, for each critical component and each decision epoch, a probabilistic representation of

Remaining Useful Life (RUL) consisting of: (i) a point prediction, and (ii) an uncertainty-aware predictive distribution. In this way, the mathematical optimization framework is evaluated on the same prognostic information interface later used by the learning-based models, ensuring controlled comparison across decision-support paradigms.

As defined in Section 3.5, each component lifetime is generated from a component-specific Weibull distribution, and prognostic uncertainty is represented through a time-varying standard deviation that contracts as more condition information becomes available. The resulting point prediction is converted into a truncated predictive distribution over non-negative time-to-failure values. This shared prognostic layer is fully reproducible under fixed random seeds and can be replaced in future applications by real condition-monitoring and prognostic data without modifying the downstream optimization structure.

For the purpose of the MILP, the continuous predictive distribution must be mapped onto the daily time grid used by the tactical planner. Let T_f denote the random time-to-failure of a component at decision epoch k , with predictive density $\phi_k(t)$ and cumulative distribution function $F_k(t)$. Since the optimization model operates over a finite horizon of H days, the continuous prognostic belief is discretised into day-wise failure probabilities. For each day bin $i \in \{0, \dots, H-1\}$, the corresponding probability mass is defined as

$$p_k(i) = \mathbb{P}(T_f \in [i, i+1)) = \int_i^{i+1} \phi_k(t) dt = F_k(i+1) - F_k(i). \quad (4.1)$$

This discretisation provides the probability that the component fails during day i of the planning horizon, conditional on the information available at epoch k . Using the discrete mass function $\{p_k(i)\}_{i=0}^{H-1}$, the model can evaluate the expected cost of delaying maintenance, the survival probability up to a candidate maintenance day, and the expected downtime incurred if failure occurs before the planned intervention. In this sense, the optimization does not rely on a single deterministic RUL value, but on a risk-aware representation of failure timing that is consistent with the probabilistic benchmark environment introduced in Chapter 3.

To simplify notation in the following subsections, the predictive distribution at decision epoch k is denoted by ϕ_k , and its discretised day-wise probability mass function by $p_k(i)$. All expected-value calculations appearing in the maintenance and optimization model are based on this discrete-time representation. Accordingly, the role of prognostics in the MILP is to provide a compact probabilistic summary of component health that can be embedded directly into the tactical planning problem without redefining the underlying prognostic model in each chapter.

4.3.2 Opportunistic maintenance

In the proposed model, an opportunistic maintenance strategy is defined to save on the use of vessels by combining multiple tasks into a single maintenance cycle. The selection of vessels is based on their suitability for different tasks and the expected weather conditions (Li et al., 2024) and will influence the expected cost of the maintenance cycle.

A detailed approach to the selection of the appropriate vessel can be found in Dalgic et al. (2015b) where SOVs and CTVs have been characterized based on their capacity and seakeeping abilities. Here, following a similar logic, vessel types are ranked from 1 to 4 based on their ability to carry out different maintenance tasks, as shown in Table 4.1.

Table 4.1: Vessel Hierarchy for Maintenance Tasks

Vessel Type	Rank	Description
CTV (Crew Transfer Vessel)	I	Suitable for most minor maintenance tasks.
SOV (Service Operation Vessel)	II	Capable of handling major maintenance activities and more extensive tasks compared to CTV.
Jack-up Vessel	III	Used for more intensive maintenance activities requiring significant equipment and stability.
Heavy Lift Vessel	IV	The vessel with the highest capacity to transport technicians, equipment, and spare parts, used for the most extensive maintenance and replacement tasks.

When maintenance activities are grouped, the vessel with the lowest rank that can perform all required tasks is chosen. This ensures that the chosen vessel has the necessary capacity to handle the combined workload.

The logic for selecting the appropriate vessel for grouped maintenance activities is summarized in Figure 4.1.

For example, if minor and major maintenance activities are grouped together, the model would select a vessel that can handle both types of tasks. Although for a minor intervention a CTV (Rank 1) might be sufficient, if a SOV (Rank 2) is needed for the major maintenance, the model will choose the SOV to ensure all tasks are completed efficiently.

The optimizer evaluates the trade-off between deploying multiple vessels, which leads to higher mobilization costs but shorter rental durations, versus utilizing a single, higher-capacity vessel for an extended period, potentially resulting in lower mobilization costs but increased rental expenses due to longer sea time.

4.3.3 Maintenance model

As highlighted in Dalgic et al. (2015b), different maintenance tasks are associated with different logistic requirements. Our model differentiates between minor, major repairs and replacements. To each maintenance type corresponds a certain lifetime extension, expressed as a percentage of the component's age at the time of repair. Let C be the set of components, A the set of actions [minor, major, replacement, Do Nothing], and T the planning-horizon length. To update the age of components that undergo maintenance, an age reduction coefficient is defined as follows:

$$\alpha_a = [0.3, 0.5], \forall a \in [\text{Minor}, \text{Major}] \quad (4.2)$$




Figure 4.1: Flowchart of vessel selection for grouped maintenance activities

So, the expected RUL of a component on which maintenance activity a is performed on day d can be computed as:

$$RUL_{\text{new}} = RUL_{\text{pred}} + d\alpha_a \quad (4.3)$$

In case of preventive or corrective replacements, a new component is installed in place of the old one, and its RUL predictions will not depend on the day on which maintenance is performed. The age reduction coefficients associated to different maintenance actions are taken from Li et al. (2022b).

The total cost of each maintenance activity includes material costs, labor costs, downtime costs, and transportation costs (Castro-Santos & Diaz-Casas, 2014).

We consider the cost of materials to be fixed and dependent only on the component and on the type of maintenance, thus, for any activity a performed on component c , there is an associated material cost $C_{\text{material}}(c, a)$.

Labor costs are calculated as follows:

$$C_{\text{tech}}(c, a) = N_{\text{tech}}(c, a) \times R_{\text{hour}} \times T(c, a) \quad (4.4)$$

where $N_{\text{tech}}(c, a)$ is the number of technicians needed for maintenance type a on component c , R_{hour} is the technicians' cost per hour, and $T(c, a)$ is the number of hours required for the task.

Downtime cost estimates are based on the expected lost production due to the turbine being non-operational. Because the proposed model targets tactical planning, we approximate power loss using a seasonal expectation rather than an hourly wind profile. In particular, the mean wind speed is treated as month-dependent and computed from the metocean dataset described in the Case Study.

$$C_{\text{dt}}(T_{\text{dt}}, k) = P(\bar{v}_{m(k)}) C_{\text{el}} T_{\text{dt}}, \quad (4.5)$$

where $m(k) \in \{1, \dots, 12\}$ denotes the calendar month associated with day k , $\bar{v}_{m(k)}$ is the corresponding monthly average wind speed at the site, $P(\bar{v}_{m(k)})$ is the turbine power curve evaluated at that monthly average, C_{el} is the electricity price, and T_{dt} is the downtime duration (hours) during which the turbine is not producing energy due to failures or maintenance.

The cost related to the utilization of vessels is computed as the sum of mobilization cost, fuel consumption and daily rate:

$$C_{\text{trans}}^j(T(c, a)) = C_{\text{mobil}}^j + F_{\text{cons}}^j \cdot 2D_{\text{travel}} + R_{\text{day}}^j \cdot \frac{T(c, a)}{24} \quad (4.6)$$

where C_{mobil}^j is the mobilization cost of vessel j , F_{cons}^j is the fuel consumption rate of vessel j , $2D_{\text{travel}}$ is the round-trip sailing distance, R_{day}^j is the daily cost of vessel j , and $T(c, a)/24$ is the time (in days) required to perform maintenance activity a on component c .

Finally, the total cost of maintenance a on component c can be expressed as:

$$C(c, a) = C_{\text{material}} + C_{\text{tech}}(c, a) + C_{\text{dt}}(T) + C_{\text{trans}}^j(T(c, a)) \quad (4.7)$$

Now that the cost items have been described, the next paragraphs show how they are used to compute the expected cost of initiating any maintenance action a on any component c on any day d of the planning horizon.

The expected cost of each possible maintenance action is computed considering the cumulative probability of the component surviving until the day of predictive maintenance d or failing beforehand, incurring into corrective replacement and potential downtime costs.

$$\begin{aligned} \mathbb{E}_a[\text{Cost}(k, d)] &= C(c, a) \sum_{i=0}^{d-1} \tilde{p}_k(i) \\ &\quad + C(c, a_{\text{replacement}}) \left(1 - \sum_{i=0}^{d-1} \tilde{p}_k(i) \right) \\ &\quad + C_{\text{dt}} \left(\sum_{i=0}^{d-1} \tilde{p}_k(i) (d - i - 1) \right), \quad \forall a \in \mathcal{A}. \end{aligned} \quad (4.8)$$

where $\mathcal{A} = \{\text{replacement}, \text{major}, \text{minor}, \text{DoNothing}\}$.

The expected cost includes the cost of the specific maintenance action $C(c, a)$, the cost of corrective replacement $C(c, a_{\text{replacement}})$ we would incur in if the component fails beforehand, and the downtime cost C_{dt} .

An additional cost is considered for minor and major repairs to acknowledge future expenses when the component will eventually be replaced. A discount factor (DF) is introduced to account for the time value of money considering that future costs are less valuable than current expenses due to potential interest earnings, as used by Sandborn et al. (2014) or by Ioannou et al. (2018) in their techno-economic model for offshore wind life-cycle.

Let η_a represent the expected RUL of the component in case of it undergoing maintenance a , converted into years:

$$\eta_a = \frac{\mathbb{E}_a[\text{Life}(k, d)] - T_{\text{elapsed}}}{365} \quad (4.9)$$

We can then use this value to calculate the discount factor at the time of the new expected failure for the component:

$$\text{DF}_a = (1 - \delta)^{\eta_a} \quad (4.10)$$

where δ is the discount rate.

In a similar fashion to how we calculated the expected cost, the expected lifetime of the components in case they undergo different maintenance tasks are computed using the cumulative probability of failure over the planning horizon. Maintenance activities impact the expected lifetime of components differently, this is characterized by the age reduction coefficient α .

For each maintenance type, the expected lifetime is estimated by considering whether the component fails before the planned maintenance day d , or survives until d and receives the scheduled intervention. For compactness, the survival probability up to day d is defined as

$$S_k(d) = 1 - \sum_{i=0}^{d-1} \tilde{p}_k(i), \quad (4.11)$$

where $\tilde{p}_k(i)$ is the discretised probability that the component fails during day i of the planning horizon, given the information available at decision epoch k .

For imperfect maintenance actions, namely minor and major repairs, the expected lifetime is computed as

$$\begin{aligned} \mathbb{E}_a[\text{Life}(k, d)] &= k + \sum_{i=0}^{d-1} i \tilde{p}_k(i) + d S_k(d) + \alpha_a(k+d) S_k(d), \\ \forall a &\in \{\text{major}, \text{minor}\}. \end{aligned} \quad (4.12)$$

The first term represents the current decision epoch. The summation accounts for the expected failure time if the component fails before the planned intervention. The term $d S_k(d)$ accounts for the case in which the component survives until the scheduled maintenance day. The final term represents the additional lifetime extension produced by imperfect maintenance, conditional on survival until day d . The coefficient α_a denotes the restorative effect of action a , with $\alpha_{\text{minor}} = 0.30$ for minor repair and $\alpha_{\text{major}} = 0.50$ for major repair.

In case of replacement, the task terminates the current component life and installs a new component. Therefore, the expected lifetime of the current component is only evaluated up to the planned replacement day d , or up to the earlier failure time if failure occurs before replacement:

$$\mathbb{E}_{\text{replacement}}[\text{Life}(k, d)] = k + \sum_{i=0}^{d-1} i \tilde{p}_k(i) + d S_k(d). \quad (4.13)$$

In case no maintenance is scheduled during the planning horizon ($a = \text{DoNothing}$), the expected lifetime is computed as the elapsed time plus the predicted RUL of the component:

$$\mathbb{E}_{\text{DN}}[\text{Life}(k)] = k + \text{RUL}_{\text{pred}}(k). \quad (4.14)$$

The expected cost/lifetime ratio is a key indicator of the O&M performance, representing the yearly cost of maintenance for a specific component. This metric allows for a direct comparison of the cost-effectiveness of different maintenance plans.

$$\mathbb{E}_a[\text{Ratio}(c, k)] = \frac{\mathbb{E}_a[\text{Cost}(k, d)]}{\mathbb{E}_a[\text{Life}(k, d)]/365}, \quad \forall a \in \mathcal{A}. \quad (4.15)$$

where $\mathcal{A} = \{\text{replacement}, \text{major}, \text{minor}, \text{DoNothing}\}$.

4.3.4 Stochastic optimization formulation

The optimization model aims to identify the optimal scheduling of maintenance activities over a defined planning horizon that minimizes this cost/lifetime ratio.

The model uses a binary decision variable $x_{c,a,k}$ to indicate, for each component c , what maintenance activity a should be scheduled on day k .

$$x_{c,a,k} = \begin{cases} 1 & \text{if maintenance activity } a \text{ is planned} \\ & \text{for component } c \text{ on day } k \\ 0 & \text{otherwise} \end{cases} \quad (4.16)$$

Additionally, the binary variable $y(k)$ is introduced to define if a maintenance cycle should be initiated on day k of the planning horizon:

$$y(k) = \begin{cases} 1 & \text{if a maintenance cycle is to be initiated on day } k \\ 0 & \text{otherwise} \end{cases} \quad (4.17)$$

The objective functions in scheduling models typically account for maintenance expenses and system utilization. For instance, Chen et al. (2021) seek to minimize the expected cost over a component's projected lifespan, whereas other models concentrate exclusively on reducing maintenance costs as in Li et al. (2016). Considering the expected lifespan in the objective function allows the optimizer to penalize early interventions and over-maintenance; in other studies, such as Jain et al. (2021), this is done by promoting a trade-off between corrective maintenance and lost RUL due to preventive maintenance.

In our case, the main goal of the model is to minimize the yearly maintenance cost and is formulated as follows:

$$\begin{aligned} \min Z = & \sum_{c \in C} \sum_{k=0}^{T-1} \sum_{a \in A} y(k) x_{c,a,k} \mathbb{E}_a[\text{Ratio}(c, k)] \\ & + \sum_{c \in C} \sum_{k=0}^{T-1} \sum_{a \in A} (1 - x_{c,a,k}) \mathbb{E}_{\text{DN}}[\text{Ratio}(c, k)]. \end{aligned} \quad (4.18)$$

where $\mathbb{E}_a[\text{Ratio}(c, k)]$ is the expected cost-over-lifetime ratio of performing maintenance action a on component c on day k , and $\mathbb{E}_{\text{DN}}[\text{Ratio}(c, k)]$ is the corresponding ratio when no maintenance action is planned.

Logical and operational constraints are applied to ensure feasibility of the solutions generated by the model.

$$\sum_{a \in A} x_{c,a,k} \leq 1 \quad \forall c \in C \quad (4.19)$$

$$\sum_{k=0}^{T-1} y(k) \leq 1 \quad (4.20)$$

$$\sum_{a \in A} x_{c,a,k} \leq \sum_{k=0}^{T-1} y(k) \quad (4.21)$$

$$M(k) = \begin{cases} 1 & \text{if technicians are available on day } k \\ 0 & \text{otherwise} \end{cases} \quad (4.22)$$

$$S(k) = \begin{cases} 1 & \text{if spare parts are available on day } k \\ 0 & \text{otherwise} \end{cases} \quad (4.23)$$

$$W(k) = \begin{cases} 1 & \text{if enough operational hours on day } k \\ 0 & \text{otherwise} \end{cases} \quad (4.24)$$

$$y(k) \leq M(k) \quad \forall c \in C, \forall k \in \{0, \dots, T-1\}, \forall a \in A \quad (4.25)$$

$$y(k) \leq S(k) \quad \forall c \in C, \forall k \in \{0, \dots, T-1\}, \forall a \in A \quad (4.26)$$

$$y(k) \leq W(k) \quad \forall c \in C, \forall k \in \{0, \dots, T-1\}, \forall a \in A \quad (4.27)$$

$$\sum_{a \in A} x_{c,a,k} \geq f_c(d) \quad (4.28)$$

Each component can undergo at most one maintenance action within the planning horizon, this is taken into account in Equation 4.19. Using Equation 4.20, we can ensure that maintenance is initiated only once during the planning horizon. Additionally, Equation 4.21 is used to ensure that a task is assigned to a component only if a maintenance cycle has been planned. Furthermore, maintenance can only be scheduled if there are enough technicians, spare parts, and if the weather conditions are favorable, considering the vessels' operational limits; the binary variables M_k , S_k , and W_k are thus defined in Equations 4.22, 4.23 and 4.24 to indicate whether technicians, spare parts, and operational hours are available on day k respectively. Constraints 4.25, 4.26 and 4.27 ensure that maintenance is scheduled only if all necessary resources are available. Finally, statement 4.28 ensures that, if a component has already failed at the time of observation d , a maintenance cycle has to be initiated as soon as possible, here $f_c(d)$ is a binary variable indicating whether component c is failed on day d :

$$f_c(d) = \begin{cases} 1 & \text{if component } c \text{ is failed on day } d \\ 0 & \text{otherwise} \end{cases} \quad (4.29)$$

Various solution methods have been employed in the literature to optimize the expected cost of repair considering failure probability. For example, exhaustive search methods, despite their high computational demands, give exact solutions, as demonstrated in Vianna & Yoneyama (2018). These methods can be particularly useful for simpler systems to use as benchmarks for more complex models, as illustrated in Verleijdonk et al. (2024).

In this chapter, similarly to Zhuang et al. (2023); Chen et al. (2021) and Camci et al. (2019), we use Mixed Integer Linear Programming (MILP) with a rolling-horizon approach, as utilized in Mitici et al. (2023) and Zhuang et al. (2023) to revise decisions as new data becomes available, an essential feature given the continuous updates provided by prognostic information.

4.3.5 MPC formulation

We formulate maintenance scheduling as a receding-horizon control problem with partial observability of component health. At each decision epoch k , the controller observes an information state and computes a plan over a finite prediction horizon H (days), but only implements the first-step decision before re-optimizing.

Let the underlying (unobserved) physical state be $x_k = \{RUL_{\text{true},c}(k)\}_{c \in C}$. The controller operates on the belief state

$$b_k = \{\mu_{c,k}, \sigma_{c,k}\}_{c \in C}, \quad (4.30)$$

where $(\mu_{c,k}, \sigma_{c,k})$ parameterize the probabilistic RUL belief for each component.

Given the current belief b_k , the optimizer predicts future beliefs by propagating: (i) component health forward in time (RUL decreases by one day if no intervention occurs), and (ii) uncertainty contraction according to (4.36). The resulting failure-time pmf $\tilde{p}_{c,k}(i)$ is computed and used inside the MILP objective.

At epoch k , we solve the MILP over days $k, \dots, k+H-1$ and apply the immediate decision of initiating a cycle or not. Then, at $k+\Delta k$ new prognostic measurements are assimilated, the belief is updated, and the MILP is re-solved. This repeated re-optimization constitutes the closed-loop MPC feedback mechanism.

If no maintenance is performed, the true RUL decreases deterministically:

$$RUL_{\text{true}}(k+\Delta k) = RUL_{\text{true}}(k) - \Delta k, \quad (4.31)$$

while if maintenance is carried out, RUL_{true} is increased according to Equation 4.3. The predicted RUL is propagated in parallel. At epoch $k+\Delta k$, a new prognostic observation is generated:

$$RUL_{\text{pred}}(k+\Delta k) = RUL_{\text{true}}(k+\Delta k) + e(k+\Delta k). \quad (4.32)$$

$$e(k+\Delta k) \sim \mathcal{N}(0, std^2(k+\Delta k)) \quad (4.33)$$

The range of possible standard deviations is defined, as in Chapter 3, by:

$$std_{\text{max}} = RUL_{\text{true}} \cdot (\text{large uncertainty}) \quad (4.34)$$

$$std_{\text{min}} = RUL_{\text{true}} \cdot (\text{small uncertainty}), \quad (4.35)$$

Where std_{max} represents the initial uncertainty in the RUL prediction and std_{min} the reduced uncertainty after new data has been accumulated. The confidence in the prediction increases with time, converging towards the lower bound of uncertainty. Here, large uncertainty denotes the conservative variance assumed when only limited condition data are available, while small uncertainty reflects the narrower variance achieved once more observations have refined the estimate. The specific values of *large* and *small uncertainty* can be chosen according to the quality of the prognostic system being simulated, allowing different levels of prediction accuracy to be represented. The transition in time from large to low uncertainty is modeled as:

$$std(k) = std_{\max} - (std_{\max} - std_{\min}) \left(\frac{1}{1 + \exp\left(-\frac{k - aRUL_{\text{true}}}{bRUL_{\text{true}}}\right)} \right) \quad (4.36)$$

where $std(k)$ is the time-dependent standard deviation of the prediction, and parameters a and b shape how quickly the uncertainty contracts.

The MILP then minimizes the expected cost/lifetime ratio using the new predicted RUL together with logistical constraints and the cycle repeats with the updated beliefs.

This iterative cycle of (i) propagating the prior, (ii) updating it with the new prognostic observation, and (iii) re-optimizing decisions constitutes the feedback mechanism that characterizes our closed-loop MPC formulation.

Figure 4.2 summarizes the overall workflow of the proposed predictive maintenance approach, from prognostic inputs and offshore accessibility information to rolling-horizon optimization, execution, and feedback-based belief updates.




Figure 4.2: Framework of the proposed prognostic-driven maintenance approach.

4.3.6 Diagnosis of rolling-horizon limitations

The rolling-horizon MILP introduced in the previous subsection determines maintenance actions over a finite planning horizon and is then re-solved as new information becomes available. While this structure is consistent with Model Predictive Control principles and is computationally tractable for tactical offshore O&M planning, it also introduces a well-known structural limitation: decisions are only optimized with respect to the current look-ahead window. As a result, economically relevant actions that lie just beyond the end of the active horizon may be ignored, even when they could have been grouped with the current maintenance cycle at lower marginal cost.

Two specific concerns motivated the additional analysis introduced here. First, the finite-horizon MILP may postpone maintenance actions that would become attractive immediately after the current planning window, thereby missing grouping opportunities that are visible only from the perspective of the next optimization window. Second, the same horizon truncation may lead to preventive replacements being scheduled too late in the lifetime of the plant, i.e. at a point where the remaining operational life before decommissioning is too short for the intervention to be economically justified.

To diagnose these shortcomings, an oracle-aided variant of the model was implemented. The oracle is not intended as a realistic decision-support policy, since it uses information that is not available to the planner at decision time. Instead, it is used as an analytical benchmark to quantify the extent to which the finite-horizon MILP is affected by horizon myopia. The comparison is therefore not between two equally deployable planners, but between the original MILP and a “powered-up” version that is allowed to inspect the consequences of current decisions beyond the active optimization horizon.

Let the current optimization window start at time τ and span a planning horizon H , so that the active MILP optimizes decisions over the interval $[\tau, \tau + H]$. After solving the standard MILP in the current window, a shadow MILP is solved from the start of the next window, i.e. from $\tau + H$, using the state that would result if the current-window decisions were implemented. The shadow MILP is restricted to an early look-ahead interval of length L . Accordingly, the shadow problem identifies components that the optimization model would schedule in the interval

$$[\tau + H + 1, \tau + H + L].$$

A component is classified as a *missed grouping opportunity* if it satisfies two conditions:

1. it is not scheduled by the current-window MILP over $[\tau, \tau + H]$;
2. it is scheduled by the shadow MILP within the first L days after the current horizon.

Formally, let $\mathcal{S}_\tau$ denote the set of components scheduled by the main MILP in the current window, and let $\mathcal{S}_{\tau+H,L}^{\text{shadow}}$ denote the set of components scheduled by the shadow MILP in the interval $[\tau + H + 1, \tau + H + L]$. The set of missed grouping opportunities is defined as

$$\mathcal{M}_\tau = \mathcal{S}_{\tau+H,L}^{\text{shadow}} \setminus \mathcal{S}_\tau. \quad (4.37)$$

The rationale is that every component in $\mathcal{M}_\tau$ is one that the model itself would want to maintain shortly after the horizon boundary, but that was not selected in the current cycle even though grouping might have been beneficial.

In the oracle-aided policy, each component in $\mathcal{M}_\tau$ is pulled back into the current maintenance cycle and assigned to the same maintenance day selected by the main MILP in the current window. The action applied to component $c \in \mathcal{M}_\tau$ is the action with minimum expected cost-to-lifetime ratio on that already selected day. If $\mathcal{A} = \{\text{minor}, \text{major}, \text{replacement}\}$ denotes the admissible preventive action set and $R_{c,a}(k)$ the expected cost-to-lifetime ratio of applying action a to component c on day k of the current window, then the oracle selects

$$a_{c,\tau}^* = \arg \min_{a \in \mathcal{A}} R_{c,a}(k_\tau^*), \quad (4.38)$$

where k_τ^* is the maintenance day already chosen by the current-window MILP. In this way, the oracle does not introduce an additional maintenance cycle; rather, it augments the cycle already selected by the optimizer.

To make the oracle analysis reproducible, the system state at the start of the next window must be generated deterministically from the current state and the decisions taken in the current window. For each scheduled component, the post-maintenance state is computed using the same virtual-age logic adopted in Chapter 3 and in the main MILP simulation loop. Unscheduled components simply age by H days. The resulting state is then used as the initial condition for the shadow MILP. This procedure ensures that the oracle diagnosis is fully consistent with the maintenance-effect model used throughout the chapter and does not introduce an alternative degradation logic. The implementation retains the relative-RUL representation across all windows and fixes the random seed to preserve comparability across Monte Carlo runs.

A second indicator is introduced to diagnose late-horizon replacement behaviour close to plant decommissioning. Let T_{EOL} denote the plant end-of-life. If the optimizer selects a preventive action for a component whose failure time would occur beyond T_{EOL} , then the economic benefit of the intervention may be questionable, since the life extension produced by the action cannot be fully exploited before decommissioning. To capture this effect, an *end-of-life regret* event is recorded whenever a preventive action is scheduled for a component whose latent failure time satisfies

$$t_c^{\text{fail}} > T_{\text{EOL}}. \quad (4.39)$$

Two KPIs are then tracked: the number of such events and the total fixed maintenance cost associated with them. The first measures how often the rolling-horizon model selects potentially non-beneficial late-life interventions, while the second quantifies their direct economic magnitude.

The oracle experiment is evaluated using the same aggregate economic KPIs adopted for the main model, together with two additional diagnostics specifically designed to expose finite-horizon shortcomings:

- yearly total maintenance cost;
- yearly downtime cost;
- missed grouping opportunities (count);
- end-of-life regrets (count);
- end-of-life regret cost.

The missed-grouping KPI captures horizon-induced failure to exploit current maintenance cycles, while the end-of-life regret KPI captures horizon-induced tendency to postpone or mis-time replacements close to decommissioning. Together, these measures provide a structured way to quantify how much of the performance gap between the original MILP and the oracle-aided variant is attributable to rolling-horizon myopia rather than to the underlying cost model itself.

In the next section, we describe the case scenario used to benchmark the results of the model against the defined oracle and an existing strategy.

4.4 Case study

The Case Study was chosen specifically to verify the results of the proposed methodology against the Multiple Age-Based Opportunity (MABO) strategy described in Li et al. (2021) so that the results could be compared and verified.

Unless otherwise stated, the numerical parameterisation adopted in this chapter follows the generic offshore wind farm benchmark introduced by Li et al. (2021). This includes the wind farm configuration, turbine specifications, component reliability parameters, and replacement costs. Weather conditions are generated from the historical metocean dataset presented by Li et al. (2015). More specifically, the wind speed and significant wave height distributions correspond to the selected North Sea location in that dataset, from which the synthetic weather time series and the mean wind-speed statistics used in this chapter are derived. Economic parameters that are not explicitly reported in the benchmark are adopted as representative scenario assumptions to define a consistent reference case.

We consider a generic offshore wind farm in the North Sea, located 25 km from an onshore base in the Netherlands. The system comprises 50 generic wind turbines with a rated power of 3 MW. The turbine specifications adopted from Li et al. (2021) are summarised in Table 4.2.

Each turbine is represented as a series system composed of five critical components: blades, main bearing, gearbox, generator, and pitch system. Failure of any component renders the turbine unavailable. The component-specific Weibull lifetime parameters and preventive and corrective replacement costs are adopted from the benchmark developed by Li et al. (2021) and are reported in Table 4.3.

The remaining operational parameters required by the optimization model are summarized in Table 4.4. The mean wind speed is computed from the historical wind-speed time series corresponding to the selected North Sea location in the metocean dataset of Li et al. (2015). The electricity price, distance from shore, and mobilization cost follow the benchmark configuration

Table 4.2: Turbine Specifications (Li et al., 2021)

Specification	Value
Rated Power Output (MW)	3
Rotor Diameter (m)	90
Hub Height (m)	80
Cut-in wind speed (m/s)	3
Cut-out Wind Speed (m/s)	25
Rated Wind Speed (m/s)	12

Table 4.3: Failure distribution and cost parameters for critical components

Component	Shape	Scale	Failure Replacement Cost [kEUR]	Preventive Replacement Cost [kEUR]
Blades	3	1847	215	55
Bearings	2	1811	60	15
Gearbox	3	1477	260	65
Generator	2	1594	90	25
Pitch	3	1144	46	10

Table 4.4: Operational Parameters

Parameter	Value
Mean Wind Speed at 10m (m/s)	8.58
Cost of Electricity (EUR/MWh)	150
Distance from Shore (km)	25
Hourly Rate for Technicians (EUR/h)	75
Mobilization cost (kEUR)	50

of Li et al. (2021). The technician hourly rate and vessel operating cost define fixed economic assumptions adopted consistently throughout the numerical experiments.

Long-term weather measurements are essential to evaluate maintenance strategies under realistic variability in accessibility and waiting times. This Chapter uses the long-term hindcast database compiled by Li et al. (2015) for five European offshore sites. Their dataset is based on simultaneous hourly wind and wave hindcast data from 2001 to 2010 and provides joint information on mean wind speed at 10 m height and significant wave height, among other parameters. This multi-year coverage enables robust characterization of seasonal patterns and inter-annual variability, which is critical for benchmarking O&M decision models.

This study requires long time series of metocean conditions to benchmark maintenance decision policies under realistic access variability, including seasonal persistence of “good” and “bad” weather windows. We implemented a weather generator that preserves the empirical temporal structure of accessibility through a month-dependent Markov process. Rather than generating continuous wind and wave time series, the generator produces vessel-specific hourly operability states (accessible vs. non-accessible) derived from historical wind speed and significant wave height records. This directly provides the binary accessibility indicators used in Equation 4.27 as constraints for the optimization.

To obtain long accessibility time series while preserving realistic seasonal patterns, we generate synthetic vessel-specific operability sequences using a month-conditioned two-state Markov process calibrated from the hindcast-derived accessibility data. Concretely, the hourly hindcast observations are first thresholded using the vessel operability limits in Table 4.5, yielding an empirical binary sequence (accessible / non-accessible) for each vessel class. Then, for each vessel class, a separate two-state Markov model is estimated for each calendar month, so that the probability of remaining in (or switching between) accessible and non-accessible states reflects the observed persistence of weather windows in that month (e.g., longer storm regimes in winter versus more frequent short interruptions in summer). By simulating these month-dependent Markov models over multiple years with fixed random seeds, we produce long, reproducible accessibility series that retain the temporal structure of access constraints. The scheduling model uses these vessel-specific binary accessibility indicators (aggregated to the daily decision grid) as exogenous feasibility inputs for maintenance execution.

The adopted approach is intentionally decision-centric: it generates the binary accessibility signal required by the O&M decision models, rather than aiming to reconstruct the full joint distribution of wind and wave processes. Representative threshold values for each considered transportation mode $(\bar{U}_v, \bar{H}_{s,v})$ are reported in Table 4.5, collected and adapted from Donnelly et al. (2024) and Hu (2019).

Therefore, in the model, only two parameters are retained as exogenous drivers of marine operability: significant wave height (H_s) and wind speed at 10 m (U_{10}), consistent with the key variables required to determine access feasibility and, where relevant, to approximate production conditions. For the purpose of this chapter, vessels and spare parts are assumed to be available whenever a maintenance action is scheduled. This assumption is justified by the planning structure adopted in the modelling environment, where maintenance activities must be planned at least 30 days in advance. This lead time is assumed to be sufficient to secure the required access asset and spare parts, so that vessel unavailability and spare-part stockouts are not modelled explicitly. As a result, weather-driven accessibility remains the only time-varying operational constraint considered in the benchmark.

Once the case scenario is set, by adapting our model to utilize the same scheduling philosophy of Li et al. (2021), we can verify the results. A brief description of the age-based strategy proposed in the above-mentioned paper is therefore provided: in this case, maintenance cycles will be initiated once a component fails, at that moment, the age consumption of every other component will be computed as:

$$u_c(d) = \frac{d}{RUL_{\text{true}}} \quad (4.40)$$

Where u_c is the age consumption of component c , d is the elapsed time and RUL_{true} is the RUL of component c , randomly generated from the Weibull distribution. Therefore, having set age thresholds for preventive replacements, major and minor activities, every component whose age will have surpassed said thresholds will undergo the necessary action. This strategy, the selection of maintenance actions and its impact on the lifetime of the repaired component are summarized in Table 4.6.

Finally, in absence of real-life CM data, a synthetic dataset of probabilistic RUL prognosis was used, the model used to generate prognostic information is described in Chapter 3.

Table 4.5: Metocean operability limits for vessels, adapted from Donnelly et al. (2024) and Hu (2019).

Asset / vessel type	H_s limit (m)	U_{10} limit (m/s)	Notes (typical use)
CTV	1.5	12	Daily port-based technician transfer; boat landing; limited payload. (Saraswati et al., 2017)
SOV	3.0	20	Offshore accommodation; walk-to-work transfer; enables longer campaigns. (Saraswati et al., 2017)
Jack-up Vessel (heavy-lift)	2.0	10	Major component replacement; stability and crane operations typically wind- and wave-limited. (Dewan & Stehly, 2017)
Helicopter	-	20	Not directly constrained by H_s ; primarily wind/visibility dependent; used for urgent transfer. (Saraswati et al., 2017)

Table 4.6: Age thresholds in MABO strategy, collected from Li et al. (2021)

Age consumption	Action	Age reduction
$u_c(d) \geq 1$	Failure Replacement	1
$u_c(d) \geq A_{max}$	Preventive Replacement	1
$A_{min} \leq u_c(d) < A_{max}$	Repair	0.7

By simulating this process over 15 years, we can take a look at the results of our model, focusing on the yearly cost of maintenance, averaged over 500 iterations. As shown in Table 4.7, the results of the two models differ by about 6%.

In the next section, we will compare the results obtained using the Multiple Age Based Opportunistic maintenance strategy against the prognostic-driven strategy proposed.

4.5 Results

In this section, we illustrate the results obtained using the prognostic-driven scheduling model, focusing on the yearly cost of maintenance. Firstly, we show how the proposed scheduling strategy outperforms age-based maintenance strategies, then, we analyze the sensitivity of our results with respect to the size of the wind farm. To evaluate the impact of prognostic-driven scheduling on the overall lifetime costs of a wind farm, we analyzed the yearly cost of O&M by simulating the scheduling process over 15 years for the OWF described in Section 4.4. The results indicated in this chapter are averaged over 100 iterations in order to let the mean value converge and improve robustness with respect to uncertainties in the RUL predictions and in the random generation of the true RUL of the components. The convergence of the average value of the yearly cost of maintenance is shown in Figure 4.3.

The proposed model results, thus, in a yearly cost of O&M of 2041kEUR, which represents a 8,72 % reduction with respect to the MABO strategy. This reduction is better highlighted in Figure 4.4.

The costs indicated in Figure 4.4 are categorized into repair costs, downtime costs and O&M costs which represent the sum of the first two. In order to test the sensitivity of our model with respect to the size of the Offshore Wind Farm, we applied the proposed model to 10, 20, and 50 wind turbines, the results are presented in Table 4.8.

To address (a) the dependence of rolling-horizon maintenance planning on the prediction horizon length and (b) the sensitivity to dominant offshore O&M drivers, we performed three one-factor-at-a-time studies: prediction horizon H , prognostic uncertainty floor (std_{min}/RUL_{true}), and (c) distance to port. Each scenario was simulated over 15 years and repeated over multiple Monte Carlo iterations. Figure 4.5a shows that the chosen prediction horizon has a strong impact on both yearly downtime cost and total O&M cost. Longer look-ahead enables the planner to anticipate upcoming

Table 4.7: Verification of the model

	MABO (Li et al., 2021)	Verification using same age thresholds	Difference [%]
Yearly cost of maintenance [kEUR/year]	2116	2236	5.6%




Figure 4.3: Yearly cost of maintenance after 100 iterations




Figure 4.4: Cost of repair, downtime and overall O&M cost using Prognostic-Driven Scheduling vs Age-Based strategy

Table 4.8: Sensitivity analysis with respect to the size of the Wind Farm

Wind Farm Size	Strategy	Yearly cost of repairs [$\frac{\text{kEUR}}{\text{year}}$]	Yearly cost of downtime [$\frac{\text{kEUR}}{\text{year}}$]	Yearly cost of O&M [$\frac{\text{kEUR}}{\text{year}}$]
10 WT	PDS	231	201	432
	MABO	229	239	468
	Difference	1%	-16%	-7.7%
20 WT	PDS	459	336	795
	MABO	469	380	849
	Difference	-2%	-11.6%	-6.4%
50 WT	PDS	1103	937	2041
	MABO	1030	1205	2236
	Difference	+7%	-22%	-8.7%

failures and schedule preventive actions earlier, which reduces exposure to costly corrective events and associated downtime.

We varied the minimum uncertainty floor used in the uncertainty-contraction model by setting $std_{\min} = \{0.01, 0.05, 0.1\} \cdot RUL_{\text{true}}$. Figure 4.5b shows that improved prognostic precision (smaller $std_{\min}$) decreases both downtime cost and total O&M cost, supporting the premise that higher-quality prognostics enable better-timed interventions and lower downtime exposure.

To quantify the influence of a key offshore logistics driver, we varied the distance to the onshore base as $D_{\text{travel}} \in \{10, 25, 50\}$ km. Figure 4.5c shows that increasing distance increases both the overall O&M cost and, more markedly, the downtime cost, consistent with longer travel times, higher transportation expenses, and delayed access.

To isolate the effect of finite-horizon myopia, the original rolling-horizon MILP was compared against the oracle-aided variant introduced in Section 4.3.6. The purpose of this comparison is not to claim that the oracle policy is implementable in practice, but to quantify how much performance is lost because the MILP only optimizes over the active planning window. Figure 4.6 illustrates the two structural issues targeted by this analysis. In the first case, a maintenance action that becomes attractive shortly after the end of the active horizon is not selected in the current cycle, even though it could have been grouped with the current intervention. In the second case, a replacement may be scheduled late in the plant lifetime, close to end-of-life, at a point where the economic return of the intervention is limited by the remaining operational horizon. These are not modeling errors in the MILP formulation itself, but consequences of the fact that a finite-horizon optimizer can only value what lies inside, or is imperfectly proxied near, the current decision window.

The comparison confirms that these effects are materially relevant. Relative to the baseline MILP, the oracle-aided variant yields a lower yearly total cost and a lower yearly downtime cost, while sharply reducing the number of missed grouping opportunities. The reduction in missed grouping opportunities is particularly informative, because it shows that part of the cost gap is not due to poor action ranking within the active window, but to the inability of the rolling-horizon model to internalize actions that lie just beyond the horizon boundary. In other words, the oracle improves the solution primarily by correcting horizon-placement effects rather than by redefining the underlying maintenance economics.

The end-of-life regret KPI provides complementary evidence. The baseline MILP records a non-negligible number of preventive interventions whose associated failure times lie beyond the plant end-of-life, together with a measurable regret cost. Under the oracle-aided variant, this KPI is reduced or eliminated, indicating that part of the baseline replacement behaviour is indeed distorted by horizon truncation. This result is consistent with the initial concern that, in a long simulation with repeated finite windows, the optimizer may defer replacement until a point where the intervention is still locally optimal within the window, but no longer globally meaningful over the remaining plant lifetime. As shown in Figure 4.7, the oracle-aided variant improves the economic KPIs while almost eliminating missed grouping opportunities, confirming that these events are a relevant source of inefficiency in the baseline rolling-horizon planner.

Taken together, these results support two conclusions. First, the rolling-horizon MILP captures the tactical value of prognostic-driven maintenance well, but it remains sensitive to horizon placement. Second, some of the remaining performance gap to a better-informed planner is specifically linked to missed grouping and late-life replacement effects.




Figure 4.5: Sensitivity analyses: yearly total O&M cost and yearly downtime cost for key tactical planning drivers.




Figure 4.6: Illustration of rolling-horizon limitations.




Figure 4.7: Comparison between the baseline MILP and the oracle-aided MILP.

4.6 Conclusions

This chapter addressed Research Question Q3, namely: *How can O&M planning be formulated as an optimization model that incorporates condition and prognostic information?* To answer this question, a mixed-integer linear programming framework was developed for tactical offshore wind maintenance planning, in which maintenance timing, intervention type, and campaign initiation are decided jointly under weather-driven accessibility, shared logistics resources, and uncertain component health information.

Building on the research gaps identified in Chapter 2 and the modeling environment established in Chapter 3, the proposed model translates probabilistic RUL information into explicit planning decisions through a rolling-horizon stochastic optimization. In this way, the chapter provides a concrete optimization-based formulation of prognostic-driven offshore wind O&M planning, moving beyond static age-threshold policies and showing how condition information, marine logistics, and accessibility constraints can be embedded within a single tactical decision-support model.

The numerical results show that this integration is operationally valuable. Over the selected case study, the prognostic-driven strategy reduces total yearly O&M cost relative to the age-based baseline, with the largest improvement observed for the 50-turbine case, where yearly O&M cost decreases by 8.7%, primarily due to a substantial reduction in downtime cost. This result is particularly relevant in the offshore context, where lost production and waiting time are shaped not only by the technical severity of failures, but also by restricted access windows and the cost of mobilising offshore resources. The sensitivity analyses further reinforce this interpretation: the performance of the model is influenced by the prediction horizon, the quality of prognostic information, and major offshore logistics drivers such as distance to port. Taken together, these results confirm that the value of prognostics in offshore wind lies not in prediction alone, but in enabling better maintenance timing under realistic operational constraints.

In addition to the comparison against benchmark heuristics, this chapter also employed an oracle-aided variant of the rolling-horizon MILP to diagnose the structural limitations introduced by the finite planning window. The purpose of this oracle was to quantify the extent to which the baseline MILP is affected by horizon myopia. More specifically, the oracle was used to identify maintenance actions that were not selected in the current optimization window, but that the model itself would schedule shortly after the horizon boundary in the next window, thereby revealing missed grouping opportunities. It was also used to expose cases in which preventive replacements were scheduled too late in the lifetime of the plant, i.e. when the remaining time before decommissioning was too short for the intervention to be economically justified. The results show that the oracle-aided variant reduces yearly total cost from approximately 1316 kEUR/year to 1294 kEUR/year, corresponding to a reduction of about 1.7%, and decreases yearly downtime by about 1.0%. More importantly, the oracle sharply reduces the diagnostic KPIs associated with horizon myopia: missed grouping opportunities decrease from 8 to 1 events, an 87.5% reduction, while end-of-life regrets and their associated regret cost are eliminated in this experiment. These results indicate that part of the remaining inefficiency of the baseline MILP does not stem from poor action ranking within the active horizon, but from the inability to fully account for economically relevant decisions just beyond the current optimization window.

At the same time, the chapter also makes visible the limits of the optimization paradigm. The MILP framework remains dependent on a finite planning horizon, a tractable representation of uncertainty, and simplifying assumptions required to preserve solvability. Although the rolling-horizon structure allows periodic replanning, uncertainty must still be compressed into forms that can be handled within a deterministic or scenario-based optimization model, and this becomes increasingly restrictive as the system grows in scale and operational complexity. The oracle analysis strengthens this point by showing that even when the model is solved optimally within each window, its decisions remain sensitive to horizon placement and to the truncation of future consequences.

The main contribution of this chapter is therefore twofold. First, it demonstrates that prognostic-driven tactical O&M planning can be formulated as a transparent and effective optimization problem under offshore-specific constraints. Second, it clearly identifies the main structural shortcomings of this approach, namely horizon myopia, missed grouping opportunities, and late-life replacement effects, and makes them measurable within the modeling environment of the thesis. These findings define the starting point for the next chapter, which explicitly studies methods to overcome these shortcomings. This provides the methodological basis for Chapter 5, where offshore wind O&M planning is formulated as a POMDP-based DRL problem.

Chapter 5

From Optimization to Learning-Based O&M Planning

This chapter addresses **Research Question Q4**: *Which modeling requirements should a learning-based offshore wind O&M framework satisfy to address the limitations identified in the optimization-based framework?*

Chapter 4 developed and evaluated the optimization-based planning paradigm of this thesis. The MILP framework showed that probabilistic RUL information can be translated into explicit maintenance decisions while accounting for weather-driven accessibility, logistics constraints, intervention costs, and downtime losses. The results also showed that this integration is operationally valuable, with the prognostic-driven strategy reducing yearly O&M cost relative to age-based maintenance. However, the same chapter diagnosed limitations associated with the rolling-horizon optimization structure, including horizon dependence, repeated re-optimization, uncertainty representation, scalability, and the representation of realistic maintenance action spaces.

The purpose of the present chapter is therefore not to re-identify the limitations of the MILP framework, but to translate them into modeling requirements for learning-based O&M planning. Deep Reinforcement Learning (DRL) is reviewed as a candidate paradigm because it is designed for sequential decision-making under uncertainty. Instead of solving a new optimization problem at each decision epoch, DRL learns a policy that maps observed system states to actions on the basis of long-term cumulative reward.

This chapter reviews DRL methods through this specific lens. It examines which learning-based formulations can support adaptive offshore wind O&M planning, how they represent uncertain component health and operational constraints, and which modeling choices are necessary to move from optimization-based scheduling to learning-based policy design. The analysis focuses on single-agent and multi-agent DRL methods, MDP and POMDP formulations, graph-based and hierarchical architectures, and the integration of domain-specific knowledge such as weather, logistics, prognostics, and economic cost structures.

The chapter is organised as follows. Section 5.1 links the limitations identified in Chapter 4 to learning-based modeling requirements. Section 5.2 reviews single-agent DRL approaches for maintenance decision-making. Section 5.3 extends the discussion to multi-agent frameworks. Section 5.4 analyses MDP, POMDP, graph-based, and hierarchical formulations. Section 5.5 discusses the integration of domain-specific knowledge in DRL-based O&M planning. Section 5.6

synthesises the benefits and limitations reported in the literature. Section 5.7 identifies multi-level maintenance actions as a key modeling gap. Finally, Section 5.8 concludes by deriving the implications for the DRL framework developed in Chapter 6.

This chapter is based on M. Borsotti, X. Jiang, and R. R. Negenborn: “Review of Deep Reinforcement Learning for Offshore Wind Farm Maintenance Planning”, published in *Wind Energy Science*, vol. 11, pp. 1185–1204, 2026, doi: 10.5194/wes-11-1185-2026.

5.1 Introduction

Offshore wind farm maintenance presents unique challenges due to harsh weather conditions, remote locations, and the intricate coordination of logistical resources (Musial et al., 2022). Storms, high winds, and unpredictable sea states create uncertainty in scheduling, while the limited accessibility of offshore sites further complicates intervention planning. Additionally, the need to allocate maintenance crews, vessels, and spare parts efficiently increases operational complexity, often leading to delays and higher costs. Maintenance strategies such as corrective or scheduled preventive, struggle to mitigate these challenges, often resulting in unplanned failures or unnecessary servicing (Fox et al., 2022).

As seen in Chapter 2, offshore wind O&M has traditionally been supported by deterministic or stochastic optimization models, rule-based policies, and predictive maintenance strategies informed by condition-monitoring and SCADA data. However, these approaches typically require predefined decision rules or restrictive modeling assumptions, limiting their ability to adapt to the dynamic and uncertain offshore environment.

In recent years, DRL has shown promising results as a data-driven approach to tackle these challenges. DRL is a class of algorithms that combine the sequential decision-making framework of reinforcement learning (RL) with the representational power of deep neural networks. In a typical RL setting, an agent interacts with an environment defined as a Markov Decision Process (MDP), observing a state s_t , selecting an action a_t according to a policy $\pi(a_t|s_t)$, and receiving a scalar reward r_t . The objective is to learn a policy that maximizes the expected cumulative discounted return $E\pi[\sum_t \gamma^t r_t]$, where $\gamma \in [0, 1]$ is the discount factor determining how future rewards are weighted relative to immediate ones (Sutton & Barto, 2018). Classical RL struggles when the state or action space is high-dimensional or continuous, as in offshore wind maintenance, where turbine states, weather, and logistics create vast combinations. DRL overcomes this limitation by using deep networks to approximate policy and value functions, enabling end-to-end learning directly from high-dimensional or partially observed inputs such as condition monitoring or weather data. These features make DRL particularly suitable for complex, stochastic decision problems in offshore wind O&M, where the agent must learn adaptive, long-horizon maintenance policies under uncertainty.

Figure 5.1 summarises four recurring challenges in offshore wind O&M, harsh and uncertain weather, unplanned failures, remote locations with limited accessibility, and complex logistics requiring resource coordination, and highlights four corresponding opportunities where data-driven decision support, and DRL in particular, can add value. First, *adaptive decision-making and learning* directly targets uncertainty: by learning policies that update actions based on newly observed information (e.g., condition-monitoring signals or revised weather forecasts), DRL can




Figure 5.1: Challenges of offshore wind maintenance and opportunities of Deep Reinforcement Learning methods

move beyond static decision rules and react to changing offshore conditions. Second, *proactive scheduling and resource allocation* addresses all challenges by exploiting forecasts and operational data to time interventions when access windows are likely (e.g., using metocean forecasts and, where available, operational sensing such as lidar-informed wind estimates) and to prioritise tasks before expected inaccessibility or risk escalation, thereby reducing weather-driven waiting time and avoidable downtime. Third, *data-driven optimization* provides a mechanism to jointly trade off competing objectives (e.g. cost, energy yield, risk, availability) under uncertainty, which is essential when failures and maintenance actions have long-horizon consequences. Finally, *scalable frameworks* respond to the growth in decision complexity as wind farms scale (more turbines, more components, more interacting constraints): function approximation, state abstraction (e.g., spatial or graph representations), and decentralised or hierarchical extensions provide pathways to maintain tractability when the state and decision spaces expand, increasing the complexity of logistics and resource coordination.

Despite these advantages, DRL also comes with well-recognised limitations that are particularly relevant for safety-critical infrastructure such as offshore wind. First, policies are usually represented by deep neural networks, which behave as black-box models and make it difficult for operators to understand or audit the rationale behind individual decisions; this lack of transparency is identified as a barrier to deployment and has motivated a dedicated line of work on explainable and safe reinforcement learning (Qing et al., 2022; Bui et al., 2024). Second, state-of-the-art DRL algorithms tend to be data- and computation-hungry, often requiring millions of interactions with an environment for training; this sample inefficiency and dependence on high-fidelity simulators are highlighted as key obstacles for applying DRL in real-world systems where experiments are costly or risky (Dulac-Arnold et al., 2021). Third, recent reviews of RL in power and energy systems emphasise the challenges of transferring policies from simulation to real assets, enforcing strict safety and reliability constraints, and encoding operational limits and human oversight in reward functions, all of which have so far limited large-scale industrial adoption (Pesántez et al., 2024; Bui et al., 2024).

A substantial body of literature has reviewed different aspects of data-driven O&M, but none provides a dedicated synthesis of Deep Reinforcement Learning (DRL) for maintenance planning. For instance, (Fox et al., 2022) review predictive and prescriptive O&M strategies, highlighting how

data-driven prognostics can support maintenance planning but without examining learning-based sequential decision frameworks. Tusar & Sarker (2022a) provide a systematic review of maintenance cost-minimisation models for offshore wind farms, focusing on optimization formulations and cost structures rather than adaptive control or online decision making. Similarly, reviews on machine-learning-based condition monitoring and prognostics, such as Stetco et al. (2019), Tautz-Weinert & Watson (2017), Pandit & Wang (2024), and Wang et al. (2026), synthesise diagnostic and RUL estimation techniques but do not address how such prognostic information interacts with maintenance scheduling policies.

Reinforcement learning itself has also been surveyed within the wind and power-systems domain. Narayanan (2023) reviews reinforcement-learning applications in wind energy, primarily in the context of control, forecasting, and wake steering. Abkar et al. (2023) survey RL approaches for wind-farm flow control, while Li et al. (2023) discuss DRL for modern power-system control problems, including frequency and voltage regulation. Yet, none of these reviews analyse O&M decision processes, nor do they compare DRL architectures, modeling assumptions, or their integration with offshore wind maintenance requirements.

To the best of our knowledge, this is the first review that focuses specifically on DRL for offshore wind-farm maintenance planning. In contrast to prior reviews, this work provides a focused synthesis of DRL approaches that have been proposed for, or are transferable to, offshore wind farm maintenance planning. Specifically, we (i) compare single- and multi-agent DRL frameworks and the algorithms they employ (value-based, policy-gradient, and actor-critic) in relation to the maintenance decision problems they target; (ii) analyse the problem formulations adopted in these studies, including MDP, POMDP, graph-based, and hierarchical representations, and how they encode uncertainty, PHM information, wake effects, weather and logistical constraints; and (iii) discuss the role of domain-specific knowledge and the remaining modeling gaps, with particular emphasis on the prevailing use of binary repair decisions instead of more realistic multi-level maintenance actions. By organising the review along these dimensions, we aim to clarify what current DRL models can and cannot do for offshore wind maintenance planning, and to outline research directions needed to move from promising simulation results toward practical adoption in real offshore wind operations.

Rather than attempting an exhaustive survey of all works in the DRL domain, we have deliberately narrowed our focus to a select group of key studies that exemplify the state-of-the-art in this area. To ensure transparency and reproducibility, the corpus analysed in this review was identified through a structured search process. Searches were conducted in Scopus, Web of Science, ScienceDirect, Researchgate, and Google Scholar using combinations of the terms: “offshore wind”, “maintenance”, “reinforcement learning”, “deep reinforcement learning”, “predictive maintenance”, and “O&M optimization”. The search window covered publications up to 2025. Inclusion criteria were: (i) studies proposing or evaluating DRL-based methods for maintenance or inspection planning of offshore or onshore wind systems; (ii) papers applying DRL to components or systems representative of wind turbine O&M; (iii) articles providing sufficient methodological detail to characterise the learning formulation. Exclusion criteria included: purely theoretical RL work, asset-management models without learning components, and high-level conceptual papers. This protocol yielded a total of 54 papers after full-text screening. The term “deliberately narrowed” refers to the intentional focus on works with explicit DRL formulations for maintenance decision-making, excluding broader O&M optimization domains (e.g., dispatch, routing, or power forecasting) unless they directly informed maintenance planning.

The primary goal of this chapter is to synthesize the advancements demonstrated by these representative models, critically assessing their strengths and identifying the limitations that still exist. By doing so, we aim to demonstrate how DRL can effectively address the challenges of offshore wind maintenance planning, while also pinpointing areas that require further refinement. This targeted review not only clarifies the current state of research in this field but also offers insights into future research directions, particularly in the development of decision frameworks that move beyond simplistic binary repair actions.

5.2 Single-agent DRL approaches for offshore wind O&M

Most DRL-based maintenance planners for offshore wind adopt a *single-agent* paradigm, where one agent learns an optimal policy for the entire system (e.g., a wind farm or a single turbine). This structure is suitable when a central decision-maker can coordinate all maintenance actions and information is aggregated at the farm or turbine level.

In DRL, two key functions define the learning objective: the *state–action value function* $Q(s, a)$, which estimates the expected cumulative reward of taking action a in state s , and the *state value function* $V(s)$, which measures the expected reward of being in state s and following the current policy. Different algorithm families approximate and use these functions in distinct ways, thus, among single-agent methods, three main families of DRL algorithms are most frequently applied:

- *Value-based*: value-based methods such as the Deep Q-Network (DQN) algorithm (Mnih et al., 2015), learn an approximation of $Q(s, a)$ using a deep neural network and select actions that maximize this value. Stability is achieved through experience replay and a target network. DQN and its variants (Double DQN, Dueling DQN) are effective for discrete maintenance decisions such as “*maintain*” versus “*not maintain*”.
- *Policy-gradient*: instead of estimating value functions, policy-gradient methods learn a parameterized policy $\pi_{\theta}(a|s)$ by adjusting the parameters θ in the direction of the performance gradient. Proximal Policy Optimization (PPO) (Schulman et al., 2017) is a widely used variant that constrains policy updates within a clipped trust region, improving stability and sample efficiency for large or continuous decision spaces, such as allocating multiple maintenance crews.
- *Actor–critic*: hybrid algorithms such as the actor-critic methods, combine both paradigms by maintaining a policy (the *actor*) and a value estimator (the *critic*). Deep Deterministic Policy Gradient (DDPG) (Lillicrap et al., 2016) and Soft Actor–Critic (SAC) (Haarnoja et al., 2018) extend DRL to continuous control using deterministic or stochastic policies with off-policy learning, while Asynchronous Advantage Actor–Critic (A3C) (Mnih et al., 2016) accelerates training through parallel environment instances.

The flowchart in Figure 5.2 illustrates these single-agent DRL approaches: all methods share common initial steps (environment setup, state observation) before diverging by learning logic: DQN (green) selects actions via ϵ -greedy exploration based on Q -values; policy-gradient methods (purple) sample actions from a learned distribution and update directly via performance gradients; actor–critic methods (orange) use a critic network to evaluate actions and guide policy improvement.

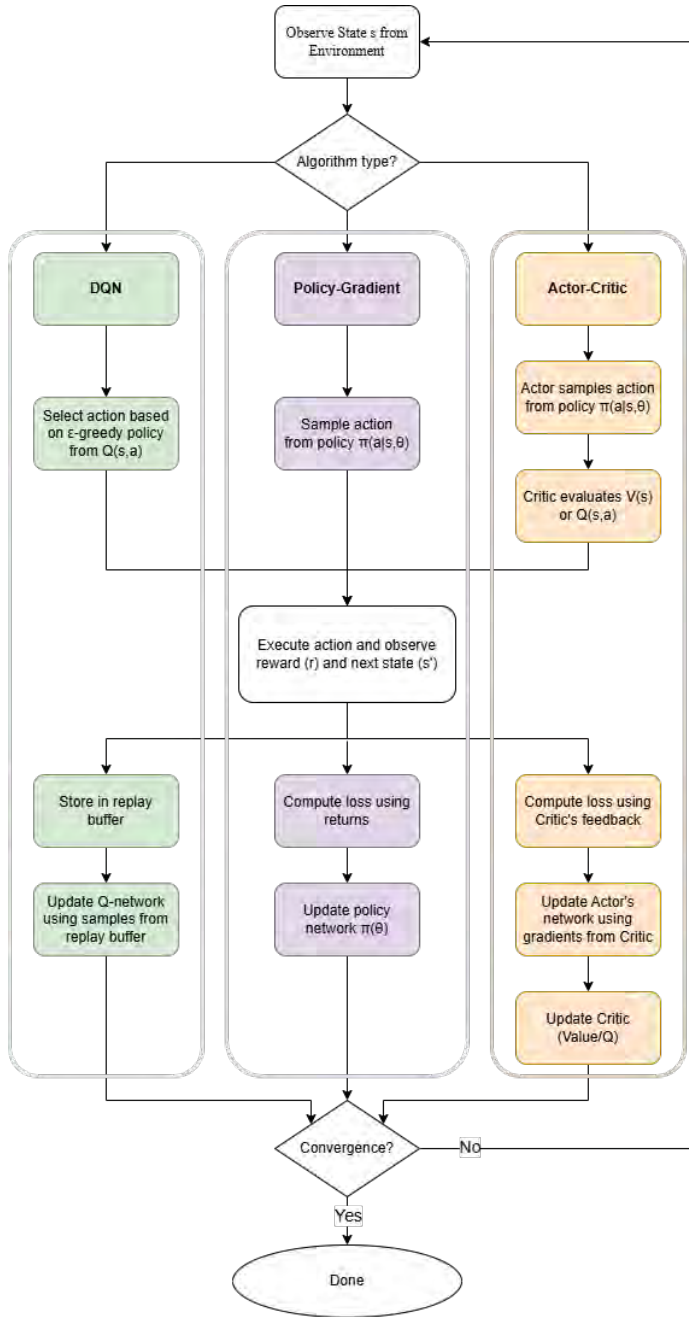


Figure 5.2: Overview of single-agent Deep Reinforcement Learning algorithm families: DQN (green), Policy-Gradient (purple) and Actor-Critic (orange).

The following subsections review how each family has been applied to specific O&M formulations and performance objectives.

5.2.1 Deep Q-networks

Value-based methods like DQN are popular for discrete maintenance decisions (e.g. whether to service a component now or later). For instance, Kerckamp et al. (2022) combined DQN with Graph Neural Networks to take into account asset topology. In their framework, a single agent uses a Graph Convolutional Network (GCN) to group maintenance actions on geographically proximate pipes, yielding more efficient schedules. The DQN+GCN approach produced more reliable networks and higher maintenance grouping compared to a plain DQN and to conventional preventive/corrective policies.

Similarly, Lee et al. (2025) developed a domain-informed DQN ensemble to schedule offshore wind farm maintenance tasks. They formulated maintenance scheduling as a MDP and incorporated wind wake effect models and weather variability into the state. By using convolutional layers to process spatial-temporal features (like turbine wake interactions), their DQN agent improved power generation by 11.1% compared to a baseline schedule.

These studies chose DQN for its stability in discrete action spaces and supplemented it with domain-informed neural architectures such as Convolutional Neural Networks (CNNs), or Graph Neural Networks (GCNs) to accelerate learning and capture dependencies. Another value-based approach is the Double DQN (DDQN) or Dueling DQN to handle large state spaces more stably, Zhang & Si (2020) tested DDQN for large state spaces in multi-component maintenance and showed better performance than simple threshold policies.

However, value-based methods can struggle when the action space or planning horizon grows large, or when the state is partially observed (e.g. uncertain component health) (Hausknecht & Stone, 2015). These challenges have led researchers to explore policy-gradient methods as well.

5.2.2 Policy gradient methods

Policy-based DRL algorithms directly optimize a parameterised policy $\pi(a|s)$ and are particularly advantageous in long-horizon, stochastic, or partially observable environments, which are characteristic of offshore maintenance planning. Because policy-gradient updates do not require explicit enumeration of all Q-values, these methods handle continuous or large multi-discrete action spaces more naturally than value-based approaches. They also tend to yield smoother and more stable optimization when rewards are sparse or delayed (Schulman et al., 2017), which helps explain the successful use of PPO in discrete maintenance-planning studies such as Pincirolì et al. (2021) and Cheng et al. (2023).

Pincirolì et al. (2021) developed a PPO-based agent to optimize maintenance dispatch for a wind farm with multiple crews. They formulated the problem as a sequential decision process and included prognostic information in the state (predicted RULs of turbines from PHM systems, and even forecasted power production for upcoming days). The PPO agent learned a policy that outperformed corrective, scheduled, and threshold-based predictive maintenance benchmarks in profit maximization. Notably, the DRL policy automatically scheduled maintenance during low-

power periods and anticipated failures using RUL predictions, something a simple RUL-threshold policy could not do. PPO effectively handles continuous decision-making under uncertainty and yields better long-horizon rewards than static policies.

In another example, Cheng et al. (2023), applied both DQN and PPO to learn cost-optimal condition-based maintenance for an offshore turbine component. They investigated policies with dynamic inspection intervals and adaptive repair thresholds, formulating the decision as an MDP and comparing DRL algorithms. Both a DQN agent and a PPO agent were able to discover optimal policies under varying conditions, outperforming fixed-interval or fixed-threshold strategies by reducing lifecycle costs although, in their case, PPO performed better than DQN.

5.2.3 Actor-critic and other methods

In problems with continuous action decisions (e.g. scheduling exact timing), actor-critic methods like deep deterministic policy gradient (DDPG) (Lillicrap et al., 2016) and Soft Actor-Critic (SAC) become relevant. In a broader manufacturing maintenance context, for example, A3C has been used to optimize resource allocation problems where quick convergence was needed (Mnih et al., 2016).

In the context of wind energy, Zhao & Zhou (2022) used a SAC algorithm to solve a maintenance scheduling problem formulated as an MDP. SAC's ability to handle continuous actions, maximizing the trade-off between reward and entropy, makes it suitable for maintenance problems requiring fine timing control. Their implementation allowed the agent to decide whether to perform maintenance on a turbine in each time period, aiming to minimize long-term cost. The SAC-based scheduler showed improved adaptability to random wind and failure events, outperforming greedy or periodic policies in simulation. Asynchronous Advantage Actor-Critic (A3C) and related algorithms have also been mentioned in the context of maintenance optimization to benefit from parallel training. While we are not aware of specific applications of A3C to offshore wind maintenance, the algorithm's ability to run multiple environment instances in parallel can accelerate training for complex simulations.

In principle, A3C/A2C could be applied to wind farm O&M planning to speed up learning across many simulated weather scenarios or failure realizations. However, stability can be an issue, so recent works have gravitated more to PPO for O&M problems.

To provide a clearer overview of how these families compare in terms of strengths, limitations and typical use cases in O&M decision-making, Table 5.1 summarises their key characteristics.

The discussion on single-agent DRL approaches has highlighted how tailored algorithms can effectively optimize maintenance decisions. However, as wind farms scale up and the complexity of interdependent maintenance decisions increases, a shift toward distributed decision-making might become necessary. The following section explores multi-agent DRL frameworks, which distribute the decision-making process across multiple agents, thereby addressing scalability challenges while enabling coordinated maintenance planning.

Table 5.1: Comparison of single-agent Deep Reinforcement Learning algorithm families for offshore wind Operations and Maintenance.

Algorithm Family	Strengths	Limitations
Value-based (DQN, DDQN, Dueling DQN)	Effective for low-dimensional discrete decisions Sample efficient with experience replay Simple and stable for small action spaces	Less stable in long-horizon problems Sensitive to partial observability Harder to scale to multi-discrete decisions
Policy-gradient (PPO)	Stable updates in long-horizon, stochastic settings Robust in partially observable environments Naturally handles multi-discrete decision elements	Less sample efficient than value-based methods Performance sensitive to reward shaping No inherent advantage in simple discrete tasks
Actor-critic (A2C/A3C, SAC, DDPG)	Combines policy stability with value-based guidance Efficient for complex states or long dependency chains	Training can be unstable without tuning Off-policy AC methods require careful replay-buffer design Computationally heavier than DQN variants

5.3 Multi-agent DRL in maintenance planning

As wind farms scale up and the environment increases in size, a multi-agent DRL (MADRL) approach can be considered to distribute the decision-making across multiple agents (Lowe et al., 2017) (e.g. one agent per turbine or per subsystem).


Figure 5.3 contrasts the single-agent and multi-agent DRL approaches. In the single-agent framework (5.3a, left), one centralized policy observes the environment and makes all maintenance decisions, whereas in the multi-agent framework (5.3b, right), each agent receives local observations and collectively coordinates actions.

Multi-agent frameworks address the curse of dimensionality that a single agent faces when managing many components simultaneously (Ogunfowora & Najjaran, 2023). In a cooperative MADRL setting, agents must learn policies that jointly optimize the overall maintenance outcome. A common approach is centralized training with decentralized execution: during learning the agents share information, but in execution each acts independently (Rashid et al., 2018).

A multi-agent perspective is taken by Andriotis and Papakonstantinou (2019), who developed a Deep Centralized Multi-Agent Actor-Critic (DCMAC) algorithm (Andriotis & Papakonstantinou, 2021). This algorithm treats each component as an “actor” with individual actions, and uses a centralized critic to evaluate the joint outcome. Their method allowed individualized component-level decisions (like a multi-agent system) but maintained a single value function for the overall



(a) Single-Agent Approach



(b) Multi-Agent Approach

Figure 5.3: Comparison of Single-Agent and Multi-Agent Deep Reinforcement Learning Approaches

system. This hybrid approach achieved strong results on high-dimensional maintenance problems, outperforming time-based and condition-based benchmarks. DCMAC can be seen as bridging single and multi-agent methods: it's centrally trained on the whole state, but action vectors are factorized per component. The success of DCMAC and similar centralized-training approaches again underlines that decomposing the action space among multiple decision-makers is a powerful strategy for scalability.

Another relevant study is Nguyen et al. (2022) (expanded by Do et al. (2024)), who optimized maintenance for a 13-component system using a Weighted QMIX algorithm. Here, each component is controlled by an agent, and a mixing network combines their action-values into a global Q-value to enforce cooperation.

The “weighted” QMIX (W-QMIX) variant addresses limitations of standard QMIX by improving credit assignment to each agent's actions (Liang et al., 2025). By customizing QMIX, they overcame the exponential growth of joint action-space and achieved cost-effective policies that significantly reduced total maintenance cost compared to independent or rule-based policies. In fact, their multi-agent policy outperformed a traditional threshold-based maintenance strategy, yielding 20% lower cost in the case study. The agents learned to coordinate, essentially performing opportunistic maintenance on other components when one component required service, something that is hard-coded in opportunistic heuristics but emerged naturally via learning. Notably, they simplified the architecture by using a branching neural network (one network outputting multi-component actions).

Figure 5.4 presents a high-level overview of how local Q-values from multiple agents are combined into a single joint Q-value through a central mixing network. At the top, global state observations and individual agents' local Q-values feed into the network. Within the mixing network, these local estimates are merged. The output at the bottom is a single joint Q-value, enabling decentralized agents to learn coordinated policies that account for global objectives.




Figure 5.4: Multi-agent mixing network architecture

While most offshore wind DRL studies to date use a single centralized agent, the multi-agent perspective is highly relevant. A wind farm can be seen as a team of turbines (or a team of maintenance crews) that could each be an agent. Multi-agent DRL can explicitly model interactions like shared resources (e.g. a vessel cannot fix two turbines at once) and learn decentralized policies. For example, one could have a DRL agent for each maintenance vessel coordinating via a mixing network to maximize farm availability. Nevertheless, multi-agent approaches introduce challenges such as non-stationarity during training and coordination complexity. To avoid these issues, for example, Pincirolì et al. (2021) effectively used a single PPO agent to dispatch multiple crews, which could be interpreted as a centralized multi-action policy rather than fully decentralized agents. An interesting future direction is to combine multi-agent RL with the physical layout of a wind farm, e.g. treating turbines as agents that learn when to request maintenance, or crews as agents learning which turbine to service, to further scale O&M optimization. In summary, multi-agent DRL is a promising direction that can address scalability and modularity in offshore wind maintenance, ensuring that solutions remain effective as the number of assets grows.

Having examined both single-agent and multi-agent DRL frameworks for maintenance planning, it is clear that the choice of algorithm is intricately linked to how the underlying problem is modeled. The next section focuses on the core methodologies and decision frameworks used, ranging from Markov Decision Processes (MDP) to Partially Observable MDPs (POMDPs), hierarchical and graph-based representations, that provide the foundation for these DRL approaches. This discussion will clarify how the characteristics of the problem formulation drive the design and performance of the DRL models.

5.4 Problem formulation

Formulating the maintenance planning problem correctly is a core requirement for DRL. In literature, we see formulations as Markov Decision Processes (MDPs), Partially Observable MDPs (POMDPs), and even graph-based representations, each chosen to capture the nature of the decision environment.

5.4.1 Markov decision process

Many works assume the system state is fully observable. For example, if one uses direct sensor readings or known component health states as the state, the decision process can be modeled as an MDP. The reward design typically combines negative maintenance costs and downtime losses into a single scalar reward (or profit) to guide the agent toward cost-optimal decisions, as in Pincirolì et al. (2021). The CBM decision can also be formulated as an MDP where the state might include the current damage level or time since last inspection as in Cheng et al. (2023). They defined actions such as whether to inspect or repair, with rewards based on cost. Similarly, the DQN-based scheduling by Lee et al. (2025) uses an MDP where the state includes turbine power outputs (affected by wake and weather) and maintenance statuses, assuming those are known to the agent. MDP formulations are simpler and allow use of standard DRL algorithms, but they rely on having a reliable estimator of the system's health state.

5.4.2 Partially observable Markov decision process

Several offshore wind O&M decision processes are more naturally modelled as POMDPs, because the agent typically does not observe the true underlying system state, but only noisy and delayed measurements (SCADA/CM signals, inspection outcomes, imperfect weather and access forecasts). In a POMDP, the agent maintains a belief to make decisions under uncertainty (Kaelbling et al., 1998).

Beyond SCADA-based monitoring, a complementary body of literature focuses on non-destructive evaluation (NDE) techniques that can provide inspection- and SHM-driven observations for wind-turbine components and structures (e.g., blades, tower, foundations). For instance, Civera & Surace (2022) review non-destructive techniques for wind-turbine condition and structural health monitoring over the last two decades, covering approaches such as visual inspection, acoustic emission, ultrasonic testing, infrared thermography, radiographic and electromagnetic methods, as well as oil monitoring, with deployments ranging from human inspections to robotic and UAV-based surveys. Such sensing and inspection modalities can enrich the observation space available to data-driven O&M decision support, but they also highlight why offshore maintenance planning is often partially observable in practice: measurements can be sparse in time (inspection-driven), noisy, and operationally constrained by access windows and logistics, motivating POMDP formulations and memory-/belief-based policy representations. In practical DRL implementations, three families of remedies are commonly used to mitigate partial observability and approximate belief-state reasoning:

- History-based state augmentation. A simple approach is to concatenate a fixed window of past observations and actions to the agent input, thereby providing short-term memory of recent dynamics. While straightforward, this can be insufficient when relevant dependencies extend over long horizons (e.g., degradation accumulation, delayed maintenance effects) (Kaelbling et al., 1998).
- Recurrent DRL (implicit belief-state via memory). Recurrent neural networks (typically LSTMs/GRUs) can compress the observation, action history into a hidden state that acts as an implicit belief surrogate. This idea has been adopted in value-based learning (e.g., Deep Recurrent Q-Networks, DRQN) (Hausknecht & Stone, 2015) and in actor-critic / policy-gradient settings, where recurrent policies are trained end-to-end to improve performance under partial observability (Williams, 1992; Schulman et al., 2017).
- Transformer-based memory (long-range dependencies). When long-range temporal structure matters, transformer architectures offer an alternative to RNNs by using attention mechanisms to retrieve relevant past information. Transformer-based RL agents have demonstrated improved stability and credit assignment in partially observable and long-horizon tasks by enabling flexible access to historical context (Parisotto et al., 2020). This is conceptually aligned with offshore maintenance planning, where optimal actions may depend on events far in the past (e.g., prior repairs, earlier inspections), although transformers may require careful regularisation and substantial training data.

Beyond implicit memory, explicit belief-state estimation can also be pursued by combining filtering with RL, for instance using Bayesian filters or particle filters when a tractable transition/observation model is available, or by learning latent-state models that infer hidden degradation

dynamics from observations before or during policy learning (Kaelbling et al., 1998; Igl et al., 2018). Overall, expanding offshore wind DRL formulations from MDP to POMDP settings primarily affects the policy representation: memory-augmented policies (recurrent or transformer-based) and/or explicit belief estimators provide practical mechanisms to cope with uncertainty in health information, accessibility, and delayed maintenance outcomes. A first example of such a formulation for an O&M planning problem is Kerckamp et al. (2022), who formulate maintenance planning as a POMDP on a graph, where the underlying deterioration states are partially observed through inspections. Similarly, Lee & Mitici (2023a) treat their predictive maintenance problem for aircraft components as a POMDP, using a CNN to process raw sensor data into an observation for the DRL agent. The choice of POMDP acknowledges that maintenance decisions must be made with imperfect information, and DRL agents in this setting are trained to be more robust to uncertainty (e.g. scheduling maintenance a bit earlier to hedge against uncertain failure times). Methodologically, solving a POMDP with DRL often means using belief state features or statistical features of uncertainty (such as predicted failure probability) in the state.

5.4.3 Graph and network-based representations

Some complex systems benefit from graph-based formulations. Kerckamp et al. (2022) is an example where the asset network structure (a sewer network in their case) is encoded as a graph, and a GCN is used to inform the DRL agent. While their domain was not wind, one can imagine an offshore wind farm graph where nodes are turbine components and edges could represent spatial proximity or electrical/functional dependencies. This approach could let the agent learn policies that consider component interactions (e.g. if neighboring turbines' maintenance can be combined). In their framework, the graph-based state and GCN embedding encouraged the agent to group maintenance geographically, improving efficiency.

5.4.4 Hierarchical and interpretable models

To tackle the black-box nature of DRL, Abbas et al. (2024) introduced a hierarchical DRL framework for turbofan engine maintenance that combines an Input-Output Hidden Markov Model (IOHMM) with a DQN-like agent. The high-level IOHMM module interprets sensor data to detect the likely fault mode or degradation state (providing a human-understandable diagnostic), while the low-level DRL module learns the optimal replacement or repair policy given that inferred state. This two-level approach achieved performance on par with end-to-end DRL, but with the added benefit that decisions could be traced to identifiable health-state estimates.

In safety-critical domains like aerospace or offshore energy, such interpretability is valuable for gaining trust in the AI's recommendations. Moreover, by narrowing the policy search to focus on critical decisions (informed by the HMM), the agent avoids spurious actions in sparse failure domains. Although turbofan engines differ from wind turbines, the concept is relevant: they use a probabilistic model to identify health states and provide an interpretable layer, and the DRL agent makes maintenance decisions at a higher level. This yields a more transparent policy, which could be valuable in offshore wind where operators demand understanding of the AI's decisions (Adadi & Berrada, 2018). Such hierarchical or hybrid frameworks (combining physics-based or expert models with learning) are a way to inject domain knowledge directly into the DRL algorithm, improving learning speed and trustworthiness.

Figure 5.5 illustrates different frameworks commonly used to represent state information in Deep Reinforcement Learning (DRL) models for maintenance decision-making. Starting from an initial state S , the flowchart shows four distinct ways the state can be represented or processed before making a decision. The MDP formulation (blue) assumes full access to the true system state. The POMDP approach (yellow) highlights the partial observability of real-world systems requiring the agent to infer underlying states from limited observations $O(S)$ of the true state. The Graph-based method (green) structures state information using nodes and edges relationships to capture spatial or topological dependencies, allowing the agent to consider how interconnected assets affect one another's maintenance needs. Finally, the Hierarchical approach (red) incorporates domain-specific insights, either from expert knowledge or physics-based models. After selecting an action A based on the chosen representation, the system receives a reward R and transitions to a new state S' , looping back to begin a new decision-making cycle.

5.4.5 Decision frameworks and objectives

Across these formulations, the decision-making frameworks can vary in objective. Some agents aim to maximize availability or energy production (profit) (Pincirolì et al., 2021; Lee et al., 2025). Others aim to minimize total cost (including repair costs, downtime costs, and possibly penalty for using resources) (Kerckamp et al., 2022). These are effectively two sides of the same coin, since maximizing uptime or output will implicitly minimize downtime losses. The reward function should, in fact, include terms for lost revenue when a turbine is down, crew/vessel dispatch costs, spare part costs, and maybe even penalties for equipment degradation. For example, Pincirolì et al. (2021) designed a reward equal to the short-term profit (energy revenue minus O&M costs) at each decision step, so the PPO agent's cumulative reward corresponds to total profit. Cheng et al. (2023) focused on cost rates, giving negative rewards for inspection or repair costs and for failures, to encourage the agent to find the policy with minimum average cost. In multi-agent settings, a




Figure 5.5: Overview of problem formulations in Deep Reinforcement Learning models: Markov Decision Problem (blue), Partially Observable Markov Decision Problem (yellow), graph/network-based formulation (green), and hierarchical approaches (pink)

global reward is often used for full cooperation (as in QMIX, where agents maximize a shared cost-saving metric). Some works also impose constraint handling in the formulation. For instance, maintenance actions might be invalid under certain weather conditions; a realistic environment simulator will simply not allow those actions (or will assign a large negative reward if attempted). Thus, ensuring decisions are feasible (e.g. not sending a crew when waves are too high) can be done by action masking or via constraint-penalty in the reward. Recent research even explores incorporating optimization constraints into neural network design (e.g. using attention masks to enforce constraints) as in Kazemian et al. (2024), though this is still emerging in O&M applications.

In summary, the methodologies range from straightforward MDP models with fully observable states to sophisticated POMDP and graph-based models that embrace the complexities of offshore wind maintenance. The trend is toward more realistic problem formulations, acknowledging partial observability, incorporating spatial and logistical structure, and aligning the reward with business metrics (cost or profit). The following section provides a structured summary of these methodologies based on their application and the domain-specific knowledge considered, finally highlighting key aspects such as agent design, algorithm choices, and problem formulations in a comparative table. This summary aims to offer a clear and concise reference for understanding the variations in DRL-based maintenance planning approaches and their defining characteristics.

5.5 Applications of DRL approaches for offshore wind O&M

A fundamental aspect of training and deploying DRL agents for maintenance planning is the integration of domain-specific knowledge. The following subsections explore how different forms of domain knowledge can enhance decision-making: (i) *wind farm aerodynamic* interactions that affect energy yield; (ii) *weather and sea state* constraints that determine accessibility; (iii) *logistics* and crew-related considerations; (iv) *prognostic and health management data* for predictive maintenance, and finally (v) *economic factors* such as market prices and budget limits.

5.5.1 Wind farm aerodynamics

In a wind farm, turbines cast aerodynamic wakes that reduce the output of downwind turbines (Vermeer et al., 2003). In Lee et al. (2025), the authors considered this by making their DRL agent wake-aware. They fed a wake interaction model into the state, and used an ensemble of DQN models to capture different wake scenarios. By incorporating multiple wake models, their agent learned maintenance decisions that minimize farm power loss due to wakes, leading to higher overall energy production. This domain knowledge is particularly important for tightly spaced wind farms where wake losses are significant; a purely self-learned agent without wake inputs might need enormous training experience to “discover” such effects, whereas providing a wake model upfront accelerates learning. The result is a policy that schedules maintenance such that either the waking effect is minimal (e.g. performing maintenance when winds are low or from directions that don’t affect other turbines) or that multiple affected turbines are maintained together to collapse the wake loss into a single period.

5.5.2 Weather and sea state

Offshore maintenance is heavily weather-dependent, high waves, strong winds, or storms can prevent crew transfers and repairs (Jenkins et al., 2021). Some DRL studies explicitly include weather in the environment. Chatterjee & Dethlefs (2021) demonstrated a DRL approach for planning vessel transfers to turbines, integrating real SCADA data and weather conditions. Their agent could prioritize critical repairs and navigate the stochastic availability of weather windows, something traditional scheduling struggles with (Borsotti et al., 2024). By training on historical weather patterns, the DRL policy learns, for example, to take advantage of a calm sea state to perform a repair even if it's slightly early, because waiting might mean a long weather delay. Similarly, Pincirolì's state included predicted power (which indirectly reflects wind forecast) to help the agent plan around periods of low wind (often correlated with calmer weather) (Pincirolì et al., 2021; Ogunfowora & Najjaran, 2023).

In practice, one can input wave height forecasts or wind speed forecasts into the DRL state; the agent will then learn not to “choose” an action that requires travel during bad weather. Domain knowledge here ensures feasibility and robustness: the agent that knows about weather will inherently develop a maintenance schedule that aligns with seasonal weather patterns, reducing cancellations and idle times.

5.5.3 Logistics

Offshore O&M involves vessels, helicopters, crews, and spare parts, logistical aspects that greatly affect cost. DRL models have started to include these. In the multi-crew PPO model by Pincirolì et al. (2021), the state and action were designed to capture crew positions and availability. The environment simulation accounted for travel times to turbines and repair durations. This domain realism meant the learned policy actually coordinates crew movements: e.g. sending Crew A to a turbine that will finish repair soon, while Crew B waits at the depot until a large failure occurs. By encoding travel time and multiple crews, the DRL agent learned to avoid wasted trips and to keep crews utilized, emulating optimal routing and scheduling decisions. In Kerkkamp et al. (2022), logistics is addressed in a spatial sense by using a GCN to group geographically close maintenance. In an offshore wind context, that could translate to handling nearby turbines in one outing to minimize transit.

Even without an explicit graph, a DRL agent can learn from cost feedback that doing maintenance on neighboring turbines back-to-back saves the transit cost of multiple separate trips. Future research is incorporating more detailed logistics, such as vessel capacity, fuel cost, and inventory of spare parts, into the DRL state/reward. The benefit of doing so is that the learned policy becomes a holistic O&M schedule that respects not just failure risks but also supply chain and labour constraints.

5.5.4 Prognostics and health management (PHM) data

Almost all DRL approaches in this area use PHM outputs (like condition monitoring and RUL predictions), which is essential domain knowledge for predictive maintenance. The difference is in how they incorporate it. For instance, one could use a discretized health state (e.g. “good”,

“degraded”, “critical”) as part of the state space, which is easier for an agent to handle than raw sensor readings. Lee & Mitici (2023a)’s approach of using a CNN on sensor data before the DRL agent is another way, essentially letting a deep learning model extract relevant features (vibration patterns, etc.) which portray the health knowledge. For different applications, Abbas et al. (2024) and Abbas (2024) integrate an Input-Output Hidden Markov Model to classify health states of turbofan engines and feed this into the DRL agent.

By combining PHM and DRL, these frameworks close the loop from condition monitoring to decision-making. The advantage is clear: the better the agent’s awareness of actual component condition (even if inferred), the better it can time maintenance. PHM can also be embedded in the model by shaping the reward as well, e.g. a large penalty for a failure is effectively encoding the idea “a catastrophic gearbox failure is very bad” which the agent must learn. PHM-driven rewards can also be used, such as giving a small penalty for running a component in a highly degraded state (reflecting higher wear or risk).

5.5.5 Economic factors

Domain knowledge also includes economic factors like energy price and maintenance cost structure. Some works incorporate dynamic electricity prices or contractual penalties into the reward. For instance, if energy price forecasts are available, a DRL agent could decide to do maintenance when energy price (hence lost revenue) is low. While not explicitly seen in the reviewed models, Attention-based approaches are starting to include market price as part of the input to O&M scheduling models (Kazemian et al., 2024). In offshore wind, including such factors could align maintenance not just with technical needs but also with business cycles (e.g. perform maintenance during low demand seasons or when subsidy/price is low). Another economic constraint is the budget or resource limit, DRL can consider these by capping certain rewards or through state variables (like remaining budget).

Integrating domain-specific information can make DRL formulations more realistic and better aligned with the physical and operational characteristics of offshore wind systems. Figure 5.6 summarises the types of knowledge commonly incorporated in existing studies, including wake interactions, weather accessibility, logistics constraints, degradation physics, and PHM-derived indicators. The role of these elements is to provide structure that may help the agent focus on features that are relevant to decision-making, nevertheless, it is worth noting that increased model complexity, additional training effort, or conflicting inputs may offset potential gains, and over-specifying the environment can make optimization more difficult rather than easier. Thus, domain knowledge can support DRL when carefully selected and validated, but its contribution depends on the quality, relevance, and reliability of the information provided.

All the studies reviewed have, in one way or another, blended domain knowledge into their DRL approach: Lee et al. incorporated wake-effect modelling and aerodynamic interactions, Chatterjee added weather data, Kerkkamp incorporated graph-based asset relationships and spatial dependencies, Pinciroli added RUL and power forecasting, Abbas added HMM interpretable layers, etc. This synergy of domain expertise and reinforcement learning is key to developing trustworthy, high-performance maintenance policies that industry can adopt.

To compare the DRL-based maintenance planning approaches reviewed, Table 5.2 highlights their key features such as agents, algorithms and problem formulations. All the mentioned ap-




Figure 5.6: Domain-specific knowledge for offshore wind

proaches aim to optimize maintenance scheduling but differ in the algorithms and formulations used. The column *agent* refers to whether one policy controls the whole system or multiple coordinating policies exist, *algorithm* shows which training method was used, *problem formulation* indicates how the maintenance decision problem is modeled for the agent, and *domain-specific knowledge* indicates which categories of offshore wind-specific knowledge are explicitly integrated in each study (e.g., wake/aerodynamic effects, weather, logistics, and PHM information).

In Figure 5.7, we provide a comparative view of the distribution of key features across the reviewed literature. 5.7a shows the proportion of single- versus multi-agent frameworks. 5.7b outlines the distribution of algorithms such as DQN, PPO, SAC, and QMIX variants, showing that DQN remains the most commonly used DRL algorithm despite the explosion of the action space, which grows with the number of components or maintenance tasks, limiting its applicability in larger-scale problems. Finally, 5.7c highlights the prevalence of MDP and POMDP-based modeling.

With a clear understanding of the diverse methodologies used to model offshore wind maintenance planning, we now turn to the practical impact of these approaches. The following section summarizes the key contributions and performance improvements achieved through the application




Figure 5.7: Comparative view of single vs. multi-agent methods, main Deep Reinforcement Learning algorithms, and problem formulations in the reviewed literature.

Table 5.2: Comparison of Deep Reinforcement Learning Approaches for Offshore Wind Farm Maintenance.

Reference	Reported gain(s)	Agent	Algorithm	Problem Formulation	Domain-specific knowledge
Lee et al. (2025)	+11.1% power generation vs. baseline schedule	Multi	✓	MDP	Aeroacoustics ✓
		Single	✓	✓	✓
Abbas et al. (2024)	NA	✓	✓	✓	✓
		✓	✓	✓	✓
Do et al. (2024)	~20% more cost-effective vs. threshold-based maintenance (total maintenance cost, case study)	✓	✓	✓	✓
		✓	✓	✓	✓
Lee & Mitici (2023a)	95.6% reduction in unplanned downtime and maintenance cost vs. periodic maintenance	✓	✓	✓	✓
		✓	✓	✓	✓
Cheng et al. (2023)	23% life-cycle cost reduction: 2.32×10^4 EUR vs. 3.01×10^4 EUR (DIAR vs. UIFR)	✓	✓	✓	✓
		✓	✓	✓	✓
Zhao & Zhou (2022)	NA	✓	✓	✓	✓
		✓	✓	✓	✓
Kerkkamp et al. (2022)	NA	✓	✓	✓	✓
		✓	✓	✓	✓
Nguyen et al. (2022)	~20% more cost-effective vs. threshold-based maintenance (total maintenance cost, case study)	✓	✓	✓	✓
		✓	✓	✓	✓
Pinciroli et al. (2021)	NA	✓	✓	✓	✓
		✓	✓	✓	✓
Chatterjee & Dethlefs (2021)	NA	✓	✓	✓	✓
		✓	✓	✓	✓
Andriotis & Papakonstantinou (2021)	NA	✓	✓	✓	✓
		✓	✓	✓	✓

of DRL, illustrating how these methodological choices translate into tangible benefits such as cost reduction, improved reliability, and enhanced operational efficiency.

5.6 Discussion

The application of DRL to offshore wind O&M has yielded notable improvements over traditional maintenance strategies, as documented in the reviewed studies. The application of DRL to offshore wind O&M has shown clear advantages over conventional maintenance strategies, achieving lower costs, higher availability, and more adaptive planning. The reviewed studies highlight five recurring benefits: (i) *cost and downtime reduction*, as agents learn to time interventions “just in time” before failure; (ii) *enhanced predictive maintenance*, through integration of RUL data and operational context; (iii) *opportunistic maintenance*, by grouping actions and exploiting favorable conditions; (iv) *improved reliability and safety*, via proactive scheduling and embedded risk constraints; and (v) *computational scalability*, enabling optimization of large, stochastic systems.

Finally, we focus on the remaining gaps regarding the use of real-case scenarios for testing the models and the inclusion of multi-level repair that should reflect real operational practices.

The following subsections discuss each benefit in detail, illustrating how DRL contributes to more cost-efficient and resilient offshore wind maintenance planning.

5.6.1 Cost and downtime reduction

A primary goal is lowering maintenance costs and turbine downtime compared to baseline strategies (reactive or scheduled). DRL agents have demonstrated the ability to significantly reduce unplanned failures and associated costs.

For example, (Cheng et al., 2023) report that their PPO-based adaptive inspection and repair policy (DIAR) achieves an expected life-cycle cost of 2.32×10^4 EUR, compared to 3.01×10^4 EUR for the best uniform-interval, fixed-threshold strategy (UIFR), i.e. a 23% reduction in their case study.

In Lee & Mitici (2023a), the DRL-based policy reduced unplanned downtime and maintenance cost in a multi-component system by 95.6% versus conventional periodic maintenance.

These improvements come from the agent’s ability to optimize the timing of maintenance: servicing components “just in time” before failure, but not too early to waste life. DRL policies effectively find this sweet spot by continuous learning and adjustment.

5.6.2 Predictive (PHM-based) strategies

Many wind operators use condition-based triggers (like RUL thresholds from prognostics) to plan maintenance. DRL can enhance these predictive strategies by adding dynamic decision-making. As noted in Pincirolì et al. (2021), the DRL policy outperformed a pure RUL-threshold policy by considering not only the component health but also external factors such as power demand and crew availability. In other words, whereas a standard predictive maintenance strategy says “replace component X when its $RUL < Y$ days,” a DRL agent might learn “replace component X

when $RUL < Y$ days and a maintenance team is idle and a low wind period is coming up,” thereby minimizing impact. This kind of contextual decision-making led to higher reward (profit) in their experiments. We see a similar theme in Lee et al. (2025): their DQN agent, augmented with wake effect knowledge, boosted energy production beyond what a wake-unaware strategy achieved. By learning the true optimal policy through trial and error, DRL approaches can exceed the performance of both corrective maintenance (which incurs high downtime) and simple predictive rules (which might be myopic or inflexible).

5.6.3 Opportunistic maintenance

DRL has shown strength in exploiting opportunistic maintenance opportunities that humans or simple policies might miss. For instance, in multi-component scenarios, an agent can coordinate maintenance on multiple turbines in one go to avoid repeated downtime. The PPO agent in Pincirolì’s work learned to wait for low wind output days to schedule maintenance, which is an opportunistic behavior yielding higher rewards. Huang et al. (2020) observed that, even when not explicitly programmed, DRL agents naturally learn to group multiple repairs together, thereby reducing repetitive downtime and sharing high logistics costs over several interventions, similar findings can be found in Dong et al. (2021) and Valet et al. (2022). Do et al. (2024) demonstrated this with their multi-agent approach where the learned policy effectively grouped maintenance tasks to save on shared downtime costs, beating a policy that treats components independently. Even in single-agent setups, agents learn to use times of low production or existing outages to perform additional repairs. In the context of offshore wind, for example, an agent might schedule a minor repair on one turbine when a vessel is already en route for a major repair on a neighboring turbine, effectively reducing additional transit costs.

5.6.4 Reliability and safety

An important point is that DRL policies can improve reliability metrics (e.g. mean time between failures, availability) by preventing failures proactively. Lee & Mitici (2023a) noted fewer failures in their DRL-maintained system than a conventional approach. Additionally, DRL can incorporate safety constraints (like not allowing maintenance deferral beyond a limit) via rewards or state features. The hierarchical approach by Abbas (2024) is aimed at safety-critical maintenance, by integrating an interpretable model, they ensure the DRL decisions for turbofan engines remain within safe bounds and can be understood by engineers. In offshore wind, ensuring that a DRL policy does not inadvertently run turbines to catastrophic failures is very important. Studies so far show that with proper reward design (heavily penalizing failures), the agents naturally learn to avoid risky deferrals.

5.6.5 Computational feasibility

Another finding across the literature is that DRL can handle high-dimensional problems that were previously intractable by brute force optimization. Maintenance scheduling for a wind farm with many turbines, each with several components, is a huge combinatorial problem over a long horizon. Some studies also accelerate learning by transfer or imitation. For example, Pincirolì et al. (2021)

initialized their PPO agent via imitation learning from a heuristic policy to shorten training time. This hybrid approach marries human insight with AI optimization. The result is a practical decision-support tool that can quickly recommend which turbine to maintain and when, given the current observations.

DRL, with its experience-driven learning, provides feasible solutions. Nevertheless, training can be computationally intensive (e.g. WQMIX took 12 hours in the 13-component case (Nguyen et al., 2022)), although once trained, the policy can execute decisions in real-time.

Model-based or brute-force methods struggle to consider all contingencies and long-term effects (Chen et al., 2025). The consensus of recent work is that DRL-based maintenance planning consistently outperforms static strategies in simulation, often by a wide margin in cost savings or uptime. These improvements stem from DRL's ability to learn optimal scheduling under uncertainty, adapt to varying conditions (weather, load demands), and coordinate multiple decisions in a way that human-designed rules cannot easily mimic.

While the performance gains of DRL-based maintenance strategies are evident, their success is further amplified by the integration of domain-specific knowledge. The next section focuses on how incorporating elements such as wind farm aerodynamics, weather constraints, logistical considerations, and PHM data not only enriches the state representations but also guides the learning process toward more realistic and reliable maintenance policies.

5.6.6 Simulation-based studies vs. real-world applications

Simulation-based research has been the predominant way to develop and evaluate DRL for wind farm maintenance. All the studies reviewed above rely on simulated environments, typically a stochastic model of turbine degradation and failure, combined with models of maintenance actions (costs, durations, effects) and often weather or logistics simulators.

For instance, Lee et al. (2025) used a simulation of wind dynamics and turbine wakes to train their DQN ensemble; Pincirolì et al. (2021) built a custom simulator for turbine failures and crew assignments to test PPO. Simulation is essential because it provides a safe and flexible sandbox to train DRL agents (which often require millions of decision steps for convergence). Researchers can speed up or repeat scenarios (like many years of operation) to expose the agent to rare events, something impossible to do quickly in the real world.

Key insights from simulation studies include the significant performance gains of DRL policies over traditional maintenance strategies. Many papers report that their learned policies yield lower cost or higher availability than periodic (time-based) or reactive (run-to-failure) maintenance. For example, Pincirolì et al.'s agent outperformed corrective and age-based schedules with fewer failures and lower cost, mirroring Andriotis & Papakonstantinou (2021)'s results against periodic policies. Such results, consistently observed in simulation, build a compelling case for applying DRL in practice.

That said, real-world applications of DRL for offshore wind O&M remain at an early stage, and published demonstrations typically rely on simplified physical and operational assumptions. A recent example is the work by Lee et al. (2025), developed in collaboration with an industry research institute (KEPCO). Their DQN-based scheduling framework illustrates how DRL can, in principle, coordinate maintenance actions and exploit wake interactions to improve farm-level

energy production. However, their study relies on the Jensen kinematic wake model, which, while computationally efficient, provides only coarse accuracy in the near-wake region and may misrepresent wake recovery dynamics. Consequently, the reported power-gain improvements (i.e. an 11.1% increase over their baseline scheduling strategy) should be interpreted as gains within the constraints of a simplified simulation environment (Lee et al., 2025). Additional assumptions, such as fixed maintenance durations, deterministic task lists, and the absence of vessel or access constraints, further limit the generalisability of the results.

These studies therefore highlight both the potential and the current limitations of DRL for maintenance planning: DRL can identify useful scheduling patterns in controlled testbeds, but its effectiveness in operational settings remains dependent on the fidelity of the underlying simulator. More comprehensive validation using higher-fidelity wake models, realistic metocean variability, and operational logistics will be essential to assess the practical applicability of DRL-based maintenance planning.

While the deployment of a trained agent in a live wind farm has not been reported yet, the involvement of an industrial stakeholder and the real-world fidelity of the simulation (including measured wake effects) indicate a step toward actual adoption.

In fields like aviation and manufacturing, some DRL-based maintenance planners have been tested on real data if not deployed directly. For instance, the turbofan case by Abbas et al. (2024) used NASA engine datasets to train and validate their approach. This kind of validation on real-world data builds confidence that the policies will translate from simulation to reality. Moreover, ongoing advances in digital twins for wind farms (Zhang et al., 2024) (high-fidelity replicas of turbines and operations) may enable DRL agents to be trained or at least fine-tuned in an environment that closely mimics reality.

5.6.7 Multi-level repairs

Traditional maintenance models in offshore wind operations typically reduce decisions to a simple binary choice, either repair or not. However, real-world observations show that turbine downtime arises from a spectrum of failures. For instance, industry data reveals that roughly 70% of downtime is caused by major repairs, about 17% by minor repairs, with the remainder due to simple resets (Carroll et al., 2017).

This suggests that maintenance actions in the field are inherently multi-level, ranging from minor fixes that temporarily restore performance to major overhauls or full component replacements.

To make the notion of *multi-level* maintenance actions more concrete in a wind context, Aafif et al. (2022) study preventive maintenance for a wind-turbine gearbox under temperature-based condition monitoring. They compare (i) a commonly adopted industrial strategy in which, whenever a threshold temperature is exceeded, the turbine is temporarily derated while the gearbox is cooled, and, after such events become frequent enough, the gearbox is ultimately renewed (replacement or overhaul), against (ii) a multi-level strategy in which each threshold exceedance triggers an imperfect preventive maintenance action that partially restores the gearbox condition by reducing its failure rate to a value between the current rate and that of a new gearbox, with renewal enforced only after N imperfect PM actions. Their numerical comparison shows that neither “renewal-only” nor “imperfect-PM-then-renewal” dominates universally: the more economical strategy depends on gearbox reliability and on the relative magnitude of production-loss, cooling, preventive-maintenance,

and renewal logistics costs. This case illustrates why binary “repair/replace” abstractions can be limiting for offshore wind O&M: introducing intermediate action levels (e.g., partial restoration actions before renewal) enables the policy to express realistic trade-offs between short-term production impacts and long-term degradation management.

While incorporating these multi-level actions enriches the state-action representation, it also exposes a significant gap in current DRL models for offshore wind O&M. Most existing studies focus on binary decision frameworks.

The lack of integration of multi-level maintenance actions is discussed in greater detail in the next Section by examining how DRL-based planning models can incorporate multi-level maintenance strategies, allowing agents to choose among different maintenance tasks. Furthermore, we will explore the implications of this refined modeling on cost, reliability, and overall operational efficiency.

5.7 Research gap

Traditional maintenance models in offshore wind often reduce decisions to a binary choice, such as whether to maintain a component or defer the intervention. However, real offshore wind O&M practice involves a wider range of intervention types, including no action, minor repair, major repair, and replacement. These actions differ not only in cost, but also in duration, logistical requirements, downtime consequences, and restorative effect on the component state.

In a DRL formulation, this limitation can be addressed by expanding the action representation beyond a simple maintain-or-wait decision. The maintenance planning problem can be formulated as a sequential decision problem in which the state includes asset and operational features, such as component condition, age, prognostic information, accessibility, and weather windows, while the action space enables the agent to choose among multiple repair options. Such an approach allows the agent to balance short-term cost and long-term reliability by, for example, selecting a less expensive interim repair when degradation is moderate, or committing to a full replacement when the risk of failure justifies the higher cost.

Incorporating minor and major repairs as viable actions can prevent small degradations from escalating into failures while avoiding unnecessary replacements. DRL agents trained on multi-level maintenance tasks can learn when low-level interventions are sufficient and when more extensive actions are required. For example, Wei et al. (2019) demonstrated that a DRL-based policy for structural maintenance maintained high reliability by optimally balancing minor and major interventions across numerous components.

Two modeling avenues are particularly relevant for enabling DRL agents to consider multi-action maintenance planning:

- **Expanded action spaces:** Researchers have begun explicitly modeling a range of actions, such as “do nothing”, “perform minor repair”, “conduct major repair”, or “replace component”. For example, Zhang et al. (2023) formulate an infinite-horizon DRL maintenance problem where a component’s health can be partially recovered through an imperfect repair or fully restored via corrective maintenance. This expanded action space helps the agent learn which level of intervention yields the optimal long-term cost and reliability trade-off.

- Parameterized and hybrid actions: To manage the complexity arising from multiple discrete repair choices combined with continuous or scheduling-related variables, advanced approaches employ parameterized action spaces. In one instance, a parameterized PPO algorithm was developed to handle mixed discrete-continuous decisions, effectively allowing the agent to adjust both the type of repair and the timing of intervention simultaneously (Zhang et al., 2023). This structured action representation helps preserve convergence stability while exploring more complex maintenance policies.

For the purpose of this thesis, this research gap is treated as a modeling requirement addressed in Chapter 6. The learning-based framework developed there adapts DRL to the offshore wind maintenance scheduling problem by combining uncertain health observations, offshore-specific costs and downtime signals, and a structured action representation that includes multiple intervention types. In this way, Chapter 6 takes a first step beyond binary DRL maintenance policies and evaluates whether a richer action space can support more realistic O&M decision-making.

Having identified the need for structured action spaces in DRL-based offshore wind maintenance scheduling, the chapter now consolidates the main modeling requirements derived from the review. The following section summarises these requirements and explains how they motivate the learning-based framework developed in Chapter 6.

5.8 Conclusions

This chapter addressed Research Question Q4: *Which modeling requirements should a learning-based offshore wind O&M framework satisfy to address the limitations identified in the optimization-based framework?* Building on the limitations diagnosed in Chapter 4, the chapter translated the observed constraints of the rolling-horizon MILP paradigm into a set of requirements for learning-based O&M planning.

The first requirement is long-term sequential decision-making. Chapter 4 showed that optimization-based planning can produce high-quality and interpretable schedules under explicit constraints, but also that a finite planning horizon may limit the model's ability to internalise actions whose consequences extend beyond the active optimization window. A learning-based framework should therefore evaluate maintenance actions through their long-term consequences, including future degradation, downtime, accessibility, and grouping opportunities.

The second requirement is fast online adaptivity. In the optimization-based framework, new RUL predictions, updated accessibility information, or changes in component condition require the planning problem to be solved again. This closed-loop re-optimization is valuable because it preserves explicit feasibility and interpretability, but it may become computationally demanding as the number of turbines, components, vessels, scenarios, and action choices increases. A learning-based framework should therefore shift part of this computational burden to offline training, so that online decisions can be obtained rapidly by evaluating a trained policy.

The third requirement is scalability. Offshore wind O&M planning becomes increasingly high-dimensional when larger wind farms, multiple components per turbine, heterogeneous vessels, weather-dependent access, and multi-level maintenance actions are represented simultaneously. The reviewed DRL literature shows that function approximation, multi-agent decomposition, graph-based representations, and hierarchical formulations can provide mechanisms to structure this

complexity. These methods are especially relevant when the decision space becomes too large for direct enumeration or repeated exact optimization.

The fourth requirement is partial observability and uncertainty-aware state representation. Offshore wind maintenance decisions are not based on direct knowledge of the true degradation state, but on imperfect observations such as condition-monitoring signals, inspection outcomes, RUL predictions, and weather forecasts. The reviewed literature therefore supports the use of POMDP formulations, belief-state features, recurrent memory, or uncertainty-aware observations to represent the fact that health and accessibility information are only partially known at the time decisions are made.

The fifth requirement is domain-informed operational feasibility. A learning-based method cannot be useful for offshore wind O&M if it learns policies that ignore access limits, resource constraints, vessel requirements, or the cost structure of offshore interventions. The review shows that DRL models become more relevant when they incorporate domain-specific knowledge such as metocean accessibility, logistics, PHM information, wake effects, and economic cost drivers. In this thesis, this requirement is particularly important because the learning-based model must remain consistent with the modeling environment defined in Chapter 3.

The sixth requirement is realistic multi-level maintenance actions. A central limitation of many existing DRL formulations is their reliance on binary repair decisions, such as maintain or wait. This abstraction is insufficient for offshore wind O&M, where interventions may involve minor repair, major repair, or full replacement, each with different costs, durations, resource requirements, and restorative effects. A learning-based O&M framework should therefore preserve the multi-level action structure already used in the optimization-based framework, rather than simplifying maintenance into a single repair action.

Taken together, these requirements clarify the role of DRL in this thesis. DRL is not introduced as a universal replacement for mathematical optimization. Rather, it is investigated as a complementary decision-making paradigm that may address specific limitations of finite-horizon optimization: sensitivity to horizon boundaries, repeated re-optimization, high-dimensional state-action spaces, and uncertainty in system observations. At the same time, DRL introduces its own challenges, including dependence on simulator fidelity, reward-design sensitivity, limited interpretability, training cost, and the need for extensive validation before practical deployment.

The literature demonstrates that DRL is a promising approach for offshore wind farm maintenance planning. By learning from interactions in a simulated environment, DRL agents can devise maintenance policies that outperform traditional corrective, time-based, and other predictive strategies on key metrics like cost, downtime, and energy production. Both single-agent and multi-agent frameworks have been explored: single-agent DRL has succeeded in optimizing complex maintenance schedules by considering long-term consequences, while multi-agent DRL offers a path to scaling these solutions to larger systems by decentralizing decisions. Moreover, the most successful studies embed domain-specific knowledge, from wake physics to weather patterns to prognostic models, into the DRL process, creating hybrid solutions that learn efficiently and behave realistically.

The reviewed works highlight several advantages of DRL in offshore wind O&M: adaptive scheduling that responds to the actual condition of turbines, opportunistic maintenance that smartly times actions to minimize impact on operations, and the ability to handle the high dimensionality of scheduling problems that defy mathematical or brute-force optimization. For example, DRL

agents have learned to schedule maintenance during low wind periods (Ogunfowora & Najjaran, 2023), to cluster repairs and save vessel trips (Kerckamp et al., 2022), and to prevent failures by reacting to prognostic alarms better than fixed rules. These capabilities translate into quantifiable gains: double-digit percentage improvements in cost savings or energy output in case studies are common (Lee et al., 2025; Nguyen et al., 2022).

However, challenges remain before DRL is commonplace in live offshore wind operations. Safety and trust are critical aspects; operators need assurance that an AI agent won't recommend catastrophic decisions. This is why interpretability (as in the hierarchical HMM approach) and extensive testing are crucial. Computational efficiency is also a concern: multi-agent or long-horizon DRL can be computationally intensive in the training phase, though improvements in algorithms and hardware mitigate this. Additionally, integration with existing maintenance management systems requires user-friendly interfaces and perhaps human-in-the-loop designs (where human planners can review or override AI suggestions).

Moreover, a critical limitation in current DRL applications is the overly simplistic treatment of repair actions. Most methods consider only one kind of repair, typically reducing the decision to whether to replace a component. In practice, maintenance is a multifaceted process that often involves choosing among various repair strategies, each with distinct implications for system performance and cost. For instance, an optimal policy should not only determine the optimal timing of an intervention but also decide which specific repair tasks to undertake based on the current condition of components. Addressing this gap requires developing agents capable of discerning a richer set of actions that reflect the complexities of real-world maintenance tasks.

For practitioners and researchers, these findings support moving toward DRL-based decision support tools for wind farm O&M. As offshore wind farms continue to grow and data from operations accumulate, DRL approaches aim to become integral in optimizing maintenance planning, ultimately lowering the cost of renewable energy and improving the reliability of wind power generation.

Despite efforts to follow a transparent and reproducible literature selection protocol as defined in Section 5.1, several threats to validity remain. First, *publication bias* may affect the evidence base: studies reporting positive performance gains of DRL approaches are more likely to be published than negative or inconclusive results, and industrial deployments are often underreported due to confidentiality. As a consequence, the reviewed corpus may over-represent successful proof-of-concept demonstrations and under-represent failure cases or practical limitations. Second, *reproducibility* is constrained by limited access to high-fidelity O&M simulators, proprietary data (e.g., SCADA/CM and maintenance logs), and incomplete reporting of experimental details (e.g., hyperparameters, reward shaping, baseline tuning, random seeds). These issues complicate rigorous replication and can make cross-paper comparisons sensitive to implementation choices rather than underlying algorithmic differences. Third, *generalisability* is limited because most reviewed studies evaluate DRL in stylised simulation environments with simplified wake, metocean, and logistics assumptions; consequently, reported gains may not transfer to real offshore operations or to farms with different layouts, failure modes, contractual constraints, and accessibility regimes. Finally, while many insights on DRL modeling choices (e.g., partial observability remedies, multi-agent coordination, multi-level actions) are transferable beyond offshore wind, their effectiveness may vary substantially across domains and should not be assumed without domain-specific validation.

The findings of this chapter directly motivate Chapter 6. Building on the requirements derived here, the next chapter formulates offshore wind O&M planning as a POMDP-based DRL problem.

The proposed framework uses probabilistic RUL information to represent uncertain component health, defines multi-level maintenance actions to reflect realistic intervention choices, and evaluates the resulting policy within the modeling environment established in Chapter 3. In this way, the thesis moves from the optimization-based planner of Chapter 4 to a learning-based policy framework designed around the modeling requirements identified in the present chapter.

Chapter 6

Learning-Based O&M Planning

This chapter addresses Research Question Q5: *How can offshore wind O&M decision-making be formulated and solved as a learning problem to derive adaptive maintenance policies under offshore constraints?*

The preceding chapters developed the methodological progression of the thesis. Chapter 3 established an offshore wind O&M modeling environment, including the wind farm configuration, critical components, maintenance action set, cost structure, accessibility assumptions, and synthetic prognostic information. Chapter 4 then formulated tactical O&M planning as a mixed-integer linear programming problem, showing that probabilistic RUL information can be translated into explicit maintenance decisions under offshore-specific constraints. However, the rolling-horizon optimization analysis also revealed structural limitations, including horizon dependence, repeated re-optimization burden, limited scalability, and sensitivity to future events outside the planning window. Chapter 5 therefore examined Deep Reinforcement Learning (DRL) as a candidate paradigm for learning adaptive maintenance policies under uncertainty, and identified three key modeling requirements for a learning-based offshore wind O&M framework: partial observability of component health, operationally meaningful maintenance actions, and reward structures that capture the trade-off between preventive interventions, corrective failures, and lost production.

The requirements identified in Chapter 5 should be interpreted as characteristics of a more comprehensive learning-based offshore wind O&M framework, rather than as features that are all implemented simultaneously in the model developed in this chapter. The subsequent model focuses specifically on uncertain component-health information, sequential decision-making, and multi-level maintenance actions, while retaining the controlled abstractions of the common modelling environment introduced in Chapter 3. Thus, the elements addressed in Chapter 5 serve both as design principles for the proposed model and as a roadmap for progressively extending learning-based O&M models toward greater operational fidelity.

Building on these requirements, this chapter formulates offshore wind O&M planning as a Partially Observable Markov Decision Process (POMDP). In this formulation, the true degradation state of each component is not directly observable. Instead, the decision-maker receives uncertain health observations derived from probabilistic RUL predictions. This reflects the practical setting in which condition-monitoring and prognostic systems provide estimates of future failure behaviour, but not perfect knowledge of the underlying component state. The maintenance planner must there-

fore learn to act under imperfect information, balancing the cost of early preventive interventions against the risk of unexpected failures and downtime.

The chapter proposes a DRL framework based on Proximal Policy Optimization (PPO). The PPO agent observes discretized health states derived from probabilistic RUL forecasts and selects maintenance actions for each component. Unlike many existing DRL formulations for maintenance planning, which reduce the decision to a binary choice between maintenance and no maintenance, the proposed framework includes a multi-level action space consisting of no action, minor repair, major repair, and replacement. This action structure is aligned with the maintenance taxonomy defined in Chapter 3 and responds directly to one of the main gaps identified in Chapter 5: learning-based O&M models should not only decide whether to maintain, but also which type of intervention is appropriate for the observed condition of the component.

This chapter demonstrates how the offshore wind O&M problem can be represented as a learning environment in which probabilistic prognostic information, component degradation, maintenance effects, and cost consequences interact. Secondly, it evaluates whether a DRL agent with access to multi-level maintenance actions can learn an adaptive policy that improves long-term O&M performance relative to conventional benchmark strategies. The proposed PPO policy is compared against corrective maintenance, an age-based policy, and a binary-action PPO variant. This comparison isolates both the value of adaptive state-aware decision-making and the additional value of allowing multiple maintenance levels.

The chapter is organized as follows. Section 6.1 provides a short overview of the literature most directly relevant to the proposed implementation, focusing on predictive-maintenance, POMDP formulations, DRL algorithms, and multi-level repair actions. Section 6.2 presents the POMDP formulation, including the true state, observation model, action space, maintenance-state transition model, and reward definition. Section 6.3 describes the experimental setup, including the offshore wind farm demonstrator, component lifetime parameters, maintenance costs, and PPO training configuration. Section 6.4 compares the learned PPO policy against the benchmark strategies in terms of total cost, maintenance cost, downtime cost, failures, and maintenance action mix. Finally, Section 6.5 concludes the chapter by discussing how the results answer Q5 and how they contribute to the thesis-level comparison between optimization- and learning-based O&M planning.

This chapter is based on M. Borsotti, R. Patrao, X. Jiang, A. Napoleone, and R. R. Negenborn: “A Deep Reinforcement Learning Framework for O&M Planning in Offshore Wind”, accepted for publication in the *Proceedings of the 16th International Conference on Computational Logistics (ICCL 2025)* and *1st EURO Mini Conference on Maritime Optimization and Logistics (EUROMar 2025)*, Lecture Notes in Computer Science, Springer, forthcoming.

6.1 Literature review

Chapter 5 provided a broad review of DRL methods for offshore wind O&M and identified the modeling requirements that motivate the formulation used in this chapter. This section summarizes the literature most directly connected to the implementation developed here: predictive-maintenance decision support, POMDP-based formulations, DRL algorithms for high-dimensional maintenance planning, and the need for multi-level repair actions.

Decision-making approaches are generally categorized into (i) model-based approaches, which rely on physical models to describe the system, the environment and the underlying processes, (ii) knowledge-based approaches, that use expert and in-house knowledge to derive optimal policies, and (iii) data-driven models that rely on data analytics (Bousdekis et al., 2018). DRL falls into the third category, and has emerged as a promising approach to optimize decision-making across diverse industries for problems such as maintenance planning, layout optimization and more. The main advantage of DRL is its ability to handle large and complex environments while adapting to uncertainties through trial and error learning (Mnih et al., 2015). On the other hand, its main limitation lies in its black-box results: the agent's decisions are hardly interpretable by humans and it is hard to understand the causality of the suggested actions (Khan & Khan, 2012).

Most studies on optimizing maintenance schedules aim to enhance cost efficiency, operational reliability, and resource allocation. As discussed in Chapter 5, the reviewed papers demonstrate the potential of DRL in O&M planning, but they often rely on binary maintenance decisions and only partially consider the logistical challenges typical of offshore wind farms. Many works frame the O&M problem as a Markov Decision Process (MDP), employing DRL techniques such as Deep Q-Networks (DQN) and PPO to navigate high-dimensional decision spaces. For instance, Lee et al. demonstrate significant improvements in offshore wind farm productivity through a domain-informed DQN ensemble, incorporating wake interactions in tightly spaced turbines to improve real-world applicability (Lee et al., 2025). Similarly, Cheng et al. (2023) use adaptive repair thresholds and dynamic inspection intervals tailored to offshore conditions to reduce lifecycle costs.

Some studies incorporate multi-agent DRL frameworks to address economic, structural, and stochastic dependencies. For example, Do et al. utilize a Weighted QMIX algorithm to optimize maintenance decisions in multi-component systems (Do et al., 2024). Additionally, Abbas et al. integrate probabilistic models with DRL to enhance interpretability in predictive maintenance for turbofan engines, providing interesting insights for safety-critical applications (Abbas et al., 2024). Moreover, Kerkkamp et al. (2022) combine DRL with Graph Neural Networks. Their methodology formulates the maintenance problem as a POMDP on a graph structure and uses DRL to learn policies that account for component interactions.

POMDP formulations are particularly relevant for predictive maintenance because the true degradation state of a component is usually not directly observable. In Lee & Mitici (2023b), the authors propose a hybrid approach that integrates Convolutional Neural Networks (CNNs) with DRL for predictive maintenance in a multi-component aircraft system. Their methodology employs CNNs to extract degradation features from sensor data, which then inform the DRL agent's decision-making within a POMDP framework. Quantitatively, their experiments demonstrate a marked reduction in unplanned downtime and maintenance costs compared to conventional strategies. This type of formulation is aligned with the offshore wind setting considered in this chapter, where probabilistic RUL predictions provide partial and uncertain observations of component health.

Regarding the learning algorithms, value-based methods like DQN have been widely used for discrete action problems, such as in Lee et al. (2025), where an ensemble DQN policy coordinates maintenance for a 105-turbine offshore wind farm. While effective in some large-scale cases, DQN architectures are typically limited to single discrete action outputs and struggle with the high-dimensional, multi-discrete action spaces characteristic of complex O&M environments. In contrast, Proximal Policy Optimization (PPO) offers greater flexibility and stability for such problems. Its clipped objective function prevents overly aggressive policy updates, supporting stable learning

across large state-action spaces (Schulman et al., 2017). PPO has been shown to scale effectively in real-world-inspired settings; for example, Pinciroli et al. (2021) applied it to a 50-turbine wind farm with multiple maintenance crews, achieving better performance than optimized heuristics under exponential combinatorial complexity. These examples demonstrate how DRL, and PPO in particular, enables tractable and effective learning in high-dimensional O&M scenarios where traditional methods become difficult to apply.

The reviewed studies demonstrate the potential of DRL to address O&M planning by providing scalable, adaptive, and efficient solutions to complex scheduling problems. However, despite these advancements, a key limitation identified in Chapter 5 is the prevalent reliance on simplistic binary maintenance decisions (maintain/not maintain) rather than capturing the practical complexity of multi-level or imperfect repairs, which are essential for realistic operational scenarios (Su et al., 2022; Zhang et al., 2023). Addressing the lack of predictive maintenance models tailored to offshore wind applications and the gap in integrating multi-level repair strategies within the DRL literature, this chapter formulates offshore wind O&M planning as a POMDP and utilizes probabilistic RUL predictions to enable dynamic maintenance policies that account for multiple levels of repair.

6.2 Methodology

We formulate the maintenance problem as a POMDP because the physical condition governing component failure is not directly observable by the decision-maker. Instead, maintenance decisions are based on imperfect prognostic observations derived from the latent degradation state. Figure 6.1 illustrates this distinction between the true state, the information available to the agent, the selected maintenance actions, and the resulting rewards. The formulation adopted here constitutes a reduced POMDP representation of the benchmark environment: it captures partial observability of component health and stochastic state evolution, while retaining the simplifying assumptions on logistics, accessibility, and component dependencies introduced in Chapter 3.

As discussed in Chapter 5, the distinguishing characteristic of a POMDP is that the physical state governing system evolution is not directly available to the decision-maker. This characteristic is central to the predictive-maintenance problem considered in this chapter. Component failure




Figure 6.1: Schematic Representation of the POMDP. True states $S_c(k)$ evolve stochastically, while observations $O_c(k)$ are partial or noisy. Actions $A_c(k)$ are selected based on $O_c(k)$, yielding rewards $R_c(k)$ that guide policy updates.

is governed by the latent degradation state, represented here through the true RUL, whereas the maintenance planner only receives an uncertain prognostic estimate of that quantity. The true RUL therefore determines how the system evolves, but cannot itself be used by the agent when selecting maintenance actions. The same discretized health observation can correspond to different values of latent true RUL because of prognostic uncertainty, and these underlying states can imply different probabilities of future failure. The POMDP formulation preserves this distinction between the physical degradation process and the information available for decision-making and is therefore considered more appropriate than treating the uncertain prognostic estimate as the true system state.

For each component c at time step k we define:

True State $S_c(k)$: The true state represents the component's actual RUL:

$$S_c(k) = \text{RUL}_{\text{true},c}(k) \quad (6.1)$$

Each component c begins with an initial RUL drawn from a Weibull distribution with shape α_c and scale β_c , the probability density function is defined as:

$$f(t; \beta_c, \alpha_c) = \frac{\alpha_c}{\beta_c} \left(\frac{t}{\beta_c} \right)^{\alpha_c - 1} \exp \left[- \left(\frac{t}{\beta_c} \right)^{\alpha_c} \right] \quad (6.2)$$

where t represents the time since the start of the observation period. The choice of a Weibull distribution is motivated by its widespread use in modeling component lifetime distributions in offshore wind turbines, due to its flexibility in capturing various failure behaviors (Li et al., 2022a; Jiang et al., 2025).

Thus, we can generate a sample for the failure time of each component using inverse transform sampling:

$$\text{RUL}_{\text{true},c}(0) = \beta_c (-\ln(U))^{\frac{1}{\alpha_c}}, \quad U \sim \text{Unif}(0, 1).$$

Observation $O_c(k)$: In practice, the true RUL of a component cannot be measured directly. Instead, we derive an observation from predictive models. Let

$$\phi_c^k(t) \sim \mathcal{N}(\text{RUL}_{\text{pred},c}(t), \sigma_c(k)) \quad (6.3)$$

represent the probability distribution of the predicted RUL for component c at time t , where $\text{RUL}_{\text{pred},c}(k)$ is the point prediction made on day k and $\sigma_c(k)$ is the standard deviation.

To quantify the risk of failure over a selected prediction horizon H , we compute the failure probability as:

$$P_{\text{fail},c}(k; k+H) = \int_k^{k+H} \phi_c^k(t) dt, \quad (6.4)$$

such probability measures the likelihood that the component's RUL will fall before H days.

Based on $P_{\text{fail},c}(k; k+H)$, the observation space for component c at time step k is discretized into one of four health states:

$$O_c(k) = \begin{cases} \text{Good} & \text{if } P_{\text{fail},c}(k; k+H) < 50\%, \\ \text{Degraded} & \text{if } 50\% \leq P_{\text{fail},c}(k; k+H) < 70\%, \\ \text{Critical} & \text{if } 70\% \leq P_{\text{fail},c}(k; k+H) < 100\%, \\ \text{Failed} & \text{if } \text{RUL}_{\text{pred},c}(k) \leq 0 \end{cases} \quad (6.5)$$

In the present formulation, an updated prognostic observation is assumed to be available to the agent at every decision epoch. Information acquisition is therefore exogenous to the maintenance decision: obtaining a new health observation does not require a separate inspection action and does not incur an additional inspection cost. This assumption corresponds to a setting in which condition-monitoring and prognostic systems continuously provide updated component-health information to the maintenance planner.

Action $A_c(k)$: At each decision epoch k , a maintenance action is chosen for component c from:

$$A_c(k) = \begin{cases} 0, & \text{if No Action,} \\ 1, & \text{if Minor Repair,} \\ 2, & \text{if Major Repair,} \\ 3, & \text{if Replacement} \end{cases} \quad (6.6)$$

A richer POMDP formulation could include inspection explicitly in the action space. In such a formulation, an inspection would not directly restore component condition, but would provide an additional observation that reduces uncertainty regarding the hidden degradation state. The decision-maker would then have to balance the expected value of obtaining more accurate health information against the inspection cost and, in an offshore setting, the vessel, technician, and weather-window requirements needed to obtain that information. This would couple condition monitoring and maintenance planning through an explicit value-of-information trade-off and would be particularly relevant for components whose condition is observed primarily through periodic offshore inspection rather than continuous monitoring.

The change in state at the next timestep $k + \Delta k$, as an effect of each action and the past state, is modeled as:

$$\text{RUL}_{\text{true},c}(k + \Delta k) = \begin{cases} \text{RUL}_{\text{true},c}(k) - \Delta k, & \text{if No Action,} \\ \text{RUL}_{\text{true},c}(k) - \Delta k + 0.3 \text{RUL}_{\text{true},c}(0), & \text{if Minor Repair,} \\ \text{RUL}_{\text{true},c}(k) - \Delta k + 0.5 \text{RUL}_{\text{true},c}(0), & \text{if Major Repair,} \\ \text{Weibull(shape, scale),} & \text{if Replacement} \end{cases} \quad (6.7)$$

where the age reduction as a result of minor and major repairs has been collected from Li et al. (2021).

Reward $R_c(k)$: The reward function includes maintenance expenses and downtime losses. Formally, for component c at time k :

$$R_c(k + \Delta k) = -(\text{Cost}_{\text{maint},c} + \text{Cost}_{\text{downtime},c}(k)) \quad (6.8)$$

where $Cost_{maint,c}$ includes material, labor, transportation costs, and the planned downtime required to execute the selected maintenance action, during which the turbine must be temporarily shut down. This shutdown duration corresponds to T_{minor} , T_{major} , or $T_{replacement}$, depending on whether the action is a minor repair, major repair, or replacement, respectively. In contrast, $Cost_{downtime,c}(k)$ accounts for unplanned production losses caused by component failures, specifically the interval between the fault occurrence and its resolution on day $k + T_{maint}/24$.

We normalize the reward using a zero-centered exponential moving average (EMA) to stabilize learning and reduce variance during policy updates. At each time step k , for a component c , we compute a running mean of the rewards using:

$$\mu_c(k) = (1 - \alpha) \cdot \mu_c(k - \Delta k) + \alpha \cdot R_c(k) \quad (6.9)$$

where $\mu_c(k)$ is the running mean, $R_c(k)$ is the immediate reward, and $\alpha \in (0, 1]$ is a smoothing factor controlling how much the current reward contributes to the mean.

We then subtract this running mean from the current reward to obtain the normalized reward:

$$\text{NormReward}_c(k) = R_c(k) - \mu_c(k) \quad (6.10)$$

This normalization approach helps to dampen the impact of outliers or sudden spikes in the reward, providing a smoother signal that leads to more stable policy updates during training.

With the environment defined as a POMDP, a DRL agent is employed to learn an optimal maintenance policy. At each time step k , the agent observes the system state through the observation vector:

$$\mathbf{O}(k) = \{O_c(k) \mid c = 1, \dots, N\} \quad (6.11)$$

where N is the total number of components. The agent then selects actions:

$$\mathbf{A}(k) = \{A_c(k) \mid c = 1, \dots, N\} \quad (6.12)$$

to maximize the cumulative reward over the whole simulation horizon T :

$$\max_{\pi} \left[\sum_{t=1}^T \sum_{c=1}^N R_c(k) \right] \quad (6.13)$$

Since each component has four discrete actions, the overall action space is multi-discrete. This complexity motivates the selection of an algorithm capable of handling such high-dimensional decision spaces, in our case a PPO algorithm is implemented to select a policy defined as follows:

$$\pi(\mathbf{A}(k) \mid \mathbf{O}(k)) \quad (6.14)$$

PPO is selected as the policy-learning algorithm based on the characteristics of the maintenance decision problem discussed in Chapter 5. First, the action is multi-discrete: for each component, the agent must select among no action, minor repair, major repair, and replacement. The resulting joint action space grows combinatorially with the number of components. A standard value-based

method such as DQN would require representing action values over this large joint space unless an additional action-factorization architecture were introduced. PPO instead directly parameterizes the policy and can represent the component-level categorical action distributions without explicitly enumerating every joint action combination.

Second, maintenance planning is a long-horizon stochastic problem in which individual interventions can have consequences many decision epochs later. PPO constrains successive policy updates through its clipped surrogate objective, reducing the likelihood of destabilizing changes to the policy during training. This property was considered advantageous for the simulated environment, in which stochastic degradation and uncertain prognostic observations produce variable trajectories and rewards. Third, although PPO is an on-policy method and is therefore generally less sample-efficient than off-policy alternatives, this disadvantage is less restrictive in the present study because experience is generated through simulation and multiple environment instances can be executed in parallel. Once trained, the resulting policy maps new observations directly to actions without requiring a new optimization problem to be solved at every decision epoch.

The PPO agent follows an actor–critic architecture, as illustrated in Figure 6.2. Both the actor and critic are parameterized by feedforward multilayer perceptrons (MLPs) consisting of two fully connected hidden layers, each with 64 neurons and ReLU activation functions. The input to both networks is a vector that encodes the current observed state of the system, including the discretized health status of each component. The actor outputs a probability distribution over four discrete maintenance actions (no action, minor repair, major repair, or replacement) for each component, enabling joint action selection across the wind farm. The critic network estimates the scalar value function $V(s)$, representing the expected cumulative cost-to-go, to evaluate the current policy. Both networks are trained using PPO updates: the actor via a policy gradient objective that is clipped in order to avoid updates that are too large, and the critic via a value loss. Policy updates are based on trajectories (i.e., state–action–reward transitions) collected during simulated wind farm operation episodes.




Figure 6.2: Proximal Policy Optimization (PPO) architecture

The physical degradation dynamics are factorized at component level. Conditional on its current state and the maintenance action applied to it, the evolution of one component does not directly depend on the degradation state of another component. This is a deliberate modelling abstraction: mechanical dependencies, load-sharing effects, correlated degradation, and common-cause failures are not represented in the transition model. Explicitly representing such dependencies would require either empirically calibrated joint transition models or additional state variables describing common operating, loading, or environmental conditions. Graph-based state representations, as discussed in Chapter 5, provide another possible approach for embedding known relationships between turbines or components.

The maintenance decisions nevertheless remain coupled through their system-level consequences. Each turbine is modelled as a series system, such that failure of any critical component causes turbine-level production loss, and the costs and downtime associated with maintenance actions contribute jointly to the farm-level objective. Where maintenance actions are grouped within the same intervention cycle, campaign-level cost consequences may also be shared. These are economic and decision-level couplings. Dynamic competition for vessels, technicians, and spare parts is not represented in the present DRL formulation, since these resources are assumed to be secured within the planning lead time described in Section 6.3. Extending these quantities into explicit state variables and constraints would allow future formulations to represent other resource-induced dependencies between otherwise physically independent components.

Within the PPO framework, the present implementation deliberately adopts a feed-forward policy of the form $\pi(A(k) \mid O(k))$. This provides a controlled baseline for evaluating whether the current uncertainty-aware prognostic observation, together with the multi-level maintenance action space, contains sufficient information for learning useful maintenance behaviour without introducing an additional state-estimation architecture. Compared with recurrent or explicit belief-state policies, the feed-forward architecture also reduces model complexity and training requirements, which is advantageous for establishing and diagnosing the behaviour of the initial learning-based demonstrator.

A further consideration is the synthetic nature of the prognostic observation model used in this thesis. As described in Chapter 3, prediction errors are generated according to a prescribed time-varying uncertainty function whose standard deviation progressively contracts as failure approaches. Providing a recurrent policy with the complete history of such observations could allow the agent to exploit regularities specific to the synthetic data-generation mechanism and reconstruct the latent degradation state more accurately than would necessarily be possible from real operational prognostic data. In that case, improvements obtained from memory could partly reflect learning the structure of the synthetic generator rather than the general value of temporal information for maintenance decision-making. Restricting the initial policy to the current observation therefore helps isolate the maintenance-policy learning problem from the additional problem of latent-state inference.

In order to compare the performance of our trained PPO model with baseline policies, we benchmark the proposed DRL-based maintenance strategy against two traditional policies and a simplified DRL variant to assess the impact of multi-level maintenance actions. The first baseline, *Corrective*, operates purely reactively: maintenance is scheduled only when a component fails, in which case the component is immediately replaced. The second baseline, referred to as the *Age-Based* policy, was adapted from the model proposed in Li et al. (2021). Here, maintenance is triggered when any component fails, initiating a maintenance cycle in which all components are

evaluated. For each component, the decision to perform maintenance and the type of action taken is based on its predicted usage level at the time of the failure event. This usage level is defined as the ratio of elapsed life to total estimated life. Based on this value, the policy applies one of three maintenance levels (minor repair, major repair, or replacement). The thresholds used were optimized via a brute-force grid search using a discrete event simulation to minimize total cost over multiple Monte Carlo episodes.

To isolate the value of allowing multiple repair levels, we also train a PPO agent using the same environment and architecture, but with a binary action space limited to either doing nothing or replacing a component. Comparing it with the full multilevel PPO policy allows us to quantify the performance gains attributable to having richer action granularity.

This benchmarking framework enables us to assess the benefits of the proposed DRL approach not only against traditional strategies, but also against constrained DRL formulations, illustrating the value of both intelligent scheduling and fine-grained decision flexibility.

6.3 Experimental setup

In this section, we describe the simulation environment and the case scenario used to evaluate the proposed DRL framework. The environment models the O&M of an offshore wind farm consisting of three 5 MW wind turbines located 25 km off the coast of the Netherlands. Specifications of the site, wind turbines and cost functions are adapted from Chapter 3. Each turbine is assumed to be a series system of its critical components, meaning that a failure in any one of these components leads to the entire turbine being inoperable.

Vessels and spare parts are assumed to be available whenever a maintenance action is selected by the agent. This assumption follows from the planning structure of the environment, where maintenance activities are planned with a minimum lead time of 20 days. The lead time is considered sufficient to secure the required access asset and the necessary spare parts. Therefore, vessel unavailability and spare-part stockouts are not included as separate state variables in the DRL formulation. Operational feasibility is instead represented through the weather-driven accessibility conditions inherited from the modelling environment described in Chapter 3.

We consider five critical components per turbine: blades, bearing, gearbox, generator and pitch system. Their RUL is modeled using a Weibull distribution, with parameters (shape and scale) as in Table 6.1.

The maintenance actions available for each component are *No Action*, *Minor Repair*, *Major Repair*, and *Replacement*. Each action is associated with a fixed intervention duration and a fixed maintenance cost, as summarized in Table 6.2. The intervention durations are adapted from the

Table 6.1: Weibull Parameters for Each Critical Component, collected from Li et al. (2021)

Component	Shape	Scale (days)
Blades	3	1847
Bearing	2	1811
Gearbox	3	1477
Generator	2	1594
Pitch System	3	1144

Table 6.2: Maintenance times and fixed costs for each component, adapted from Carroll et al. (2016) and Li et al. (2021).

Component	$T_{\text{replacement}}$ [h]	T_{major} [h]	T_{minor} [h]	$\text{Cost}_{\text{repl}}^{\text{fix}}$ [EUR]	$\text{Cost}_{\text{major}}^{\text{fix}}$ [EUR]	$\text{Cost}_{\text{minor}}^{\text{fix}}$ [EUR]
Blades	288	18	9	285040	23807	8898
Bearing	36	12	6	43755	10455	3982
Gearbox	231	16	8	249511	25190	9360
Generator	81	14	7	89698	14072	5321
Pitch System	36	12	6	38755	9205	3532

empirical offshore repair-time data reported by Carroll et al. (2016), while the action hierarchy and component-specific maintenance modelling follow the benchmark formulation introduced in Chapter 3. The fixed intervention costs combine the component replacement costs adopted from Li et al. (2021) with the economic assumptions of the case study.

To emulate realistic condition monitoring data, we adopt the synthetic RUL prediction framework introduced in Chapter 3. In this approach, each component’s predicted RUL is modeled as a noisy estimate of its true RUL, with uncertainty that decreases over time as failure approaches. Specifically, the prediction error is drawn from a normal distribution whose standard deviation shrinks sigmoidally from a maximum to a minimum bound, capturing the increasing confidence in predictions as more degradation information becomes available. The resulting distributions are used to construct discrete health observations for each component (Equation 6.5), forming the input to the agent.

Table 6.3 presents the training hyperparameters used for the PPO agent. The learning rate, clipping range, and batch size follow standard values commonly used in the default configuration of the Stable-Baselines3 PPO implementation (Schulman et al., 2017). The number of total timesteps and parallel environments, and the reward smoothing factor α were selected iteratively to ensure stability of the training process and convergence of the policy within a reasonable training duration. These experiments were conducted on a laptop equipped with a 12th Gen Intel(R) Core(TM) i7-1265U CPU and 16 GB RAM. Although we did not perform an exhaustive hyperparameter grid search, we validated this setup through empirical trial-and-error runs.

The complexity of the problem is driven by the high-dimensional and combinatorial nature of the observation and action spaces. Each of the three turbines has five critical components, resulting

Table 6.3: Training Hyperparameters for PPO (Schulman et al., 2017)

Parameter	Value
Total Timesteps	1,000,000
Number of Environments	8
Number of Steps per Update	1024
Learning Rate	3×10^{-4}
Discount Factor (γ)	1 (no discounting)
PPO Clipping Range	0.2
Batch Size	64
α (Reward Smoothing)	0.01
Prediction Horizon H	180 days
Time step Δk	20 days
Simulation length	10 years

in 15 components to monitor and control. For each component, the agent observes a discretized health state and must select one of four possible actions: no action, minor repair, major repair, or replacement. This leads to a discrete action space of size $4^{15} = 1.07 \times 10^9$, assuming independent actions for each component. The observation space is similarly structured, with each component's health classified into four discrete levels, resulting in an observation space of size 4^{15} . The large size of both spaces makes exact dynamic programming methods intractable and highlights the need for centralized policy optimization methods.

6.4 Results and comparative study

This section evaluates the maintenance policies developed and used throughout the thesis. The comparison includes the Age-based policy, the purely Corrective policy, the multi-level DRL policy, the Binary DRL policy, and the MILP rolling-horizon optimization model. The Binary DRL policy is included only for the 3-turbine case, where it was trained with the same learning structure as the DRL policy but restricted to two actions: no action and replacement. The 5- and 10-turbine comparisons include Age-based, Corrective, DRL, and MILP. The purpose of the section is to verify that the learning-based policy is grounded in the benchmark environment defined in Chapter 4 and behaves consistently with the maintenance logic defined in this chapter. Secondly, it connects the learning-based and optimization-based workstreams by comparing how the different policies balance maintenance expenditure, downtime reduction, failure prevention, and intervention effort.

6.4.1 Training

Figure 6.3 reports the training behaviour of the DRL agent. The total loss and value loss stabilise after the initial learning phase, while the maintenance cost decreases over training. This indicates that the policy learns to associate the discretised health observations with preventive maintenance actions that reduce the long-term cost of failures and downtime.

A Monte Carlo convergence analysis is conducted to verify that the reported performance metrics are stable across repeated stochastic evaluations. Figure 6.4 shows the moving average of the total cost over repeated simulations. Considering the relative stability over multiple iterations, all results are averaged over 100 episodes.

6.4.2 Model comparison

This subsection compares the learning-based policy developed in this chapter with the benchmark maintenance policies and with the optimization-based model developed in Chapter 4. For brevity, the learning-based policy and the optimization-based model are referred to as DRL and MILP, respectively, in the tables, figures, and subsequent discussion. This comparison is included to explicitly connect the learning-based framework of this chapter with the optimization-based model developed in Chapter 4. Both approaches are evaluated under the modeling environment introduced in Chapter 3, using the same component degradation logic, maintenance action taxonomy, cost structure, and stochastic evaluation procedure. The comparison is therefore not intended only as a performance ranking, but also as a way to examine how the two modeling paradigms make different




Figure 6.3: DRL training performance.




Figure 6.4: Monte Carlo convergence.

trade-offs between preventive intervention, failure avoidance, downtime reduction, and maintenance effort.

A detailed report of the 3-turbine comparison is presented in Table 6.4. The learning-based policy obtains the lowest total cost, with 1.24 M€, followed closely by the optimization-based model and Binary DRL. Corrective maintenance is the most expensive policy because it reacts only after failure and therefore accumulates a large downtime cost. The Age-based policy improves substantially over Corrective, but still produces more failures and higher total cost than DRL.

Table 6.4: Three-turbine policy comparison.

Policy	Total cost [M€]	Std. total [M€]	Maint. cost [M€]	Down. cost [M€]	Failures [-]	Cycles [-]	Minor [-]	Major [-]	Repl. [-]
Age-based	2.27	0.38	1.32	0.95	16.30	67.16	8.90	83.20	16.22
Corrective	8.73	1.14	1.08	7.65	34.04	29.14	0.00	0.00	33.70
DRL	1.24	0.14	1.11	0.13	2.90	22.36	64.16	9.78	28.24
Binary DRL	1.81	0.14	1.81	0.01	7.08	22.54	0.00	0.00	51.88
MILP	1.79	1.11	1.42	0.37	20.12	21.60	194.00	25.38	20.16

The comparison with Binary DRL isolates the value of the multi-level action space. Both learning-based policies use state-dependent decisions, but Binary DRL can only replace components. As a result, it has fewer available ways to react to intermediate degradation states. It performs better than Corrective maintenance because it can still act preventively, but it remains more expensive than the multi-level DRL policy. The reason is that replacement is a strong and costly intervention: when the policy cannot select minor or major repairs, it must either delay action or perform a full renewal. The multi-level DRL policy avoids this binary trade-off by using minor and major repairs as intermediate actions.

The comparison with the optimization-based model provides a different perspective. Whereas the learning-based policy applies a learned mapping from observed health states to maintenance actions, the optimization-based model repeatedly solves an explicit optimization problem over a rolling horizon. The two methods therefore use the same modeling environment, but differ in how decisions are generated. The learning-based policy emphasizes adaptive state-dependent behaviour learned from repeated simulation, while the optimization-based model preserves an explicit cost-minimisation structure and campaign-level discipline at each planning epoch.

Figure 6.5 shows the total-cost comparison for all wind farm sizes. In the 3-turbine case, the learning-based policy is the most cost-efficient. Binary DRL and MILP are close in total cost, but for different reasons: Binary DRL reduces failures through learned preventive replacements, while MILP controls intervention effort through explicit optimization and maintenance grouping. Corrective maintenance remains dominated because its low decision complexity is offset by high failure-related downtime.

The cost and reliability indicators are summarized across all cases in Table 6.5. For 3 and 5 turbines, the learning-based policy achieves lower total cost than the optimization-based model. In the 3-turbine case, DRL reduces the total cost by approximately 31% relative to MILP. In the 5-turbine case, the reduction remains approximately 27%. These improvements are mainly caused by lower downtime cost and fewer unexpected failures. This indicates that, at small scale, the learning-based policy is effective at using probabilistic health observations to anticipate failures and reduce downtime.

Maintenance and downtime costs are shown in Figure 6.6. The Corrective policy has relatively low maintenance cost in direct terms, but this apparent saving is outweighed by downtime. In contrast, the learning-based policy deliberately increases preventive intervention effort to avoid failure-induced production losses. This is beneficial at 3 and 5 turbines, where the additional maintenance cost remains limited and the avoided downtime is substantial.

The 10-turbine case changes the interpretation. The learning-based policy still reduces failures compared with the optimization-based model, but the total cost increases to 7.63 M€, compared




Figure 6.5: Total cost by policy and wind farm size.

Table 6.5: Cost and reliability summary.

WTs	Policy	Total cost [M€]	Std. total [M€]	Maint. cost [M€]	Down. cost [M€]	Failures [-]
3	Age-based	2.27	0.38	1.32	0.95	16.30
3	Corrective	8.73	1.14	1.08	7.65	34.04
3	DRL	1.24	0.14	1.11	0.13	2.90
3	Binary DRL	1.81	0.14	1.81	0.01	7.08
3	MILP	1.79	1.11	1.42	0.37	20.12
5	Age-based	3.76	0.44	2.17	1.59	27.36
5	Corrective	14.52	1.40	1.79	12.73	56.60
5	DRL	2.12	0.19	1.98	0.14	3.58
5	MILP	2.92	1.00	2.32	0.60	34.20
10	Age-based	7.37	0.49	4.22	3.15	54.40
10	Corrective	28.81	1.88	3.62	25.19	113.02
10	DRL	7.63	0.70	5.58	2.05	23.46
10	MILP	5.64	0.95	4.48	1.16	68.96

with 5.64 M€ for MILP. The cost advantage therefore changes as the wind farm grows. This is not because the learned policy stops preventing failures; rather, it prevents failures through an increasingly intensive maintenance pattern. The result is a higher maintenance cost and, because many interventions are performed, a higher maintenance-related downtime cost. The comparison with MILP is therefore important for the coherence of the thesis: it shows that the learning-based policy is stronger in adaptive failure prevention, while the optimization-based model remains more disciplined in controlling intervention frequency and campaign structure as the system grows.

Figure 6.7 shows the average number of failures per episode. Corrective maintenance naturally produces the highest number of failures because every replacement is triggered by a failure event. Age-based maintenance reduces failures relative to Corrective, but it remains tied to fixed intervention logic and does not use the probabilistic health state as flexibly as the learning-based policy. DRL produces the fewest failures in all wind farm sizes for which it is evaluated. This confirms that




Figure 6.6: Maintenance and downtime cost.




Figure 6.7: Failures by policy and wind farm size.

the learned policy is strongly reliability-oriented and uses the available health observations to act before failure.

Maintenance cycles and action counts are listed in Table 6.6. It is worth mentioning that Corrective and Binary DRL are replacement-only policies by construction. Age-based and DRL both use a mixed action set, while MILP selects actions through an explicit rolling-horizon objective. This distinction is central for interpreting the comparison with Chapter 4: the optimization-based model explicitly optimizes maintenance timing and grouping, whereas the learning-based policy selects a component-level action vector from the learned policy.

The same action mix is shown in Figure 6.8. The learning-based policy uses all three intervention levels. At small scale, this is beneficial because minor and major repairs allow the agent to prevent

Table 6.6: Maintenance effort and action mix.

WTs	Policy	Cycles	Minor	Major	Repl.
3	Age-based	67.16	8.90	83.20	16.22
3	Corrective	29.14	0.00	0.00	33.70
3	DRL	22.36	64.16	9.78	28.24
3	Binary DRL	22.54	0.00	0.00	51.88
3	MILP	21.60	194.00	25.38	20.16
5	Age-based	87.28	14.98	139.18	27.34
5	Corrective	43.38	0.00	0.00	55.90
5	DRL	35.78	105.48	22.56	51.22
5	MILP	22.08	326.46	37.74	34.56
10	Age-based	104.88	28.70	280.90	54.22
10	Corrective	70.80	0.00	0.00	112.22
10	DRL	122.00	836.40	206.18	118.36
10	MILP	23.80	647.54	72.66	69.54




Figure 6.8: Maintenance action mix.

failures without immediately committing to full replacement. Binary DRL lacks this flexibility and therefore relies entirely on replacements. This explains why it is useful as a diagnostic comparison: it shows that the benefit of the learning-based framework does not come only from learning, but also from allowing the agent to choose among operationally meaningful intervention depths.

The scaling limitation of the current learning-based formulation becomes visible at 10 turbines. The DRL policy initiates 122 maintenance cycles and performs a very large number of minor and major repairs. MILP also performs many component-level actions, but concentrates them into fewer maintenance cycles. This distinction is central: offshore O&M cost is not determined only by the number of repaired components, but also by how interventions are grouped into campaigns. The MILP formulation explicitly structures this grouping, whereas the learned policy acts through a component-level action vector and can overuse low-cost actions when the number of components increases.

Figure 6.9 illustrates this trade-off between failure prevention and intervention effort. DRL is consistently positioned at lower failure levels than MILP, confirming its preventive behaviour.




Figure 6.9: Failure-prevention and intervention effort.

However, for 10 turbines, it shifts strongly toward higher maintenance-cycle counts. This explains why the reliability improvement does not translate into lower total cost at larger scale.

The total cost normalised by the number of turbines is shown in Figure 6.10. This view is useful because it separates absolute system size from policy scalability. Corrective maintenance remains the least attractive option at all scales. Age-based maintenance is more stable, but does not exploit health information sufficiently to minimise failures. MILP shows the most stable cost-per-turbine trend as the wind farm grows. DRL is efficient at 3 and 5 turbines, but its cost per turbine increases sharply at 10 turbines, reflecting the excessive intervention effort discussed above.

Overall, the results show that the models are grounded in the expected maintenance logic. Corrective maintenance minimises planning sophistication but produces large downtime costs. Age-based maintenance improves over Corrective by introducing preventive logic, but remains less responsive to component-specific health information. Binary DRL confirms the value of learning a




Figure 6.10: Total cost per turbine.

state-dependent policy, but also shows the limitation of restricting the action space to replacement only. The full learning-based policy performs best at small scale because it combines health-aware decisions with a richer action space. The optimization-based model becomes stronger at larger scale because it preserves campaign-level structure and avoids excessive intervention frequency.

6.5 Conclusions

This chapter addressed Research Question Q5 by formulating offshore wind O&M decision-making as a learning problem and evaluating a DRL-based maintenance policy under uncertain component health information. Building on the modeling environment introduced in Chapter 3 and the learning requirements identified in Chapter 5, the proposed framework represented maintenance planning as a POMDP in which the true RUL of each component is hidden and the agent observes discretized health states derived from probabilistic RUL predictions. This formulation reflects the practical role of prognostics in O&M planning: RUL estimates do not provide perfect knowledge of component condition, but they provide uncertain information that can be used to guide preventive, corrective, and opportunistic maintenance decisions.

The main methodological contribution of this chapter is the development of a PPO-based DRL maintenance policy with a multi-level action space. Rather than restricting the agent to a binary maintain-or-do-not-maintain decision, the proposed framework allows the selection of no action, minor repair, major repair, and replacement. Minor and major repairs may provide economically attractive intermediate options between inaction and full replacement, particularly when degradation is detected early or when full replacement is not yet justified.

The results show that the DRL policy can learn an adaptive maintenance strategy that reduces total cost and substantially limits unexpected failures compared with conventional Age-based and Corrective benchmark strategies. The learned policy achieves this by using probabilistic health observations to schedule preventive interventions before failures occur, thereby reducing downtime and avoiding the failure-driven cost escalation observed under Corrective maintenance. The comparison with the Binary DRL policy further confirms the value of the multi-level action space. Both agents use the same learning paradigm and health information, but Binary DRL is restricted to no action or replacement. Its performance shows that learning alone is not sufficient: the action space must also represent realistic maintenance options. This supports one of the main modeling requirements identified in Chapter 5, that learning-based O&M frameworks should move beyond binary repair abstractions.

The comparison with the optimization-based framework developed in Chapter 4 provides the main synthesis of the chapter. For the smaller wind-farm cases, DRL achieves lower total cost than MILP by suppressing failures and downtime more effectively. This indicates that a learned policy can exploit uncertain health observations to produce adaptive preventive behaviour without resolving an optimization problem at every decision epoch. However, the results also show that this advantage does not automatically scale. In the 10-turbine case, DRL remains effective at reducing failures, but it does so through a much larger number of maintenance interventions. The resulting increase in maintenance effort and downtime causes DRL to become more expensive than MILP. By contrast, MILP preserves a better economic discipline as the number of turbine-component decisions increases.

These results also refine the interpretation of the optimization limitations identified in Chapters 4 and 5. The motivation for introducing a learning-based formulation was not that DRL should necessarily dominate MILP in all benchmark cases, but that DRL provides a different way of addressing sequential decision-making under uncertainty and partial observability. The comparison in this chapter shows that these advantages must be interpreted together with the current limitations of the learning formulation.

First, with respect to scalability, the MILP results in Chapter 4 show that optimization remains highly competitive when the problem structure is sufficiently controlled and when campaign-level decisions can be represented explicitly. The present DRL implementation scales less favourably because a single centralized agent selects component-level actions for an increasing number of turbine-component pairs. As the wind farm grows, this leads to a much larger action space and to an increased tendency to overuse low-cost preventive actions. Therefore, the scalability limitation does not apply only to MILP; it also appears in the current single-agent DRL architecture, although in a different form. Furthermore, with respect to uncertainty handling, DRL is able to learn preventive behaviour from uncertain RUL-derived observations and substantially reduce failures. However, the results show that reducing failures is not equivalent to minimizing total O&M cost. In the 10-turbine case, the learned policy reacts strongly to degradation risk and achieves low failure counts, but this reliability-oriented behaviour increases the number of interventions and maintenance-related downtime. This indicates that the reward structure and action representation must be further refined so that the agent balances failure prevention against campaign efficiency and maintenance effort. With respect to adaptivity, the DRL policy has the advantage that, once trained, it can map updated observations directly to actions without solving a new optimization problem at each decision epoch. This is valuable for online decision support. Nevertheless, the MILP remains superior in representing explicit campaign structure, feasibility constraints, and cost trade-offs. In the present benchmark, this structure allows MILP to maintain more stable cost performance as the number of turbines increases. With respect to interpretability and constraint control, MILP provides transparent schedules and explicit constraint satisfaction, while the DRL policy requires behavioural analysis through action counts, failures, maintenance cycles, and cost components. This limits the immediate operational deployability of DRL, especially in safety-critical and high-cost offshore settings where planners must justify maintenance decisions.

An additional limitation concerns the treatment of partial observability. The POMDP formulation distinguishes between the latent true RUL and the uncertain prognostic observation available to the decision-maker, but the feed-forward PPO policy implemented in this chapter acts only on the latest observation. It does not maintain a belief state or exploit the complete observation–action history. The results should therefore be interpreted as demonstrating the performance of a reactive learning policy under uncertain health information, rather than as a solution to the general POMDP problem described in Section 5.4.2. Similarly, prognostic information is assumed to be available at each decision epoch, whereas a richer formulation could include inspection as a separate action and explicitly optimise the value of obtaining additional condition information.

More broadly, the offshore O&M challenges identified in Figure 5.1 are represented at different levels of abstraction in the present model. Stochastic component degradation and uncertain RUL observations represent health uncertainty and failure risk, and the common environment retains weather-driven accessibility effects. However, component degradation is factorized, abrupt or common-cause failure mechanisms are not modelled, and detailed resource competition, routing, vessel positioning, technician allocation, and spare-part inventory are outside the DRL state and

action spaces. The results therefore provide evidence that learning-based policies can exploit uncertain prognostic information and multi-level maintenance actions within a controlled offshore-wind benchmark; they should not be interpreted as demonstrating that the current PPO implementation captures the full operational complexity of offshore wind O&M.

Therefore, the comparison in Figure 6.5 should not be read as a failure of the learning-based paradigm. DRL is promising for adaptive failure prevention under uncertain health information, whereas MILP remains strong when explicit constraint handling, campaign grouping, and economic discipline dominate the decision problem. The results indicate that future learning-based offshore O&M models show promise but should incorporate more structure, for example through hierarchical policies, action masking, multi-agent decomposition, campaign-level decision layers, or hybrid optimization-learning architectures.

In conclusion, the learning-based policy is effective at using uncertain health observations to prevent failures, especially in smaller systems. The optimization-based model is more disciplined as the wind farm size increases, because intervention grouping and cost trade-offs are explicitly encoded in the optimization problem. This comparison strengthens the coherence of the thesis by showing how the optimization-based and learning-based approaches address the same O&M decision problem from different methodological perspectives. Learning provides adaptive preventive behaviour under partial observability, while optimization provides transparency, structure, and control over intervention effort. Future learning-based extensions should therefore focus on scalability mechanisms such as hierarchical policies, action masking, campaign-level decision layers, or hybrid optimization-learning architectures.

These findings support the comparison between optimization-based and learning-based O&M planning, nevertheless, several limitations should be acknowledged. First, the scalability analysis remains limited to relatively small wind farms. The results for 10 turbines already indicate that the current single-agent DRL formulation may diverge as the number of assets increases. Larger wind-farm cases would therefore require additional architectural structure, such as hierarchical decision-making, action masking, multi-agent formulations, or hybrid optimization-learning mechanisms. Second, while the framework is aligned with the modeling environment defined in Chapter 3, the current implementation still simplifies some offshore operational details, particularly weather-window execution, vessel availability, routing, and inter-turbine logistical coupling. Third, the DRL policy remains less transparent than the MILP solutions of Chapter 4. Although the analysis of action mixes, failures, maintenance cycles, and cost components provides insight into the learned behaviour, further work is required to make DRL policies auditable and acceptable for operational decision support.

These limitations define the bridge toward the final synthesis of the thesis. Chapter 7 consolidates the findings from the optimization-based and learning-based workstreams, comparing their respective strengths and weaknesses in terms of feasibility, adaptivity, scalability, interpretability, and robustness.

Chapter 7

Conclusions and Future Research

This thesis investigated how maintenance decision-making for offshore wind farms can be improved by combining condition and prognostic information with explicit representations of offshore accessibility, logistics, and cost consequences. Offshore wind maintenance decisions are economically significant because failures, access delays, vessel mobilisation, and lost production directly affect lifetime cost and availability. At the same time, condition monitoring and prognostic methods increasingly provide information about component health and RUL. However, this information only becomes valuable when it can be translated into feasible maintenance actions that account for weather-driven accessibility, logistics, shared resources, and uncertainty. The central motivation of the thesis is therefore to move towards decision-making frameworks that use prognostic information to support adaptive and economically meaningful offshore O&M planning.

The central research question of the thesis is:

What are the trade-offs between optimization-based and learning-based decision-making for offshore wind maintenance?

To answer this question, the thesis developed two methodological decision-support paradigms. Chapter 2 reviewed the state of the art in offshore wind O&M decision support and identified the need for more integrated models that jointly represent asset health, accessibility, logistics, and maintenance timing over multiple time horizons. Chapter 3 then established an offshore wind O&M environment, including a fixed-bottom wind farm representation, critical components, access assets, metocean-driven accessibility, synthetic RUL predictions, maintenance actions, and benchmark outputs. Chapter 4 formulated tactical maintenance planning as a rolling-horizon MILP problem and demonstrated the operational value of incorporating probabilistic RUL information into optimization-based scheduling. Chapter 5 analysed the limitations of this optimization paradigm and reviewed the emerging role of DRL for offshore wind O&M. Finally, Chapter 6 formulated offshore wind maintenance planning as a POMDP and developed a PPO-based DRL framework with multi-level maintenance actions.

This final chapter consolidates these findings. Section 7.1 answers the research questions. Section 7.2 synthesises the trade-offs between optimization-based and learning-based maintenance planning. Section 7.3 summarises the scientific contributions of the thesis. Section 7.4 discusses practical implications for offshore wind O&M decision support. Section 7.5 reflects on the main

limitations of the work and outlines future research directions, and Section 7.6 provides final remarks.

7.1 Answers to the research questions

Q1: What are the state-of-the-art, state-of-practice, and research gaps in offshore wind O&M decision support, particularly regarding the integration of accessibility, logistics, and condition-informed planning?

Chapter 2 showed that offshore wind O&M decision support has developed across three main decision-making layers: strategic, tactical, and operational. Strategic models support long-term decisions such as maintenance policy design, fleet concepts, and lifecycle cost evaluation. Tactical models address medium-term planning questions such as fleet selection, spare-parts policies, and campaign planning. Operational models focus on short-term routing, scheduling, diagnostics, and response to failures. Across these layers, the literature provides a rich set of modeling approaches, including simulation, mathematical optimization, stochastic programming, Markov models, Petri nets, routing heuristics, and emerging learning-based methods.

However, the review also identified several gaps that motivated the rest of the thesis. First, offshore wind O&M models often treat decision-making processes that belong to different time-horizons separately, even though these elements are coupled in practice. Second, although condition monitoring and prognostic methods are increasingly able to estimate component health and RUL, their outputs are still not consistently integrated into maintenance decision-making. Many existing planning models continue to rely on static failure assumptions, fixed age thresholds, or simplified health indicators, rather than using evolving and uncertain RUL information to update maintenance timing and intervention selection. Third, real operational, failure, maintenance, and prognostic datasets remain scarce in offshore wind. This limits direct validation of advanced O&M decision-support models and motivates the use of controlled benchmark environments in which optimization- and learning-based methods can be evaluated under consistent assumptions.

The answer to Q1 is therefore that offshore wind O&M decision support is methodologically mature in several separate areas, but still lacks integrated, condition-informed, and logistics-aware decision frameworks that can be compared consistently across modeling paradigms. This gap provided the basis for the modeling environment developed in Chapter 3 and for the optimization- and learning-based models developed in Chapters 4 and 6.

Q2: How to model the dynamics of offshore wind farm O&M to support maintenance decision-making?

Chapter 3 answered Q2 by defining an offshore wind O&M modeling environment so that alternative maintenance planners can be compared, they must interact with the same representation of the wind farm, weather, logistics, component degradation, maintenance actions, and performance metrics.

The environment developed in Chapter 3 consists of five main elements. First, it defines a modular fixed-bottom offshore wind farm representation, including turbine layout, distances, power

production, and turbine-level availability. Second, it identifies a tractable set of critical components whose failures drive turbine unavailability and maintenance demand. Third, it models transportation and access assets, including CTVs, SOVs, jack-up vessels, and helicopters, together with representative operability thresholds. Fourth, it generates synthetic metocean-driven accessibility states using a month-dependent Markov-chain process calibrated on historical wind and wave data. Fifth, it generates synthetic RUL predictions by combining Weibull-based latent lifetimes with time-varying prognostic uncertainty.

The environment also defines maintenance actions as multi-level interventions: minor repair, major repair, and replacement. These actions differ in restorative effect, duration, resource requirements, and cost implications. Minor and major repairs are modelled as imperfect maintenance through a virtual-age formulation, whereas replacement resets the component to an as-good-as-new state by sampling a new lifetime. The resulting benchmark stores both hidden true states and observable prognostic information, enabling optimization- and learning-based methods to operate under the same information structure.

The answer to Q2 is therefore that a decision-support model for offshore wind O&M should be built on a modular environment that links component health, uncertain prognostics, weather-driven access, marine logistics, maintenance effects, and cost consequences. In this thesis, that environment forms the common interface through which different decision-support paradigms can be evaluated on a like-for-like basis.

Q3: How can O&M planning be formulated as an optimization model that incorporates condition and prognostic information?

Chapter 4 answered Q3 by formulating offshore wind O&M planning as a rolling-horizon MILP problem. The optimization model receives probabilistic RUL information from the environment and converts it into day-wise failure probabilities over the planning horizon. These probabilities are then used to evaluate the expected cost and lifetime consequences of maintenance actions. The model jointly decides whether to initiate a maintenance cycle, which components to maintain, which action type to apply, and when to schedule the intervention, subject to accessibility, resource, and feasibility constraints.

The optimization model demonstrated that probabilistic prognostic information can be translated into explicit maintenance decisions. Compared with an age-based benchmark, the prognostic-driven strategy reduced yearly O&M cost, with improvements observed in a 50-turbine case, where total yearly O&M cost decreased by 8.7%. This improvement was primarily driven by a reduction in downtime cost, confirming that prognostics are most valuable when they support better timing of interventions under offshore access and logistics constraints. The sensitivity analyses further showed that performance depends on prediction horizon length, prognostic uncertainty, and distance to port, highlighting the interaction between health information and offshore operational drivers.

At the same time, Chapter 4 also showed that optimization-based maintenance planning is sensitive to finite-horizon effects. The oracle-aided analysis demonstrated that a rolling-horizon MILP may miss grouping opportunities just beyond the active planning window and may schedule preventive actions too close to the end of the plant lifetime. The oracle-aided variant reduced yearly downtime and missed grouping opportunities, and eliminated end-of-life regret events in the experiment. These results do not invalidate the MILP formulation; rather, they clarify its structural limitations.

The answer to Q3 is therefore that O&M planning can be formulated as a transparent, constraint-aware MILP that embeds probabilistic RUL information into maintenance timing and action selection. This provides interpretable and operationally feasible schedules, but it remains affected by horizon dependence, modeling simplifications, and the computational burden of repeated re-optimization under uncertainty.

Q4: Which modeling requirements should a learning-based offshore wind O&M framework satisfy to address the limitations identified in the optimization-based framework?

Chapter 5 answered Q4 by analysing the limitations of the optimization-based paradigm and reviewing DRL methods for offshore wind O&M and related maintenance planning problems. The chapter showed that learning-based frameworks are attractive because offshore wind O&M is naturally sequential, stochastic, and partially observable. Decisions must be made repeatedly as health estimates, weather conditions, accessibility, and resource constraints evolve. A learning-based policy can, in principle, map these evolving observations directly to actions without solving a new optimization problem at every decision epoch.

However, Chapter 5 also argued that DRL is not automatically suitable for offshore wind O&M. For a learning-based framework to be credible, it must satisfy several modeling requirements. First, it must account for partial observability: the true health state of components is not directly known, and the decision-maker observes uncertain prognostic information rather than perfect RUL. Second, it must include operationally meaningful actions. Binary maintain/do-not-maintain abstractions are insufficient for offshore wind because real interventions differ in cost, duration, logistics, and restorative effect. Third, the reward function must reflect the relevant trade-offs between preventive maintenance, corrective failures, downtime, lost production, and intervention cost. Fourth, the environment must include enough realism to make the learned policy meaningful, especially with respect to access constraints and maintenance effects.

The answer to Q4 is therefore that a learning-based offshore wind O&M framework should be formulated as a sequential decision problem under partial observability, with state features derived from uncertain prognostics, a multi-level maintenance action space, and a reward structure aligned with long-term O&M cost and downtime consequences. These requirements were implemented in Chapter 6 through a POMDP-based DRL formulation.

Q5: How can offshore wind O&M decision-making be formulated and solved as a learning problem to derive adaptive maintenance policies under offshore constraints?

Chapter 6 answered Q5 by formulating offshore wind O&M planning as a POMDP and solving it using a PPO-based DRL agent. In this formulation, the true component RUL is hidden, and the agent observes discretised health states derived from probabilistic RUL predictions. The action space includes no action, minor repair, major repair, and replacement, allowing the agent to learn both when to intervene and which intervention depth is appropriate. The reward function penalises maintenance cost, downtime, and failure-related consequences, thereby aligning the learning problem with the long-term economic objective of O&M planning.

The results showed that the PPO policy can learn an adaptive maintenance strategy that reduces total cost and substantially limits unexpected failures compared with corrective and age-based

benchmark strategies. The learned policy uses probabilistic health observations to schedule preventive interventions before failures occur, while avoiding the excessive reliance on replacements observed in corrective and binary-action baselines. The comparison with the binary PPO policy is particularly important: both agents use the same learning paradigm and health information, but the multi-level PPO policy performs better because it can exploit a richer and more realistic action space.

The answer to Q5 is therefore that offshore wind O&M decision-making can be formulated as a learning problem by representing maintenance planning as a POMDP in which uncertain prognostic observations guide sequential action selection. A PPO-based agent can learn adaptive maintenance policies from simulation, especially when the action space captures realistic intervention levels. However, the learning-based approach also introduces new challenges, including reduced interpretability, dependence on simulator fidelity, training effort, and the need for further scalability analysis.

7.2 Answer to main research question

The main research question asked what trade-offs exist between optimization-based and learning-based approaches for deriving maintenance policies. The findings of this thesis show that neither paradigm dominates the other in all dimensions. Instead, they provide different forms of decision support and are suitable for different stages of methodological development and practical deployment.

The MILP formulation provides explicit decision variables, constraints, and objective terms. This makes it transparent: a planner can inspect why an action is scheduled, which constraint is binding, how costs are computed, and how resource limitations affect the solution. This is a major advantage in safety-critical and high-cost offshore settings, where operators must justify decisions to asset owners, insurers, contractors, and regulatory stakeholders. The MILP also provides a natural way to include hard feasibility constraints, such as spare-part availability, technician availability, or weather-window requirements. The DRL policy, by contrast, produces decisions through a learned mapping from observations to actions. This allows fast execution once trained, but the internal decision logic is less transparent. Analysis of action frequencies, cost components, and failure reductions can help interpret the policy at an aggregate level, but it does not provide the same direct explanation as a mathematical programme. Therefore, learning-based methods require additional tools for policy interpretation, validation, and safety assurance before they can be used as operational decision-support systems.

The MILP adapts through rolling-horizon re-optimization: as new prognostic information becomes available, the model is solved again with updated inputs. This allows the planner to respond to changing health states and accessibility conditions. However, adaptivity is achieved by repeatedly solving an optimization problem, and the quality of each decision depends on the active planning horizon, scenario representation, and terminal assumptions. The DRL approach learns an adaptive policy directly. Once trained, the policy can react to the observed state without solving a new optimization problem at each decision epoch. This is conceptually well aligned with repeated decision-making under uncertainty. The POMDP formulation also makes the role of imperfect health information explicit: the agent does not require access to the true RUL, but learns

from uncertain prognostic observations. This makes DRL attractive for online decision support, provided that the training environment adequately represents the operational system.

The MILP can incorporate uncertainty through probabilistic inputs, expected-cost calculations, scenarios, or robust formulations. In this thesis, probabilistic RUL information was discretised into failure probabilities over the planning horizon and embedded into the expected cost/lifetime calculations. This provides a tractable and interpretable way to use prognostic uncertainty. However, richer uncertainty representations typically increase the size and complexity of the optimization problem. The DRL framework handles uncertainty differently. Instead of explicitly enumerating future scenarios during each decision, the agent learns from repeated interactions with a stochastic environment. This allows uncertainty to be internalised through training experience. However, this shifts the challenge from optimization tractability to simulator fidelity: if the training environment does not represent the relevant uncertainty mechanisms accurately, the learned policy may not transfer reliably to real operations.

The MILP provides strong guarantees under its assumptions, but its computational complexity can grow rapidly with the number of turbines, components, actions, vessels, time periods, and uncertainty scenarios. Rolling-horizon solution strategies mitigate this burden, but introduce horizon myopia, as shown by the oracle analysis in Chapter 4. The MILP is therefore particularly useful as a transparent benchmark and tactical planner, but scaling it to high-dimensional, highly stochastic, and real-time settings remains challenging. The DRL approach incurs substantial computational effort during training, but policy execution is fast once training is complete. This makes learning-based methods promising for repeated online decisions. However, the DRL implementation in this thesis was demonstrated on a smaller case study than the MILP. Its scalability to utility-scale offshore wind farms, more detailed logistics, multi-vessel routing, and multi-agent coordination remains an important open question. Thus, DRL shifts the scalability challenge from repeated exact optimization to training efficiency, state/action representation, and generalisation.

Both paradigms depend on the quality of the underlying environment. The MILP can include operational details explicitly, but each added detail increases model complexity. The DRL framework can in principle learn from complex environments, but only if the simulator is sufficiently realistic and computationally efficient. In both cases, the modeling environment developed in Chapter 3 is therefore central: it provides a controlled representation of weather-driven access, component degradation, prognostics, maintenance actions, and cost consequences; methodological comparison is only meaningful when different decision-support paradigms are evaluated under common assumptions. The benchmark makes it possible to separate the effect of the decision logic from the effect of the modeling environment.

The trade-off between optimization-based and learning-based approaches can be summarised as follows. Optimization-based methods are transparent, constraint-aware, and well suited for tactical planning when the problem can be formulated at a tractable level of detail. They provide interpretable schedules and are valuable as benchmarks, but they can suffer from horizon dependence, scalability limitations, and difficulty representing deep uncertainty without increasing complexity. Learning-based methods are more naturally suited to adaptive sequential decision-making under uncertainty and can exploit probabilistic health observations through learned policies. They reduce the need for repeated online optimization, but introduce challenges related to interpretability, simulator dependence, training stability, safety assurance, and transferability. Therefore, the central conclusion of this thesis is that optimization-based and learning-based approaches offer different strengths and face different limitations, which makes them suitable for different decision contexts.

Optimization models provide transparency, explicit constraint handling, and interpretable schedules, making them particularly valuable for tactical planning, engineering validation, and settings in which explainability and feasibility guarantees are central. Learning-based methods, by contrast, are better suited to repeated sequential decision-making problems in which policies must adapt to evolving asset states, uncertain prognostic information, and long-term consequences that are difficult to encode explicitly in a finite-horizon optimization model.

7.3 Contributions

This thesis contributes to offshore wind O&M decision support by developing, testing, and critically comparing modeling approaches that connect asset health information, offshore logistics, weather-driven accessibility, and maintenance decision-making. Rather than treating these elements as separate modeling problems, the thesis examines how they interact across the full decision-support chain: from the synthesis of existing methods, to the construction of a benchmark environment, to the evaluation of optimization- and learning-based maintenance policies.

A first contribution is the structured synthesis of offshore wind O&M decision-support literature across strategic, tactical, and operational decision layers. The review clarifies how existing models represent maintenance strategies, metocean conditions, logistics, routing, spare parts, diagnostics, prognostics, and cost. It also identifies recurring gaps in the literature, including the fragmented treatment of asset health and logistics, the limited integration of prognostic information into planning models, weak coupling across decision horizons, and the scarcity of real operational data. This synthesis provides the conceptual foundation for the modeling choices developed in the rest of the thesis and is primarily associated with Borsotti et al. (2026c).

Building on this review, the thesis develops an offshore wind O&M modeling environment for controlled comparison of decision-support methods. The environment defines the wind farm layout, critical component set, access assets, metocean-driven accessibility, RUL prediction structure, maintenance-action framework, cost model, and benchmark outputs. Its main contribution is not only reproducibility, but comparability: different decision-support paradigms can be evaluated under the same assumptions, interfaces, and performance indicators. This makes it possible to distinguish differences caused by the decision method itself from differences caused by inconsistent case-study definitions.

Within this environment, the thesis then formulates prognostic-driven tactical maintenance planning as a rolling-horizon MILP. This contribution shows how probabilistic RUL information can be translated into explicit maintenance schedules under offshore-specific constraints, including access limitations, vessel requirements, maintenance durations, and downtime costs. The results demonstrate that integrating prognostic information into maintenance planning can reduce O&M cost, primarily by reducing downtime and improving the timing of preventive interventions. These contributions are mainly reflected in Borsotti et al. (2026b) and in the earlier MPC-based framework presented in Borsotti et al. (2024).

The optimization-based analysis also provides a further contribution by making the limitations of rolling-horizon planning explicit and measurable. Instead of only reporting aggregate performance, the thesis examines why certain inefficiencies remain, particularly through horizon-boundary effects, missed grouping opportunities, and end-of-life regret. This provides a structured diagnosis

of optimization behaviour under uncertainty and helps explain where rolling-horizon MILP models remain valuable, and where their assumptions and computational structure become restrictive.

The final methodological contribution is the formulation of offshore wind O&M planning as a learning problem. The thesis develops a POMDP-based representation in which the true health state of each component is partially observable and the agent acts on probabilistic RUL-derived observations rather than perfect degradation information. This formulation is then implemented in a PPO-based DRL framework with a multi-level action space. By allowing the agent to choose between no action, minor repair, major repair, and replacement, the framework moves beyond binary maintenance abstractions and better reflects the range of decisions available in offshore maintenance practice. These developments are primarily associated with Borsotti et al. (2026a) and with the DRL framework presented in Borsotti et al. (2025).

Taken together, these contributions provide a coherent comparison between optimization and learning-based approaches for offshore wind O&M planning. The thesis does not present one paradigm as universally superior to the other; instead, it clarifies their respective strengths, limitations, and practical implications. optimization provides transparency, structure, and explicit constraint handling, while learning-based methods offer a way to derive adaptive policies under uncertainty and partial observability. The main contribution of the thesis is therefore a decision-support perspective in which these paradigms are understood through their suitability for different modeling requirements, planning contexts, and stakeholder needs.

7.4 Practical implications

The results of this thesis have practical implications for offshore wind operators, O&M planners, technology developers, and researchers developing decision-support tools. The central implication is that offshore wind maintenance planning should not be treated as a purely technical prediction problem or as a purely logistical scheduling problem. Its value emerges from the interaction between component health information, access constraints, vessel requirements, production losses, and the available set of maintenance actions.

A direct implication concerns the role of prognostics in O&M planning. RUL predictions should be evaluated not only by their statistical accuracy, but by their decision value. A prognostic estimate is useful when it improves maintenance timing, reduces avoidable downtime, prevents unnecessary interventions, or enables better use of vessels and weather windows. This means that condition monitoring and prognostic models should be designed with downstream planning decisions in mind. Prognostic information that is accurate but difficult to translate into maintenance actions has limited operational value, whereas uncertain but decision-relevant information can still support better planning if its uncertainty is represented explicitly.

The thesis also reinforces the importance of marine logistics and weather-driven accessibility as defining features of offshore wind O&M. The economic value of a maintenance action depends not only on the condition of the component, but also on whether the turbine can be accessed, which vessel is required, how long the task takes, whether other tasks can be grouped, and how much production is lost while waiting. Decision-support systems that simplify logistics too strongly risk misrepresenting the value of preventive, corrective, and opportunistic maintenance. For practical

planning, accessibility and vessel constraints should therefore be embedded in the decision model rather than added as secondary feasibility checks.

Another practical implication concerns the representation of maintenance actions. The distinction between minor repair, major repair, and replacement is not merely a modeling detail; it changes the structure of the decision problem. Binary action spaces can force unrealistic choices between inaction and full replacement, while intermediate actions allow planners or agents to trade off cost, restorative effect, downtime, and logistics burden more realistically. This is particularly relevant for offshore wind, where mobilisation costs are high and imperfect repairs may provide operationally valuable extensions of component life.

The comparison between optimization- and learning-based methods also has implications for how decision-support tools should be positioned in practice. MILP-type models remain highly relevant when planners require auditable schedules, explicit constraint handling, and transparent explanations of why an intervention is selected. However, rolling-horizon implementations should be designed carefully when grouping opportunities, long-term replacement timing, or horizon-boundary effects are important. Diagnostic indicators such as missed grouping opportunities and end-of-life regret can help planners assess when a schedule is locally optimal but strategically suboptimal.

Learning-based methods, by contrast, offer a promising route for exploring adaptive maintenance policies under uncertainty, but their practical use requires careful validation. Before DRL-based policies can support operational decisions, the training environment must be credible, the learned behaviour must be interpretable enough for human decision support, and safety and feasibility constraints must be enforced. In the near term, DRL may therefore be most valuable as a policy-learning, scenario-exploration, and stress-testing tool rather than as an autonomous maintenance planner. Its practical contribution lies in revealing adaptive decision patterns that can inform planners, benchmark heuristic rules, or support the design of future decision-support systems.

A further practical implication concerns the intended users of the different model types. The models developed in this thesis are not intended for a single homogeneous decision maker. Offshore wind O&M involves several stakeholders whose objectives overlap but are not identical. An OWF owner or asset manager is primarily concerned with lifetime profitability, availability, risk exposure, and the long-term value of the asset. For this stakeholder, the relevant objective is balance maintenance cost, downtime, production losses, contractual obligations, and residual asset value. A maintenance or O&M service provider, by contrast, may be more directly concerned with reducing intervention cost, vessel days, technician workload, and campaign execution risk. Vessel operators may benefit from models that improve fleet utilisation and campaign planning, while turbine OEMs and monitoring-system providers may use such models to quantify the operational value of improved diagnostics and prognostics.

The same modelling framework can also be adapted to different stakeholder objectives by modifying the objective or reward function. For an OWF owner, the objective could be formulated around profit or net present value, including lost production, electricity price variation, availability penalties, and long-term asset value. For an O&M provider, the objective could emphasise maintenance execution cost, vessel mobilisation, technician utilisation, and contractual service-level compliance. For a vessel provider, the model could focus on fleet utilisation, chartering decisions, and campaign regularity. For a monitoring-system provider, the model can be used to estimate the decision value of improved RUL predictions by comparing maintenance outcomes under different

levels of prognostic uncertainty. Thus, the contribution of the thesis is a modelling basis through which different stakeholders can evaluate how condition information, logistics constraints, and maintenance decisions interact under their own operational and economic priorities.

7.5 Limitations and future research

The thesis has several limitations that should be considered when interpreting the results. These limitations are primarily related to the level of abstraction required to build a controlled benchmark, the simplified representation of offshore execution dynamics, and the current scale of the learning-based case study.

The benchmark environment is synthetic. Although it is designed to be realistic, statistically grounded, and consistent across decision-support paradigms, it does not reproduce the full operational complexity of a specific offshore wind farm. The RUL generator, maintenance effects, metocean accessibility model, and cost parameters are therefore intended primarily for controlled comparison and methodological evaluation. This synthetic design is also what allows the thesis to compare optimization- and learning-based approaches under common assumptions, but it means that the numerical results should not be interpreted as direct estimates for a specific wind farm without calibration.

At the same time, the framework was deliberately designed to be data-compatible rather than purely synthetic. Its modules can be replaced or calibrated using site-specific information when such data are available. Real turbine coordinates can replace the procedural grid layout, measured metocean time series can replace the synthetic accessibility generator, condition-monitoring or SCADA-based prognostic outputs can replace the synthetic RUL predictions, and project-specific cost, vessel, and maintenance-duration data can replace the benchmark parameters. In this sense, the synthetic benchmark provides a controlled starting point, while the underlying modeling structure remains compatible with real-life deployment after appropriate calibration and validation.

A second limitation concerns the representation of weather and accessibility. The metocean model is decision-centric: it produces accessibility signals that indicate whether a given class of intervention can be performed, rather than modeling full continuous wind and wave trajectories and the detailed execution of offshore operations. This abstraction is appropriate for comparing decision models, but it limits the analysis of more detailed marine-operation dynamics such as interrupted tasks, sea-state-dependent transfer risk, visibility constraints, task continuation across multiple weather windows, and correlated multi-day offshore campaigns.

The optimization model also simplifies some aspects of offshore logistics in order to remain tractable. Routing, vessel-task coupling, technician assignment, spare-part logistics, and campaign execution could all be represented in greater operational detail. However, adding such detail would substantially increase problem size and computational burden, particularly in a rolling-horizon setting with uncertainty and repeated re-optimization. The MILP results should therefore be interpreted as evidence of the value and limitations of prognostic-driven tactical planning under a controlled logistics abstraction, rather than as a complete industrial scheduling tool.

The DRL case study is likewise limited in scale. The PPO framework demonstrates the feasibility and value of a POMDP-based learning formulation with probabilistic RUL observations and multi-level maintenance actions, but it does not yet establish scalability to utility-scale offshore wind

farms with many turbines, detailed vessel routing, multiple crews, spare-part dependencies, and richer operational constraints. Scaling the learning framework would likely require additional modeling choices, particularly multi-agent decomposition, hierarchical policies, graph-based state representations, action masking, and stronger integration of domain knowledge.

Finally, both modeling paradigms depend on the quality of the information available to them. In practice, RUL predictions may be biased, sparse, component-specific, or unavailable for some subsystems. Although the thesis represents uncertainty explicitly, real-world deployment would require calibration against operational condition-monitoring data, SCADA signals, inspection records, maintenance-management systems, and site-specific logistics data. The results should therefore be read as a methodological demonstration of how prognostic information can be embedded into offshore wind O&M decision support, rather than as a claim of immediate deployability without further validation.

Building on these limitations, the findings of this thesis suggest several directions for future research. These can be divided into short-term extensions that directly build on the contributions developed in this thesis, and longer-term research directions that broaden the scope of offshore wind O&M decision support toward wider operational, lifecycle, and sustainability challenges.

7.5.1 Short-term future research directions

The following research directions represent the most immediate extensions of the work developed in this thesis. They build directly on the proposed benchmark environment, optimization framework, and DRL formulation, and focus on improving realism, scalability, robustness, and practical applicability within the same offshore wind O&M decision-support context.

- Integration with real operational and prognostic data

The most direct next step is to replace the synthetic RUL and accessibility layers with real operational data where available. This includes SCADA data, condition-monitoring signals, inspection reports, CMMS records, vessel logs, weather forecasts, and actual maintenance histories. Such integration would allow the benchmark to be calibrated to specific wind farms and would enable validation of both optimization- and learning-based policies against observed operational behaviour.

- Hybrid optimization–learning frameworks

A promising direction is the development of hybrid approaches that combine the strengths of MILP and DRL. Optimization could be used to enforce hard feasibility constraints, generate expert demonstrations, validate learned actions, or provide short-term schedule refinement. DRL could be used to learn long-term policies, approximate value functions, identify promising maintenance opportunities, or guide optimization through learned heuristics. Such hybrid methods could reduce the weaknesses of each paradigm: the horizon myopia of rolling optimization and the limited interpretability of pure DRL.

- Scalable learning architectures and multi-agent decomposition

Future DRL work should address scalability to utility-scale wind farms. The PPO framework developed in this thesis demonstrates the feasibility of a POMDP-based learning formulation

with probabilistic RUL observations and multi-level maintenance actions, but it remains centralized and limited in scale. Scaling this approach to larger wind farms would require architectures that reduce the dimensionality of the state and action spaces while preserving coordination across turbines, components, vessels, and maintenance resources.

One promising direction is multi-agent decomposition, where the maintenance problem is distributed across multiple interacting agents. These agents could represent turbines, components, vessels, crews, or maintenance decision units, depending on the level of operational detail required. Such a formulation could allow agents to learn local maintenance policies while coordinating through shared constraints such as vessel availability, technician capacity, weather windows, and farm-level cost or availability objectives. This would extend the structured action-space formulation developed in Chapter 6 toward larger and more operationally realistic offshore wind farms.

Several architectural choices could support this direction, including graph neural networks to represent turbine layouts and component dependencies, centralized-training/decentralized-execution schemes to coordinate local policies, hierarchical policies to separate strategic, tactical, and operational decisions, and action masking to enforce feasibility. Transfer learning and domain randomisation could also help train policies that generalise across wind farm sizes, layouts, weather regimes, and component reliability assumptions.

- Richer logistics and execution modeling

Further work should include more detailed representations of offshore logistics. This includes routing between turbines, vessel chartering and mobilisation, technician shifts, spare-part availability, interrupted tasks, multi-day campaigns, simultaneous vessel operations, and port constraints. Including these details would make both MILP and DRL models more operationally realistic, but would also require careful decomposition or abstraction to remain computationally manageable.

- Improved uncertainty modeling and robustness

The thesis models prognostic uncertainty through synthetic RUL distributions and time-varying confidence. Future work should examine biased predictions, missing data, delayed inspections, false alarms, imperfect fault detection, and correlated component degradation. Robust optimization, distributionally robust learning, Bayesian RL, and risk-sensitive reward formulations may help produce maintenance policies that perform reliably under imperfect information.

- Belief-state reasoning and active information acquisition

The present PPO implementation responds to the latest health observation without explicitly retaining the history of previous observations and maintenance actions. Future work should investigate memory- and belief-based policy representations that exploit temporal information to infer the latent degradation state more accurately. Possible approaches include recurrent policies, attention-based memory, Bayesian filtering, and learned latent-state estimators. A complementary extension is to treat inspection as an explicit decision rather than assuming that updated health information is always available. This would allow the agent to optimise the value of information by balancing the expected benefit of reducing condition uncertainty against inspection cost, vessel and technician requirements, and weather-driven accessibility.

Such extensions would more fully exploit the POMDP formulation and connect maintenance planning directly with inspection and condition-monitoring strategy.

- Interpretability, safety, and decision support

For learning-based methods to be accepted in offshore wind O&M, interpretability and safety must be improved. Future research should develop methods to explain why a learned policy recommends a specific intervention, quantify confidence in the recommendation, detect out-of-distribution states, and ensure that unsafe or infeasible actions are never selected. Human-in-the-loop decision support may be particularly relevant, where the model proposes actions but planners retain final authority.

7.5.2 Longer-term future research directions

Beyond the direct extensions of the proposed models, offshore wind O&M decision support also opens broader research opportunities related to lifecycle management, sustainability, and cross-horizon coordination. The following directions represent wider research themes that could form the basis for future interdisciplinary projects and next-generation decision-support frameworks.

- Cross-horizon and lifecycle integration

Future models should further integrate strategic, tactical, and operational decisions. Maintenance timing is linked to fleet contracts, spare-part policies, inspection strategies, end-of-life planning, and repowering or decommissioning decisions. A full lifecycle framework would allow short-term decisions to be evaluated in terms of long-term economic and operational consequences, reducing the risk of locally optimal but globally suboptimal maintenance behaviour.

- Environmental and societal objectives

Finally, offshore wind O&M decision support should expand beyond cost and downtime. Vessel emissions, ecological disturbance, biodiversity-sensitive areas, noise, circularity, waste streams, and stakeholder acceptance can all be affected by maintenance logistics and scheduling. Future decision-support frameworks should therefore incorporate environmental and social objectives alongside economic performance, enabling O&M strategies that support both reliability and sustainability.

7.6 Final remarks

Offshore wind is becoming a central pillar of the energy transition, but its long-term value depends on the ability to operate and maintain increasingly large and remote assets reliably and efficiently. Maintenance planning is therefore not a secondary operational task; it is a core component of offshore wind performance. Decisions about when to intervene, which action to perform, and which resources to mobilise determine not only direct O&M cost, but also downtime, energy production, vessel utilisation, and the practical value of condition-monitoring technologies.

The work presented in this thesis shows that intelligent maintenance decision-making requires more than better predictions. It requires models that connect uncertain health information to feasible

offshore actions under weather and logistics constraints. Optimization-based approaches provide transparency and explicit constraint handling, while learning-based approaches offer adaptivity and the potential to learn long-term policies under uncertainty.

The final message of this thesis is therefore that the future of offshore wind O&M decision support should be framed as the design of integrated decision systems in which physical understanding, operational constraints, and prognostic information, are considered jointly. Such systems can help transform condition monitoring from a diagnostic capability into an operational decision-making tool, supporting offshore wind farms that are not only larger and more intelligent, but also more reliable, cost-effective, and sustainable.

Bibliography

- Aafif, Y., A. Chelbi, L. Mifdal, S. Dellagi, I. Majdouline (2022) Optimal preventive maintenance strategies for a wind turbine gearbox, *Energy Reports*, 8, pp. 803–814, supplement 9.
- Abbas, A. (2024) *A Hierarchical Framework for Interpretable, Safe, and Specialised Deep Reinforcement Learning*, Doctoral thesis, Technological University Dublin.
- Abbas, A. N., G. C. Chasparis, J. D. Kelleher (2024) Hierarchical framework for interpretable and specialized deep reinforcement learning-based predictive maintenance, *Data & Knowledge Engineering*, 149, p. 102240.
- Abkar, M., N. Zehtabiyani-Rezaie, A. Iosifidis (2023) Reinforcement learning for wind-farm flow control: Current state and future actions, *Theoretical and Applied Mechanics Letters*, 13(6), p. 100475.
- Adadi, A., M. Berrada (2018) Peeking inside the black-box: A survey on explainable artificial intelligence (XAI), *IEEE Access*, 6, pp. 52138–52160.
- Ade Irawan, C., S. Starita, H. K. Chan, M. Eskandarpour, M. Reihaneh (2023) Routing in offshore wind farms: A multi-period location and maintenance problem with joint use of a service operation vessel and a safe transfer boat, *European Journal of Operational Research*, 307(1), pp. 328–350.
- Ambühl, S., J. D. Sørensen (2017) Sensitivity of risk-based maintenance planning of offshore wind turbine farms, *Energies*, 10(4), p. 505.
- Andriotis, C. P., K. G. Papakonstantinou (2021) Deep reinforcement learning driven inspection and maintenance planning under incomplete information and constraints, *Reliability Engineering & System Safety*, 212, p. 107551.
- Aria, M., C. Cuccurullo (2017) *Bibliometrix: An R-tool for comprehensive science mapping analysis*, r package version 3.1.4.
- Asgarpour, M., J. D. Sørensen (2018) Bayesian based diagnostic model for condition based maintenance of offshore wind farms, *Energies*, 11(2), p. 300.
- Bhuria, R., K. S. Gill, K. Rajput, V. Singh (2024) Enhancing wind turbine lifespan using novel vibration sensor data techniques for prognosis, in: *Proceedings of the 2024 IEEE 3rd World Conference on Applied Intelligence and Computing (AIC 2024)*, pp. 1217–1222, conference held in Gwalior, India, 27–28 July 2024.

- Borsotti, M., X. Jiang, R. R. Negenborn (2026a) Review of deep reinforcement learning for offshore wind farm maintenance planning, *Wind Energy Science*, 11(4), pp. 1185–1204.
- Borsotti, M., X. Jiang, R. R. Negenborn (2026b) Stochastic optimization of predictive maintenance scheduling for offshore wind farms, *Ocean Engineering*, 357, p. 125424.
- Borsotti, M., R. R. Negenborn, X. Jiang (2024) Model predictive control framework for optimizing offshore wind O&M, in: *Advances in Maritime Technology and Engineering*, CRC Press, pp. 533–546, presented at the 7th International Conference on Maritime Technology and Engineering (MARTECH 2024), Lisbon, Portugal, 14–16 May 2024.
- Borsotti, M., R. R. Negenborn, X. Jiang (2026c) A review of multi-horizon decision-making for operation and maintenance of fixed-bottom offshore wind farms, *Renewable and Sustainable Energy Reviews*, 226, p. 116450.
- Borsotti, M., R. L. Patrão, X. Jiang, A. Napoleone, R. R. Negenborn (2025) A deep reinforcement learning framework for O&M planning in offshore wind, URL <https://iccl2025.org/>, accepted for publication in the proceedings of the joint 16th International Conference on Computational Logistics (ICCL 2025) and 1st EURO Mini Conference on Maritime Optimization and Logistics (EUROMar 2025), Springer Lecture Notes in Computer Science.
- Bousdekis, A., B. Magoutas, D. Apostolou, G. Mentzas (2018) Review, analysis and synthesis of prognostic-based decision support methods for condition-based maintenance, *Journal of Intelligent Manufacturing*, 29(6), pp. 1303–1316.
- Bui, V.-H., S. Das, A. Hussain, G. V. Hollweg, W. Su (2024) A critical review of safe reinforcement learning techniques in smart grid applications, URL <https://arxiv.org/abs/2409.16256>.
- BVG Associates (2014) Uk offshore wind supply chain: Capabilities and opportunities, Tech. Rep. BIS/14/578, UK Department for Business, Innovation & Skills and Department of Energy & Climate Change.
- Camci, F., K. Medjaher, V. Atamuradov, A. Berdinyazov (2019) Integrated maintenance and mission planning using remaining useful life information, *Engineering Optimization*, 51(10), pp. 1794–1809.
- Canet, H., A. Guílloré, C. L. Bottasso (2023) The eco-conscious wind turbine: Design beyond purely economic metrics, *Wind Energy Science*, 8(5), pp. 1029–1047.
- Carrillo, C., A. F. Obando Montaña, J. Cidrás, E. Díaz-Dorado (2013) Review of power curve modelling for wind turbines, *Renewable and Sustainable Energy Reviews*, 21, pp. 572–581.
- Carroll, J., A. McDonald, I. Dinwoodie, D. McMillan, M. Revie, I. Lazakis (2017) Availability, operation and maintenance costs of offshore wind turbines with different drive train configurations, *Wind Energy*, 20(2), pp. 361–378.
- Carroll, J., A. McDonald, D. McMillan (2016) Failure rate, repair time and unscheduled o&m cost analysis of offshore wind turbines, *Wind Energy*, 19(6), pp. 1107–1119.
- Castro-Santos, L., V. Diaz-Casas (2014) Life-cycle cost analysis of floating offshore wind farms, *Renewable Energy*, 66, pp. 41–48.

- Centeno-Telleria, M., J. I. Aizpurua, M. Penalba (2023) Computationally efficient analytical O&M model for strategic decision-making in offshore renewable energy systems, *Energy*, 285, p. 129374.
- Centeno-Telleria, M., B. Jenkins, J. Carroll, J. I. Aizpurua, M. Penalba (2024a) Sensitivity analysis of major replacement rates and costs for floating wind farms, in: *Proceedings of the 6th International Conference on Renewable Energies Offshore (RENEW 2024)*, pp. 989–1004, conference held in Lisbon, Portugal, 19–21 November 2024.
- Centeno-Telleria, M., H. Yue, J. Carroll, M. Penalba, J. I. Aizpurua (2024b) Impact of operations and maintenance on the energy production of floating offshore wind farms across the north sea and the iberian peninsula, *Renewable Energy*, 224, p. 120217.
- Centeno-Telleria, M., H. Yue, J. Carroll, M. Penalba, J. I. Aizpurua (2025) Assessing heavy maintenance alternatives for floating offshore wind farms: Towing vs. onsite replacement strategies, *Applied Energy*, 377, p. 124437.
- Chatterjee, J., N. Dethlefs (2021) Scientometric review of artificial intelligence for operations & maintenance of wind turbines: The past, present and future, *Renewable and Sustainable Energy Reviews*, 144, p. 111051.
- Chen, C., C. Wang, N. Lu, B. Jiang, Y. Xing (2021) A data-driven predictive maintenance strategy based on accurate failure prognostics, *Eksploatacja i Niezawodność – Maintenance and Reliability*, 23(2), pp. 387–394.
- Chen, M., Y. Kang, K. Li, P. Li, Y.-B. Zhao (2025) Deep reinforcement learning for maintenance optimization of multi-component production systems considering quality and production plan, *Quality Engineering*, 37(2), pp. 1–12.
- Cheng, J., Y. Liu, W. Li, T. Li (2023) Deep reinforcement learning for cost-optimal condition-based maintenance policy of offshore wind turbine components, *Ocean Engineering*, 283, p. 115062.
- Civera, M., C. Surace (2022) Non-destructive techniques for the condition and structural health monitoring of wind turbines: A literature review of the last 20 years, *Sensors*, 22(4), p. 1627.
- Colone, L., N. Dimitrov, D. Straub (2019) Predictive repair scheduling of wind turbine drive-train components based on machine learning, *Wind Energy*, 22(9), pp. 1230–1242.
- Commission, I. E. (2018) Failure modes and effects analysis (FMEA and FMECA), URL <https://webstore.iec.ch/en/publication/26359>, edition 2.0.
- Commission, I. E. (2022) Wind energy generation systems – part 12-1: Power performance measurements of electricity producing wind turbines, URL <https://webstore.iec.ch/en/publication/68499>.
- da Costa, P., P. Verleijdsdonk, S. Voorberg, A. Akçay, S. Kapodistria, W. van Jaarsveld, Y. Zhang (2023) Policies for the dynamic traveling maintainer problem with alerts, *European Journal of Operational Research*, 305(3), pp. 1141–1152.
- Dai, L., M. Stålhane, I. B. Utne (2015) Routing and scheduling of maintenance fleet for offshore wind farms, *Wind Engineering*, 39(1), pp. 15–30.

- Dalgic, Y., I. Dinwoodie, I. Lazakis, D. McMillan, M. Revie, J. Majumder (2015a) The influence of multiple working shifts for offshore wind farm O&M activities – StrathOW-OM tool, in: *Proceedings of the Royal Institution of Naval Architects International Conference on Design and Operation of Offshore Wind Farm Support Vessels*, London, UK, pp. 1–12, conference held in London, United Kingdom, 28–29 January 2015.
- Dalgic, Y., I. A. Dinwoodie, I. Lazakis, D. McMillan, M. Revie (2014) Optimum CTV fleet selection for offshore wind farm O&M activities, in: Nowakowski, T., M. Młyńczak, A. Jodejko-Pietruczuk, S. Werbińska-Wojciechowska, eds., *Safety and Reliability: Methodology and Applications*, CRC Press, pp. 1177–1185, proceedings of the 24th European Safety and Reliability Conference (ESREL 2014), Wrocław, Poland, 14–18 September 2014.
- Dalgic, Y., I. Lazakis, I. Dinwoodie, D. McMillan, M. Revie (2015b) Advanced logistics planning for offshore wind farm operation and maintenance activities, *Ocean Engineering*, 101, pp. 211–226.
- Dalgic, Y., I. Lazakis, I. Dinwoodie, D. McMillan, M. Revie (2015c) Advanced logistics planning for offshore wind farm operation and maintenance activities, *Ocean Engineering*, 101, pp. 211–226.
- Dalgic, Y., I. Lazakis, O. Turan (2015d) Investigation of optimum crew transfer vessel fleet for offshore wind farm maintenance operations, *Wind Engineering*, 39(1), pp. 31–52.
- Dalgic, Y., I. Lazakis, O. Turan, S. Judah (2015e) Investigation of optimum jack-up vessel chartering strategy for offshore wind farm o&m activities, *Ocean Engineering*, 95, pp. 106–115.
- Dao, C. D., B. Kazemtabrizi, C. J. Crabtree (2020) Offshore wind turbine reliability and operational simulation under uncertainties, *Wind Energy*, 23(10), pp. 1919–1938.
- Darling, S. B., F. You, T. Veselka, A. Velosa (2011) Assumptions and the levelized cost of energy for photovoltaics, *Energy & Environmental Science*, 4(9), pp. 3133–3139.
- Dawid, R., D. McMillan, M. Revie (2018) Decision support tool for offshore wind farm vessel routing under uncertainty, *Energies*, 11(9), p. 2190.
- de Pater, I., M. Mitici (2021) Predictive maintenance for multi-component systems of repairables with remaining-useful-life prognostics and a limited stock of spare components, *Reliability Engineering & System Safety*, 214, p. 107761.
- de Pater, I., A. Reijns, M. Mitici (2022) Alarm-based predictive maintenance scheduling for aircraft engines with imperfect remaining useful life prognostics, *Reliability Engineering & System Safety*, 221, p. 108341.
- Dewan, A., T. Stehly (2017) Mapping operation & maintenance strategy for u.s. offshore wind farms, Tech. Rep. NREL/PR-6A20-68494, National Renewable Energy Laboratory (NREL), Golden, CO, USA, presented at the 2nd U.S. Offshore Wind Conference and Expo, Long Island, New York, USA, 9 May 2017.
- Dinwoodie, I., O.-E. V. Endrerud, M. Hofmann, R. Martin, I. B. Sperstad (2015) Reference cases for verification of operation and maintenance simulation models for offshore wind farms, *Wind Engineering*, 39(1), pp. 1–14.

- Dinwoodie, I., D. McMillan, M. Revie, I. Lazakis, Y. Dalgic (2013) Development of a combined operational and strategic decision support model for offshore wind, *Energy Procedia*, 35, pp. 157–166.
- Dinwoodie, I. A., D. McMillan (2014) Operational strategies for offshore wind turbines to mitigate failure rate uncertainty on operational costs and revenue, *IET Renewable Power Generation*, 8(4), pp. 359–366.
- Dinwoodie, I. A., D. McMillan, F. Quail (2012) Analysis of offshore wind turbine operation & maintenance using a novel time domain meteo-ocean modeling approach, in: *Proceedings of ASME Turbo Expo 2012: Turbine Technical Conference and Exposition*, pp. 847–857, paper GT2012-68985; conference held in Copenhagen, Denmark, 11–15 June 2012.
- Do, P., V.-T. Nguyen, A. Voisin, B. Iung, W. Alves Ferreira Neto (2024) Multi-agent deep reinforcement learning-based maintenance optimization for multi-dependent component systems, *Expert Systems with Applications*, 245, p. 123144.
- Dong, W., T. Zhao, Y. Wu (2021) Deep reinforcement learning based preventive maintenance for wind turbines, in: *Proceedings of the 2021 IEEE 5th Conference on Energy Internet and Energy System Integration (EI2 2021)*, pp. 2860–2865, conference held in Taiyuan, China, 22–25 October 2021.
- Donnelly, O., F. Anderson, J. Carroll (2024) Operation and maintenance cost comparison between 15 mw direct-drive and medium-speed offshore wind turbines, *Wind Energy Science*, 9(6), pp. 1345–1362.
- Douard, F., C. Domecq, W. Lair (2012) A probabilistic approach to introduce risk measurement indicators to an offshore wind project evaluation: Improvement to an existing tool Ecume, *Energy Procedia*, 24, pp. 255–262.
- Dulac-Arnold, G., N. Levine, D. J. Mankowitz, J. Li, C. Paduraru, S. Gowal, T. Hester (2021) Challenges of real-world reinforcement learning: Definitions, benchmarks and analysis, *Machine Learning*, 110(9), pp. 2419–2468.
- Elusakin, T., M. Shafiee, T. Adedipe, F. Dinmohammadi (2021) A stochastic petri net model for O&M planning of floating offshore wind turbines, *Energies*, 14(4), p. 1134.
- Endrerud, O.-E. V., K. R. Austreng, N. Keseric, J. P. Liyanage (2015) New vessel concepts for operations and maintenance of offshore wind farms, in: *Proceedings of the 25th International Ocean and Polar Engineering Conference (ISOPE 2015)*, International Society of Offshore and Polar Engineers, pp. 637–643, conference held in Kona, Hawaii, USA, 21–26 June 2015.
- Endrerud, O.-E. V., J. P. Liyanage, N. Keseric (2014) Marine logistics decision support for operation and maintenance of offshore wind parks with a multi method simulation model, in: *Proceedings of the 2014 Winter Simulation Conference (WSC 2014)*, pp. 1712–1722, conference held in Savannah, Georgia, USA, 7–10 December 2014.
- European Commission (2020) An EU strategy to harness the potential of offshore renewable energy for a climate neutral future, Tech. Rep. COM(2020) 741 final, European Commission, accessed: 2026-01-09.

- Fan, D., Y. Ren, Q. Feng, B. Zhu, Y. Liu, Z. Wang (2019) A hybrid heuristic optimization of maintenance routing and scheduling for offshore wind farms, *Journal of Loss Prevention in the Process Industries*, 62, p. 103949.
- Ferdinand, R., A. Monti, K. Labusch (2018) Determining spare part inventory for offshore wind farm substations based on FMEA analysis, in: *Proceedings of the 2018 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe 2018)*, pp. 1–6, conference held in Sarajevo, Bosnia and Herzegovina, 21–25 October 2018.
- Fernandez-Navamuel, A., N. Gorostidi, D. Pardo, V. Nava, E. Chatzi (2025) Gaussian mixture autoencoder for uncertainty-aware damage identification in a floating offshore wind turbine, *Wind Energy Science*, 10(5), pp. 857–885.
- Fox, H., A. Pillai, D. Friedrich, M. Collu, T. Dawood, L. Johanning (2022) A review of predictive and prescriptive offshore wind farm operation and maintenance, *Energies*, 15(2), p. 504.
- Frederiksen, T. J., A. Karyotakis, J. Santos, S. Faulstich (2024) Cost-effectiveness of predictive maintenance for offshore wind farms: A case study, *Energies*, 17(13), p. 3147.
- Froger, A., M. Gendreau, J. E. Mendoza, É. Pinson, L.-M. Rousseau (2018) Solving a wind turbine maintenance scheduling problem, *Journal of Scheduling*, 21(1), pp. 53–76.
- Garcia-Teruel, A., G. Rinaldi, P. R. Thies, L. Johanning, H. Jeffrey (2022) Life cycle assessment of floating offshore wind farms: An evaluation of operation and maintenance, *Applied Energy*, 307, p. 118067.
- Ghane, M., A. R. Nejad, M. Blanke, Z. Gao, T. Moan (2018) Condition monitoring of spar-type floating wind turbine drivetrain using statistical fault diagnosis, *Wind Energy*, 21(7), pp. 575–589.
- Gintautas, T., J. D. Sørensen (2017) Improved methodology of weather window prediction for offshore operations based on probabilities of operation failure, *Journal of Marine Science and Engineering*, 5(2), p. 20.
- Global Wind Energy Council (2025) Global offshore wind report 2025, Tech. rep., Global Wind Energy Council (GWEC), accessed: 2026-01-09.
- Gutierrez-Alcoba, A., E. M. T. Hendrix, G. Ortega, E. E. Halvorsen-Weare, D. Haugland (2019) On offshore wind farm maintenance scheduling for decision support on vessel fleet composition, *European Journal of Operational Research*, 279(1), pp. 124–131.
- Haarnoja, T., A. Zhou, P. Abbeel, S. Levine (2018) Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, in: *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, vol. 80 of *Proceedings of Machine Learning Research*, PMLR, pp. 1861–1870, conference held in Stockholm, Sweden, 10–15 July 2018.
- Hall, R., E. Topham, E. João (2022) Environmental impact assessment for the decommissioning of offshore wind farms, *Renewable and Sustainable Energy Reviews*, 165, p. 112580.
- Halvorsen-Weare, E. E., C. Gundegjerde, I. B. Halvorsen, L. M. Hvattum, L. M. Nonås (2013) Vessel fleet analysis for maintenance operations at offshore wind farms, *Energy Procedia*, 35, pp. 167–176, deepWind 2013 – Selected papers from the 10th Deep Sea Offshore Wind R&D Conference, Trondheim, Norway, 24–25 January 2013.

- Hausknecht, M., P. Stone (2015) Deep recurrent Q-learning for partially observable MDPs, in: *Proceedings of the 2015 AAAI Fall Symposium Series*, AAAI Press, pp. 29–37, symposium held in Arlington, Virginia, USA, 12–14 November 2015.
- He, J., S. Khebbache, M. F. Anjos, M. Hadji (2023a) Optimization of maintenance for complex manufacturing systems using stochastic remaining useful life prognostics, *Computers & Industrial Engineering*, 182, p. 109348.
- He, R., Z. Tian, Y. Wang, Y. Chen, M. J. Zuo (2025) Predictive maintenance for offshore wind farms with incomplete and biased prognostic information, *Ocean Engineering*, 322, p. 120541.
- He, R., Z. Tian, Y. Wang, M. J. Zuo, Z. Guo (2023b) Condition-based maintenance optimization for multi-component systems considering prognostic information and degraded working efficiency, *Reliability Engineering & System Safety*, 234, p. 109167.
- Hernandez, O. M., M. Shadman, M. M. Amiri, C. Silva, S. F. Estefen, E. La Rovere (2021) Environmental impacts of offshore wind installation, operation and maintenance, and decommissioning activities: A case study of Brazil, *Renewable and Sustainable Energy Reviews*, 144, p. 110994.
- HM Government (2022) British energy security strategy, Tech. rep., HM Government, published by the Department for Business, Energy & Industrial Strategy; accessed: 2026-01-09.
- Hofmann, M., I. B. Sperstad (2013) NOWIcob – a tool for reducing the maintenance costs of offshore wind farms, *Energy Procedia*, 35, pp. 177–186, deepWind 2013 – Selected papers from the 10th Deep Sea Offshore Wind R&D Conference, Trondheim, Norway, 24–25 January 2013.
- Hofmann, M., I. B. Sperstad (2014) Technical documentation of the NOWIcob tool (D5.1–53), Tech. Rep. TR A7374, SINTEF Energy Research, Trondheim, Norway.
- Hörmann, M. (2016) Offshore warranty cover, WindEurope Summit 2016 presentation, URL https://windeurope.org/summit2016/conference/allfiles/532_WindEurope2016presentation.pdf.
- Hu, B. (2019) Offshore wind access 2019, Tech. Rep. TNO 2019 R10633, TNO, Petten, The Netherlands.
- Huang, J., Q. Chang, J. Arinez (2020) Deep reinforcement learning based preventive maintenance policy for serial production lines, *Expert Systems with Applications*, 160, p. 113701.
- Igl, M., L. Zintgraf, T. A. Le, F. Wood, S. Whiteson (2018) Deep variational reinforcement learning for POMDPs, in: *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, vol. 80 of *Proceedings of Machine Learning Research*, PMLR, pp. 2117–2126, conference held in Stockholm, Sweden, 10–15 July 2018.
- International Energy Agency (2024a) Renewables 2024 – global overview, URL <https://www.iea.org/reports/renewables-2024/global-overview>, accessed: 2026-01-09.
- International Energy Agency (2024b) Renewables 2024: Analysis and forecast to 2030, Tech. rep., International Energy Agency, accessed: 2026-01-09.
- International Energy Agency (2024c) Wind: Tracking clean energy progress, URL <https://www.iea.org/energy-system/renewables/wind>, accessed: 2026-01-09.

- International Renewable Energy Agency (2024) Renewable power generation costs in 2023, Tech. rep., International Renewable Energy Agency, Abu Dhabi, accessed: 2026-01-09.
- Ioannou, A., A. Angus, F. Brennan (2018) A lifecycle techno-economic model of offshore wind energy for different entry and exit instances, *Applied Energy*, 221, pp. 406–424.
- Irawan, C. A., M. Eskandarpour, D. Ouelhadj, D. Jones (2021) Simulation-based optimisation for stochastic maintenance routing in an offshore wind farm, *European Journal of Operational Research*, 289(3), pp. 912–926.
- Irawan, C. A., D. Ouelhadj, D. Jones, M. Stålhane, I. B. Sperstad (2017) Optimisation of maintenance routing and scheduling for offshore wind farms, *European Journal of Operational Research*, 256(1), pp. 76–89.
- Irawan, C. A., S. Salhi, H. K. Chan (2022) A continuous location and maintenance routing problem for offshore wind farms: Mathematical models and hybrid methods, *Computers & Operations Research*, 144, p. 105825.
- Jain, A. K., M. Dhada, M. Perez Hernandez, M. Herrera, A. K. Parlikad (2021) A comprehensive framework from real-time prognostics to maintenance decisions, *IET Collaborative Intelligent Manufacturing*, 3(2), pp. 175–183.
- Jenkins, B., I. Belton, J. Carroll, D. McMillan (2022) Estimating the major replacement rates in next-generation offshore wind turbines using structured expert elicitation, in: *Proceedings of the 19th Deep Sea Offshore Wind R&D Conference (EERA DeepWind 2022)*, vol. 2362, IOP Publishing, p. 012020, conference held in Trondheim, Norway and online, 19–21 January 2022.
- Jenkins, B., A. Prothero, M. Collu, J. Carroll, D. McMillan, A. McDonald (2021) Limiting wave conditions for the safe maintenance of floating wind turbines, in: *Proceedings of the 18th Deep Sea Offshore Wind R&D Conference (EERA DeepWind 2021)*, vol. 2018 of *Journal of Physics: Conference Series*, p. 012023, conference held online, hosted from Trondheim, Norway, 13–15 January 2021.
- Jiang, G.-J., J.-X. Zhou, H.-H. Sun, P. Huang (2025) Reliability prediction method for offshore wind turbines based on time series and data migration models, *Quality and Reliability Engineering International*, 41(2), pp. 102–115.
- Judge, F., F. D. McAuliffe, I. B. Sperstad, R. Chester, B. Flannery, K. Lynch, J. Murphy (2019) A lifecycle financial analysis model for offshore wind farms, *Renewable and Sustainable Energy Reviews*, 103, pp. 370–383.
- Kaelbling, L. P., M. L. Littman, A. R. Cassandra (1998) Planning and acting in partially observable stochastic domains, *Artificial Intelligence*, 101(1–2), pp. 99–134.
- Kang, J., J. Sobral, C. G. Soares (2019) Review of condition-based maintenance strategies for offshore wind energy, *Journal of Marine Science and Application*, 18(1), pp. 1–16.
- Katsouris, G., L. B. Savenije (2016) Offshore wind access 2017, Tech. Rep. ECN-E-16-013, ECN part of TNO, Petten, The Netherlands, accessed: 2026-01-09.

- Kazemian, I., M. Yildirim, P. Ramanan (2024) Attention is all you need to optimize wind farm operations and maintenance, URL <https://arxiv.org/abs/2410.24052>.
- Kerckamp, D., Z. A. Bukhsh, Y. Zhang, N. Jansen (2022) Grouping of maintenance actions with deep reinforcement learning and graph convolutional networks, in: *Proceedings of the 14th International Conference on Agents and Artificial Intelligence (ICAART 2022)*, vol. 2, SciTePress, pp. 574–585, conference held online, 3–5 February 2022.
- Khan, M. E., F. Khan (2012) A comparative study of white box, black box and grey box testing techniques, *International Journal of Advanced Computer Science and Applications*, 3(6), pp. 12–15.
- Kong, S., Y. Liu (2025) Model predictive control for risk-based inspection and maintenance planning of offshore wind turbines, *Ships and Offshore Structures*, pp. 1–12.
- Konuk, E.-B., M. Centeno-Telleria, A. Zarketa-Astigarraga, J.-I. Aizpurua, G. Giorgi, G. Bracco, M. Penalba (2023) On the definition of a comprehensive technology-informed accessibility metric for offshore renewable energy site selection, *Journal of Marine Science and Engineering*, 11(9), p. 1702.
- Kou, L., Y. Li, F. Zhang, X. Gong, Y. Hu, Q. Yuan, W. Ke (2022) Review on monitoring, operation and maintenance of smart offshore wind farms, *Sensors*, 22(8), p. 2822.
- Lambe, A. B., J. R. R. A. Martins (2012) Extensions to the design structure matrix for the description of multidisciplinary design, analysis, and optimization processes, *Structural and Multidisciplinary Optimization*, 46, pp. 273–284.
- Lazakis, I., S. Khan (2021) An optimization framework for daily route planning and scheduling of maintenance vessel activities in offshore wind farms, *Ocean Engineering*, 225, p. 108752.
- Le, B., J. Andrews (2016) Modelling wind turbine degradation and maintenance, *Wind Energy*, 19(4), pp. 571–591.
- Lee, J., M. Mitici (2023a) Deep reinforcement learning for predictive aircraft maintenance using probabilistic remaining-useful-life prognostics, *Reliability Engineering & System Safety*, 230, p. 108908.
- Lee, J., M. Mitici (2023b) Predictive maintenance of multi-component aircraft system using convolutional neural networks and deep reinforcement learning, in: *Proceedings of the 33rd European Safety and Reliability Conference (ESREL 2023)*.
- Lee, N., J. Woo, S. Kim (2025) A deep reinforcement learning ensemble for maintenance scheduling in offshore wind farms, *Applied Energy*, 377(Part A), p. 124431.
- Li, H., W. Peng, C.-G. Huang, C. Guedes Soares (2022a) Failure rate assessment for onshore and floating offshore wind turbines, *Journal of Marine Science and Engineering*, 10(12), p. 1965.
- Li, L., Z. Gao, T. Moan (2015) Joint distribution of environmental condition at five European offshore sites for design of combined wind and wave energy devices, *Journal of Offshore Mechanics and Arctic Engineering*, 137(3), p. 031901.

- Li, M. (2023) *Towards Closed-Loop Maintenance Logistics for Offshore Wind Farms: Approaches for Strategic and Tactical Decision-Making*, Ph.D. thesis, Delft University of Technology.
- Li, M., X. Jiang, J. Carroll, R. R. Negenborn (2022b) A multi-objective maintenance strategy optimization framework for offshore wind farms considering uncertainty, *Applied Energy*, 321, p. 119284.
- Li, M., X. Jiang, J. Carroll, R. R. Negenborn (2022c) A multi-objective maintenance strategy optimization framework for offshore wind farms considering uncertainty, *Applied Energy*, 321, p. 119284.
- Li, M., X. Jiang, J. Carroll, R. R. Negenborn (2024) Operation and maintenance management for offshore wind farms integrating inventory control and health information, *Renewable Energy*, 231, p. 120970.
- Li, M., X. Jiang, R. R. Negenborn (2021) Opportunistic maintenance for offshore wind farms with multiple-component age-based preventive dispatch, *Ocean Engineering*, 231, p. 109062.
- Li, M., M. Wang, J. Kang, L. Sun, P. Jin (2020) An opportunistic maintenance strategy for offshore wind turbine system considering optimal maintenance intervals of subsystems, *Ocean Engineering*, 216, p. 108067.
- Li, Q., T. Lin, Q. Yu, H. Du, J. Li, X. Fu (2023) Review of deep reinforcement learning and its application in modern renewable power system control, *Energies*, 16(10), p. 4143.
- Li, Z., J. Guo, R. Zhou (2016) Maintenance scheduling optimization based on reliability and prognostics information, in: *Proceedings of the 2016 Annual Reliability and Maintainability Symposium (RAMS 2016)*, pp. 1–5, conference held in Tucson, Arizona, USA, 25–28 January 2016.
- Liang, J., H. Miao, K. Li, J. Tan, X. Wang, R. Luo, Y. Jiang (2025) A review of multi-agent reinforcement learning algorithms, *Electronics*, 14(4), p. 820.
- Lillicrap, T. P., J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, D. Wierstra (2016) Continuous control with deep reinforcement learning, in: *Proceedings of the 4th International Conference on Learning Representations (ICLR 2016)*, conference held in San Juan, Puerto Rico, 2–4 May 2016.
- Liu, L., Y. Fu, S. Ma, L. Huang, S. Wei, L. Pang (2019) Optimal scheduling strategy of O&M task for OWF, *IET Renewable Power Generation*, 13(14), pp. 2580–2586.
- Liu, Y., R. Ferrari, P. Wu, X. Jiang, S. Li, J.-W. van Wingerden (2021) Fault diagnosis of the 10 MW floating offshore wind turbine benchmark: A mixed model and signal-based approach, *Renewable Energy*, 164, pp. 391–406.
- Lowe, R., Y. Wu, A. Tamar, J. Harb, P. Abbeel, I. Mordatch (2017) Multi-agent actor-critic for mixed cooperative-competitive environments, in: Guyon, I., U. von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. V. N. Vishwanathan, R. Garnett, eds., *Advances in Neural Information Processing Systems 30: Proceedings of the 31st Annual Conference on Neural Information Processing Systems (NIPS 2017)*, vol. 30, Curran Associates, Inc., conference held in Long Beach, California, USA, 4–9 December 2017.

- Martin, R., I. Lazakis, S. Barbouchi, L. Johanning (2016) Sensitivity analysis of offshore wind farm operation and maintenance cost and availability, *Renewable Energy*, 85, pp. 1226–1236.
- May, A., D. McMillan, S. Thöns (2015) Economic analysis of condition monitoring systems for offshore wind turbine sub-systems, *IET Renewable Power Generation*, 9(8), pp. 900–907.
- Mazidi, P., Y. Tohidi, M. A. Sanz-Bobi (2017) Strategic maintenance scheduling of an offshore wind farm in a deregulated power system, *Energies*, 10(3), p. 313.
- McCoy, A., W. Musial, R. Hammond, D. Mulas Hernando, P. Duffy, P. Beiter, P. Pérez, R. Baranowski, G. Reber, P. Spitsen (2024) Offshore wind market report: 2024 edition, Tech. Rep. NREL/TP-5000-90525, National Renewable Energy Laboratory (NREL), Golden, CO, USA.
- Mitici, M., I. de Pater, A. Barros, Z. Zeng (2023) Dynamic predictive maintenance for multiple components using data-driven probabilistic RUL prognostics: The case of turbofan engines, *Reliability Engineering & System Safety*, 234, p. 109199.
- Mnih, V., A. P. Badia, M. Mirza, A. Graves, T. P. Lillicrap, T. Harley, D. Silver, K. Kavukcuoglu (2016) Asynchronous methods for deep reinforcement learning, in: *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)*, vol. 48 of *Proceedings of Machine Learning Research*, PMLR, pp. 1928–1937, conference held in New York City, New York, USA, 19–24 June 2016.
- Mnih, V., K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, D. Hassabis (2015) Human-level control through deep reinforcement learning, *Nature*, 518(7540), pp. 529–533.
- Munich Re (2023) Offshore wind park insurance, Company webpage, URL <https://www.munichre.com/en/solutions/for-industry-clients/risk-transfer-solutions-for-on-and-off-shore-wind-power.html>.
- Musial, W., P. Spitsen, P. Duffy, P. Beiter, M. Marquis, R. Hammond, M. Shields (2022) Offshore wind market report: 2022 edition, Tech. Rep. NREL/TP-5000-83544, National Renewable Energy Laboratory (NREL), Golden, CO, USA.
- Nachimuthu, S., M. J. Zuo, Y. Ding (2019) A decision-making model for corrective maintenance of offshore wind turbines considering uncertainties, *Energies*, 12(8), p. 1408.
- Narayanan, S. (2023) Reinforcement learning in wind energy: A review, *International Journal of Green Energy*, 20(5), pp. 443–465.
- Neves-Moreira, F., J. Veldman, R. H. Teunter (2021) Service operation vessels for offshore wind farm maintenance: Optimal stock levels, *Renewable and Sustainable Energy Reviews*, 146, p. 111158.
- Nguyen, V.-T., P. Do, A. Voisin, B. Iung (2022) Weighted-QMIX-based optimization for maintenance decision-making of multi-component systems, in: *Proceedings of the 7th European Conference of the Prognostics and Health Management Society (PHME 2022)*, vol. 7, pp. 360–367, conference held in Turin, Italy, 6–8 July 2022.

- Nielsen, J. S., J. D. Sørensen (2014) Methods for risk-based planning of O&M of wind turbines, *Energies*, 7(10), pp. 6645–6664.
- Obdam, T. S., H. Braam, L. W. M. M. Rademakers (2011) User guide and model description of ECN O&M tool version 4, Tech. rep., Energy Research Centre of the Netherlands (ECN), technical report; no persistent public report URL found, metadata corroborated through secondary indexed references.
- Ogunfowora, O., H. Najjaran (2023) Reinforcement and deep reinforcement learning-based solutions for machine maintenance planning, scheduling policies, and optimization, *Journal of Manufacturing Systems*, 70, pp. 244–263.
- Pandit, R., D. Infield (2018) Gaussian process operational curves for wind turbine condition monitoring, *Energies*, 11(7), p. 1631.
- Pandit, R., J. Wang (2024) A comprehensive review on enhancing wind turbine applications with advanced SCADA data analytics and practical insights, *IET Renewable Power Generation*, 18(4), pp. 722–742.
- Papatzimos, A. K., P. R. Thies, T. Dawood (2019) Offshore wind turbine fault alarm prediction, *Wind Energy*, 22(12), pp. 1779–1788.
- Parisotto, E., H. F. Song, J. W. Rae, R. Pascanu, C. Gulcehre, S. M. Jayakumar, M. Jaderberg, R. L. Kaufman, A. Clark, S. Noury, M. M. Botvinick, N. Heess, R. Hadsell (2020) Stabilizing transformers for reinforcement learning, in: *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*, vol. 119 of *Proceedings of Machine Learning Research*, PMLR, pp. 7487–7498, conference held virtually, 13–18 July 2020.
- Pattison, D., M. Segovia Garcia, W. Xie, F. Quail, M. Revie, R. I. Whitfield, I. Irvine (2016) Intelligent integrated maintenance for wind power generation, *Wind Energy*, 19(3), pp. 547–562.
- Pérez, J. M. P., F. P. G. Márquez, A. Tobias, M. Papaelias (2013) Wind turbine reliability analysis, *Renewable and Sustainable Energy Reviews*, 23, pp. 463–472.
- Pesántez, G., W. Guamán, J. Córdova, M. Torres, P. Benalcázar (2024) Reinforcement learning for efficient power systems planning: A review of operational and expansion strategies, *Energies*, 17(9), p. 2167.
- Pincirolì, L., P. Baraldi, G. Ballabio, M. Compare, E. Zio (2021) Deep reinforcement learning based on proximal policy optimization for the maintenance of a wind farm with multiple crews, *Energies*, 14(20), p. 6743.
- Pires Leite, G. d. N., A. M. Araujo, P. A. Carvalho Rosas (2018) Prognostic techniques applied to maintenance of wind turbines: A concise and specific review, *Renewable and Sustainable Energy Reviews*, 81(2), pp. 1917–1925.
- Qing, Y., Y. Tong, Z. Qi, Y. Li (2022) A survey on explainable reinforcement learning: Concepts, algorithms, challenges, URL <https://arxiv.org/abs/2211.06665>.
- Rahimilarki, R., Z. Gao, N. Jin, A. Zhang (2022) Convolutional neural network fault classification based on time-series analysis for benchmark wind turbine machine, *Renewable Energy*, 185, pp. 916–931.

- Raknes, N. T., K. Ødeskaug, M. Stålhan, L. M. Hvattum (2017) Scheduling of maintenance tasks and routing of a joint vessel fleet for multiple offshore wind farms, *Journal of Marine Science and Engineering*, 5(1), p. 11.
- Rashid, T., M. Samvelyan, C. Schroeder de Witt, G. Farquhar, J. N. Foerster, S. Whiteson (2018) QMIX: Monotonic value function factorisation for deep multi-agent reinforcement learning, in: *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, vol. 80 of *Proceedings of Machine Learning Research*, PMLR, pp. 4295–4304, conference held in Stockholm, Sweden, 10–15 July 2018.
- Ren, Z., A. S. Verma, Y. Li, J. J. E. Teuwen, Z. Jiang (2021) Offshore wind turbine operations and maintenance: A state-of-the-art review, *Renewable and Sustainable Energy Reviews*, 144, p. 110886.
- Rinaldi, G., A. C. Pillai, P. R. Thies, L. Johanning (2020) Multi-objective optimization of the operation and maintenance assets of an offshore wind farm using genetic algorithms, *Wind Engineering*, 44(4), pp. 390–409.
- Rinaldi, G., P. R. Thies, L. Johanning (2021) Current status and future trends in the operation and maintenance of offshore wind turbines: A review, *Energies*, 14(9), p. 2484.
- Romero, A., S. Soua, T.-H. Gan, B. Wang (2018) Condition monitoring of a wind turbine drive train based on its power dependant vibrations, *Renewable Energy*, 123, pp. 817–827.
- Rowell, D., B. Jenkins, J. Carroll, D. McMillan (2022) How does the accessibility of floating wind farm sites compare to existing fixed bottom sites?, *Energies*, 15(23), p. 8946.
- Safaei, F., E. Chatelet, J. Ahmadi (2020) Optimal age replacement policy for parallel and series systems with dependent components, *Reliability Engineering & System Safety*, 197, p. 106798.
- Sandborn, P., G. Haddad, A. Kashani-Pour, X. Lei (2014) Development of a maintenance option model to optimize offshore wind farm sustainment, in: *Proceedings of the 32nd ASME Wind Energy Symposium*, held at AIAA SciTech Forum, National Harbor, Maryland, USA, 13–17 January 2014.
- Saraswati, N., T. Stehly, A. Dewan, A. Delmarre (2017) Operation and maintenance map of u.s. offshore wind farms, Tech. Rep. ECN-E-17-028, Energy research Centre of the Netherlands (ECN) and National Renewable Energy Laboratory (NREL).
- Sarker, B. R., T. Ibn Faiz (2016) Minimizing maintenance cost for offshore wind turbines following multi-level opportunistic preventive strategy, *Renewable Energy*, 85, pp. 104–113.
- Scheu, M. N., A. Kolios, T. Fischer, F. Brennan (2017) Influence of statistical uncertainty of component reliability estimations on offshore wind farm availability, *Reliability Engineering & System Safety*, 168, pp. 28–39.
- Scheu, M. N., L. Tremps, U. Smolka, A. Kolios, F. Brennan (2019) A systematic failure mode effects and criticality analysis for offshore wind turbine systems towards integrated condition based maintenance strategies, *Ocean Engineering*, 176, pp. 118–133.

- Schouten, T. N., R. Dekker, M. Hekimoğlu, A. S. Eruguz (2022) Maintenance optimization for a single wind turbine component under time-varying costs, *European Journal of Operational Research*, 300(3), pp. 979–991.
- Schulman, J., F. Wolski, P. Dhariwal, A. Radford, O. Klimov (2017) Proximal policy optimization algorithms, URL <https://arxiv.org/abs/1707.06347>.
- SGS (2026) End of warranty inspection, Company webpage, URL <https://www.sgs.com/en/services/end-of-warranty-inspection>.
- Shafiee, M. (2015) Maintenance logistics organization for offshore wind energy: Current progress and future perspectives, *Renewable Energy*, 77, pp. 182–193.
- Si, G., T. Xia, N. Gebraeel, D. Wang, E. Pan, L. Xi (2025a) Holistic opportunistic maintenance scheduling and routing for offshore wind farms, *Renewable and Sustainable Energy Reviews*, 207, p. 114991.
- Si, G., T. Xia, D. Wang, N. Gebraeel, E. Pan, L. Xi (2025b) Maintenance scheduling and vessel routing for offshore wind farms with multiple ports considering day-ahead wind-wave predictions, *Applied Energy*, 379, p. 124915.
- Sikorska, J. Z., M. Hodkiewicz, L. Ma (2011) Prognostic modelling options for remaining useful life estimation by industry, *Mechanical Systems and Signal Processing*, 25(5), pp. 1803–1836.
- Sinha, Y., J. A. Steel (2015) Failure prognostic schemes and database design of a software tool for efficient management of wind turbine maintenance, *Wind Engineering*, 39(4), pp. 453–477.
- Sinha, Y., J. A. Steel, J. A. Andrawus, K. Gibson (2013) A SMART software package for maintenance optimisation of offshore wind turbines, *Wind Engineering*, 37(6), pp. 569–577.
- Sperstad, I. B., M. Stålhane, I. Dinwoodie, O.-E. V. Endrerud, R. Martin, E. Warner (2017) Testing the robustness of optimal access vessel fleet selection for operation and maintenance of offshore wind farms, *Ocean Engineering*, 145, pp. 334–343.
- Stålhane, M., E. E. Halvorsen-Weare, L. M. Nonås (2016) A decision support system for vessel fleet analysis for maintenance operations at offshore wind farms, Working paper, SINTEF.
- Stålhane, M., E. E. Halvorsen-Weare, L. M. Nonås, G. Pantuso (2019) Optimizing vessel fleet size and mix to support maintenance operations at offshore wind farms, *European Journal of Operational Research*, 276(2), pp. 495–509.
- Stehly, T., P. Duffy, D. Mulas Hernando (2024) Cost of wind energy review: 2024 edition, Presentation NREL/PR-5000-91775, National Renewable Energy Laboratory.
- Stetco, A., F. Dinmohammadi, X. Zhao, V. Robu, D. Flynn, M. Barnes, J. Keane, G. Nenadic (2019) Machine learning methods for wind turbine condition monitoring: A review, *Renewable Energy*, 133, pp. 620–635.
- Stock-Williams, C., S. K. Swamy (2019) Automated daily maintenance planning for offshore wind farms, *Renewable Energy*, 133, pp. 1393–1403.

- Su, J., J. Huang, S. Adams, Q. Chang, P. A. Beling (2022) Deep multi-agent reinforcement learning for multi-level preventive maintenance in manufacturing systems, *Expert Systems with Applications*, 192, p. 116323.
- Sutton, R. S., A. G. Barto (2018) *Reinforcement Learning: An Introduction*, 2 ed., MIT Press, Cambridge, MA.
- Taboada, J. V., V. Díaz-Casas, X. Yu (2021) Reliability and maintenance management analysis on offshore wind turbines (OWTs), *Energies*, 14(22), p. 7662.
- Tautz-Weinert, J., S. J. Watson (2017) Using SCADA data for wind turbine condition monitoring: A review, *IET Renewable Power Generation*, 11(4), pp. 382–394.
- Tracht, K., J. Westerholt, P. Schuh (2013) Spare parts planning for offshore wind turbines subject to restrictive maintenance conditions, *Procedia CIRP*, 7, pp. 563–568, forty-Sixth CIRP Conference on Manufacturing Systems 2013.
- Turnbull, A., J. Carroll (2021) Cost benefit of implementing advanced monitoring and predictive maintenance strategies for offshore wind farms, *Energies*, 14(16), p. 4922.
- Tusar, M. I. H., B. R. Sarker (2022a) Maintenance cost minimization models for offshore wind farms: A systematic and critical review, *International Journal of Energy Research*, 46(4), pp. 3739–3765.
- Tusar, M. I. H., B. R. Sarker (2022b) Spare parts control strategies for offshore wind farms: A critical review and comparative study, *Wind Engineering*, 46(5), pp. 1629–1656.
- uit het Broek, M. A. J., J. Veldman, S. Fazi, R. Greijdanus (2019) Evaluating resource sharing for offshore wind farm maintenance: The case of jack-up vessels, *Renewable and Sustainable Energy Reviews*, 109, pp. 619–632.
- Valet, A., T. Altenmüller, B. Waschneck, M. C. May, A. Kuhnle, G. Lanza (2022) Opportunistic maintenance scheduling with deep reinforcement learning, *Journal of Manufacturing Systems*, 64, pp. 518–534.
- Van Horenbeek, A., L. Pintelon (2012) A dynamic prognostic maintenance policy for multi-component systems, *IFAC Proceedings Volumes*, 45(31), pp. 115–120, 2nd IFAC Workshop on Advanced Maintenance Engineering, Services and Technology (AMEST 2012), Sevilla, Spain, 22–23 November 2012.
- Verleijdonk, P., W. van Jaarsveld, S. Kapodistria (2024) Scalable policies for the dynamic traveling multi-maintainer problem with alerts, *European Journal of Operational Research*, 319(1), pp. 121–134.
- Vermeer, L. J., J. N. Sørensen, A. Crespo (2003) Wind turbine wake aerodynamics, *Progress in Aerospace Sciences*, 39(6–7), pp. 467–510.
- Vianna, W. O. L., T. Yoneyama (2018) Predictive maintenance optimization for aircraft redundant systems subjected to multiple wear profiles, *IEEE Systems Journal*, 12(2), pp. 1170–1181.

- Wan, Y.-H., E. Ela, K. Orwig (2010) Development of an equivalent wind plant power-curve: Preprint, in: *Proceedings of WindPower 2010 Conference and Exhibition*, NREL/CP-550-48146, National Renewable Energy Laboratory (NREL), Dallas, TX, USA, conference held in Dallas, Texas, USA, 23–26 May 2010.
- Wang, S., Y. Vidal, F. Pozo (2026) Recent advances in wind turbine condition monitoring using SCADA data: A state-of-the-art review, *Reliability Engineering & System Safety*, 267, p. 111838.
- Wei, S., X. Jin, Y. Bao, H. Li (2019) Reinforcement learning in maintenance of civil infrastructures, in: *Workshop on Reinforcement Learning for Real Life, co-located with the 36th International Conference on Machine Learning (ICML 2019)*, workshop held in Long Beach, California, USA, June 2019.
- Williams, R. J. (1992) Simple statistical gradient-following algorithms for connectionist reinforcement learning, *Machine Learning*, 8(3–4), pp. 229–256.
- WindEurope (2024) Wind energy in europe: 2023 statistics and the outlook for 2024–2030, Tech. rep., WindEurope, Brussels, Belgium, accessed: 2026-01-09.
- Wu, M. (2014) Numerical analysis of docking operation between service vessels and offshore wind turbines, *Ocean Engineering*, 91, pp. 379–388.
- Xia, J., G. Zou (2023) Operation and maintenance optimization of offshore wind farms based on digital twin: A review, *Ocean Engineering*, 268, p. 113322.
- Xie, L., X. Rui, S. Li, X. Hu (2019) Maintenance optimization of offshore wind turbines based on an opportunistic maintenance strategy, *Energies*, 12(14), p. 2650.
- Xie, L., X. Rui, S. Li, X. Hu (2020) An opportunistic maintenance strategy for offshore wind turbine based on accessibility evaluation, *Wind Engineering*, 44(5), pp. 455–468.
- Yang, W., R. Court, J. Jiang (2013) Wind turbine condition monitoring by the approach of SCADA data analysis, *Renewable Energy*, 53, pp. 365–376.
- Ying, F., L. Huang, Y. Fu (2026a) A generation model of uncertain environmental scenario sets for offshore wind farm maintenance, *Sustainable Energy, Grids and Networks*, 46, p. 102161.
- Ying, F., L. Huang, Y. Liu, L. Liu, Y. Fu (2026b) Condition-based opportunistic maintenance strategy for offshore wind farm based on accurate condition assessment results, *Reliability Engineering & System Safety*, 269, p. 112107.
- Ying, F., L. Huang, M. Song, Y. Fu (2025) Optimization of an opportunistic maintenance strategy for wind turbines based on subsystem fault chains, *Ocean Engineering*, 339, p. 122106.
- Zhang, C., Y.-F. Li, D. W. Coit (2023) Deep reinforcement learning for dynamic opportunistic maintenance of multi-component systems with load sharing, *IEEE Transactions on Reliability*, 72(3), pp. 863–877.
- Zhang, E., F. Shen, S. Liu, G. Chen, F. Zhang, S. Li (2024) Offshore wind power digital twin modeling system for intelligent operation and maintenance applications, in: *E3S Web of Conferences*, vol. 546, p. 02010, conference-proceedings article; event details not uniquely identifiable from the available metadata.

- Zhang, K., Y. Feng, Y. Cui, Z. Wang, F. Yang (2018) Quantity optimization of spare parts for offshore wind farm based on component updating, *IOP Conference Series: Earth and Environmental Science*, 186(5), p. 012024, conference paper published in IOP Conference Series: Earth and Environmental Science.
- Zhang, N., W. Si (2020) Deep reinforcement learning for condition-based maintenance planning of multi-component systems under dependent competing risks, *Reliability Engineering & System Safety*, 203, p. 107094.
- Zhao, F. J., Y. Zhou (2022) Wind farm maintenance scheduling using soft actor-critic deep reinforcement learning, in: *Proceedings of the 2022 Global Reliability and Prognostics and Health Management Conference (PHM-Yantai 2022)*, pp. 1–6, conference held in Yantai, China, 13–16 October 2022.
- Zhu, W., M. Fouladirad, C. Bérenguer (2016) A multi-level maintenance policy for a multi-component and multifailure mode system with two independent failure modes, *Reliability Engineering & System Safety*, 153, pp. 50–63.
- Zhuang, L., A. Xu, X.-L. Wang (2023) A prognostic driven predictive maintenance framework based on bayesian deep learning, *Reliability Engineering & System Safety*, 234, p. 109181.

Glossary

Conventions

The following conventions are used throughout this thesis:

- Time is generally indexed by k , representing the current decision or observation step.
- Component-specific quantities are indexed by c .
- Turbine-specific quantities are indexed by i .
- Maintenance actions are indexed by a .
- Predicted quantities are denoted using the subscript “pred”, while true latent quantities are denoted using the subscript “true”.
- Probability distributions are represented using standard statistical notation, e.g., $\mathcal{N}(\mu, \sigma^2)$ for the normal distribution and $\text{Weibull}(\alpha, \lambda)$ for the Weibull distribution.
- Remaining Useful Life (RUL) predictions are represented probabilistically through predictive distributions.
- Maintenance costs are represented as the sum of material, labour, transportation, and down-time costs.
- Offshore accessibility constraints are represented through weather-dependent operability limits associated with each vessel class.
- In the reinforcement learning formulations, the environment is modelled as a Markov Decision Process (MDP) or Partially Observable Markov Decision Process (POMDP).
- Reinforcement learning policies are denoted by π , rewards by R , observations by O , actions by A , and states by S .
- Expectations are denoted by $\mathbb{E}[\cdot]$ and probabilities by $P(\cdot)$.

List of Symbols and Notations

α_c	Weibull shape parameter of component c
β_c	Weibull scale parameter of component c
γ	Discount factor
Δk	Decision timestep length
$\mathbb{E}[\cdot]$	Expectation operator
$\mathcal{N}(\mu, \sigma^2)$	Normal probability distribution
$\mu_c(k)$	Running reward mean for component c
$\phi_c^k(t)$	Predicted RUL probability distribution for component c
π	Policy function of the DRL agent
$\sigma_c(k)$	Standard deviation of RUL prediction uncertainty
$\mathbf{A}(k)$	Vector of maintenance actions for all components
$\mathbf{O}(k)$	Observation vector for all components
a	Maintenance action index
$a_{i,h}$	Availability indicator of turbine i at hour h
$A_c(k)$	Maintenance action applied to component c at timestep k
c	Component index
C_{dt}	Downtime cost
C_{maint}	Maintenance cost
$C_{\text{mobil}}^{j(c,a)}$	Vessel mobilization cost
$C_{\text{trans}}^{j(c,a)}$	Transportation cost for vessel $j(c,a)$
d	Planning day index
D_{ij}	Distance between turbines i and j
E_h^{farm}	Total wind farm energy production at hour h
E_h^{lost}	Lost energy production at hour h
$E_{i,h}^{\text{avail}}$	Available energy production of turbine i at hour h
$f(t; \beta_c, \alpha_c)$	Weibull probability density function
$F_{\text{cons}}^{j(c,a)}$	Fuel consumption rate of vessel $j(c,a)$
h	Hourly timestep index
H	Prediction horizon
H_s	Significant wave height

i	Wind turbine index
j	Vessel or transportation asset index
$j(c, a)$	Vessel required for maintenance activity a on component c
k	Discrete decision timestep
$M(k)$	Technician availability at timestep k
N	Total number of components
N_{WT}	Number of wind turbines
$O_c(k)$	Observed health state of component c
$P(v)$	Wind turbine power output at wind speed v
$P_{\text{fail},c}(k; k+H)$	Failure probability of component c over horizon H
$R_c(k)$	Reward associated with component c at timestep k
$R_{\text{day}}^{j(c,a)}$	Daily vessel charter rate
R_{hour}	Hourly technician rate
$S(k)$	Spare-part availability at timestep k
$S_c(k)$	True degradation state of component c
t	Continuous time variable
T	Total simulation horizon
$T(c, a)$	Duration of maintenance activity a on component c
T_{dt}	Downtime duration
T_{elapsed}	Elapsed operational time
T_{major}	Duration of major repair
T_{minor}	Duration of minor repair
$T_{\text{replacement}}$	Duration of replacement
$\text{NormReward}_c(k)$	Normalized reward
$\text{RUL}_{\text{pred},c}(k)$	Predicted Remaining Useful Life of component c
$\text{RUL}_{\text{true},c}(k)$	True Remaining Useful Life of component c
$U \sim \text{Unif}(0, 1)$	Uniform random variable
v	Wind speed
v_{in}	Cut-in wind speed
v_{out}	Cut-out wind speed
v_r	Rated wind speed

$V(s)$	State-value function
$W(k)$	Weather-window availability at timestep k
$x_{c,a,k}$	Binary maintenance decision variable

List of Abbreviations

A2C	Advantage Actor-Critic
A3C	Asynchronous Advantage Actor-Critic
BDL	Bayesian Deep Learning
CBM	Condition-Based Maintenance
CNN	Convolutional Neural Network
CTV	Crew Transfer Vessel
DCMAC	Deep Centralized Multi-Agent Actor-Critic
DDPG	Deep Deterministic Policy Gradient
DDQN	Double Deep Q-Network
DQN	Deep Q-Network
DRL	Deep Reinforcement Learning
EMA	Exponential Moving Average
FMECA	Failure Modes, Effects, and Criticality Analysis
GCN	Graph Convolutional Network
HLV	Heavy Lift Vessel
IOHMM	Input-Output Hidden Markov Model
KPI	Key Performance Indicator
LCoE	Levelized Cost of Energy
LSTM	Long Short-Term Memory
MARL	Multi-Agent Reinforcement Learning
MDP	Markov Decision Process
MILP	Mixed-Integer Linear Programming
MLP	Multilayer Perceptron
MPC	Model Predictive Control
O&M	Operations and Maintenance
OWF	Offshore Wind Farm
PDF	Probability Density Function
PHM	Prognostics and Health Management
POMDP	Partially Observable Markov Decision Process
PPO	Proximal Policy Optimization

QMIX	Multi-Agent Value-Based Reinforcement Learning Algorithm
RES	Renewable Energy Sources
RL	Reinforcement Learning
RUL	Remaining Useful Life
SAC	Soft Actor-Critic
SCADA	Supervisory Control and Data Acquisition
SOV	Service Operation Vessel
VRP	Vehicle Routing Problem
W-QMIX	Weighted QMIX

Summary

Offshore wind is expected to play an essential role in the energy transition, but its long-term competitiveness depends strongly on the ability to operate and maintain wind farms reliably and cost-effectively. Operation and maintenance (O&M) is particularly challenging offshore because maintenance decisions are shaped not only by component failures, but also by weather-driven accessibility, vessel availability, logistics constraints, and the cost of lost production. At the same time, condition monitoring and prognostic methods increasingly provide information about the health and remaining useful life (RUL) of wind turbine components. However, this information only becomes valuable if it can be translated into feasible maintenance decisions. This thesis investigates how offshore wind farm maintenance planning can be supported by decision-support models that jointly account for uncertain component health information, metocean-driven accessibility, marine logistics, and economic consequences. The central research question concerns the trade-offs between optimization-based and learning-based decision-making for offshore wind maintenance.

To address this question, the thesis first reviews the state of the art in offshore wind O&M decision support and identifies a lack of integrated models that combine health information, accessibility, logistics, and maintenance timing in a consistent way. Based on these gaps, a modeling environment is developed to represent the main dynamics of offshore wind O&M. This environment includes a representative fixed-bottom offshore wind farm, critical components, transportation and access assets, weather-dependent operability, synthetic RUL predictions, multi-level maintenance actions, and cost components. It provides a common basis for comparing different decision-making approaches under the same assumptions. Within this environment, the thesis first develops an optimization-based maintenance planning model. A rolling-horizon mixed-integer linear programming formulation integrates probabilistic RUL information into maintenance planning, deciding when maintenance should be initiated, which components should be maintained, and which type of intervention should be performed. The results show that prognostic information can reduce O&M cost and downtime when it is embedded into a decision model that accounts for offshore access and logistics. At the same time, the analysis highlights limitations of optimization-based planning, including horizon dependence, missed grouping opportunities, increasing computational complexity, and the difficulty of representing uncertainty in detail.

The thesis then investigates learning-based maintenance planning. Offshore wind O&M is formulated as a partially observable sequential decision problem, and a deep reinforcement learning framework is developed using a PPO agent. The agent observes uncertain RUL predictions and learns to select among a set of different actions. The results show that learning-based policies can reduce unexpected failures and adapt maintenance decisions to evolving component conditions, especially in smaller systems. However, the learning-based approach also introduces challenges related to scalability, interpretability, simulator fidelity, and operational validation.

Overall, this thesis shows that optimization-based and learning-based approaches have different strengths and limitations. Optimization models provide transparency, explicit constraint handling, and interpretable schedules, making them valuable for tactical planning and engineering validation. Learning-based methods offer adaptivity and fast policy execution under uncertainty, but require further development before operational deployment. The thesis therefore contributes to the development of integrated offshore wind O&M decision-support systems in which prognostic information, offshore logistics, accessibility constraints, and economic trade-offs are considered jointly.

Samenvatting

Offshore windenergie zal een belangrijke rol spelen in de energietransitie, maar de langetermijnconcurrentiekracht ervan hangt sterk af van het betrouwbaar en kosteneffectief kunnen exploiteren en onderhouden van windparken. Operatie en onderhoud (O&M) is offshore bijzonder uitdagend, omdat onderhoudsbeslissingen niet alleen worden bepaald door componentstoringen, maar ook door weersafhankelijke toegankelijkheid, beschikbaarheid van vaartuigen, logistieke beperkingen en de kosten van productieverlies. Tegelijkertijd leveren conditiebewaking en prognostische methoden steeds meer informatie over de gezondheid en resterende levensduur (RUL) van windturbinecomponenten. Deze informatie is echter pas waardevol wanneer zij kan worden vertaald naar haalbare onderhoudsbeslissingen.

Dit proefschrift onderzoekt hoe onderhoudsplanung voor offshore windparken kan worden ondersteund door beslissingsondersteunende modellen die gezamenlijk rekening houden met onzekere gezondheidsinformatie van componenten, metocean-gedreven toegankelijkheid, mariene logistiek en economische gevolgen. De centrale onderzoeksvraag richt zich op de afwegingen tussen optimalisatiegebaseerde en learning-based besluitvorming voor offshore windonderhoud.

Om deze vraag te beantwoorden, wordt in dit proefschrift eerst de stand van de techniek op het gebied van O&M-beslissingsondersteuning voor offshore wind onderzocht. Hierbij wordt een gebrek aan geïntegreerde modellen geïdentificeerd die gezondheidsinformatie, toegankelijkheid, logistiek en onderhoudstiming op consistente wijze combineren.

Op basis van deze lacunes wordt een modelleringsomgeving ontwikkeld die de belangrijkste dynamieken van offshore wind O&M weergeeft. Deze omgeving omvat een representatief fixed-bottom offshore windpark, kritieke componenten, transport- en toegangsmiddelen, weersafhankelijke operabiliteit, synthetische RUL-voorspellingen, onderhoudsacties op meerdere niveaus en kostencomponenten. Hiermee ontstaat een gemeenschappelijke basis om verschillende besluitvormingsmethoden onder dezelfde aannames te vergelijken.

Binnen deze omgeving ontwikkelt het proefschrift eerst een optimalisatiegebaseerd onderhoudsplanningsmodel. Een rolling-horizon mixed-integer linear programming-formulering integreert probabilistische RUL-informatie in onderhoudsplanung en bepaalt wanneer onderhoud moet worden gestart, welke componenten moeten worden onderhouden en welk type interventie moet worden uitgevoerd. De resultaten tonen aan dat prognostische informatie O&M-kosten en downtime kan verminderen wanneer deze wordt ingebed in een beslissingsmodel dat rekening houdt met offshore toegankelijkheid en logistiek. Tegelijkertijd laat de analyse beperkingen van optimalisatiegebaseerde planning zien, waaronder horizonafhankelijkheid, gemiste groepeeringsmogelijkheden, toenemende computationele complexiteit en de moeilijkheid om onzekerheid gedetailleerd te modelleren.

Vervolgens onderzoekt het proefschrift learning-based onderhoudsplanung. Offshore wind O&M wordt geformuleerd als een partieel observeerbaar sequentieel beslissingsprobleem, waarvoor een deep reinforcement learning-raamwerk wordt ontwikkeld met een PPO-agent. De agent observeert onzekere RUL-voorspellingen en leert te kiezen uit een reeks verschillende acties. De resultaten tonen aan dat learning-based beleid onverwachte storingen kan verminderen en onderhoudsbeslissingen kan aanpassen aan veranderende componentcondities, vooral in kleinere systemen. Deze benadering brengt echter ook uitdagingen met zich mee op het gebied van schaalbaarheid, interpreteerbaarheid, simulatorgetrouwheid en operationele validatie.

Over het geheel genomen laat dit proefschrift zien dat optimalisatiegebaseerde en learning-based benaderingen verschillende sterke punten en beperkingen hebben. Optimalisatiemodellen bieden transparantie, expliciete beperkingsmodelleren en interpreteerbare plannings, waardoor zij waardevol zijn voor tactische planning en technische validatie. Learning-based methoden bieden adaptiviteit en snelle beleidsuitvoering onder onzekerheid, maar vereisen verdere ontwikkeling voordat zij operationeel kunnen worden ingezet. Dit proefschrift draagt daarmee bij aan de ontwikkeling

van geïntegreerde beslissingsondersteunende systemen voor offshore wind O&M, waarin prognostische informatie, offshore logistiek, toegankelijkheidsbeperkingen en economische afwegingen gezamenlijk worden beschouwd.

About the author

Marco Borsotti was born on the 10th of April 1995 in Rome, Italy. He studied mechanical engineering at Sapienza University of Rome where he graduated with honors and developed a growing interest in renewable energy.

Before beginning his PhD, Marco worked for two years in industry, gaining experience in asset-management across photovoltaic and wind projects. That period exposed him to the operational realities of the sector and left a lasting mark on how he approaches research. It taught him to value methods that are not only technically robust, but also practical, interpretable, and grounded in the realities of implementation.

He later joined the Department of Maritime & Transport Technology at Delft University of Technology, where he started his PhD in October 2022 under the supervision of Prof. Dr. Rudy R. Negenborn and Dr. Xiaoli Jiang. His doctoral research focuses on decision support for offshore wind farm operation and maintenance, with particular emphasis on integrating condition information, uncertainty, and marine logistics into maintenance planning models. Through this work, he hopes to contribute to making offshore wind more reliable, more efficient, and better equipped to support the much needed energy transition.

Alongside his research, Marco has enjoyed contributing to the academic environment more broadly, including through the organisation of two editions of the WindED summer school, which brought together speakers and participants from different countries, backgrounds and disciplines. Experiences like this strengthened his appreciation for collaboration, exchange, and the human side of research. Looking ahead, he aspires to continue working on problems that combine analytical depth with practical relevance, whether in academia, industry, or ideally somewhere between the two.

Outside research, Marco used to be an enthusiastic rower, in many ways, this sport mirrors the PhD itself: progress is usually the result of many smaller strokes.



Marco Borsotti

Publications

1. **Borsotti, M., R. R. Negenborn, and X. Jiang (2024).** *Model predictive control framework for optimizing offshore wind O&M*. In: *Advances in Maritime Technology and Engineering*, CRC Press, pp. 533–546. doi: 10.1201/9781003508762-65.
2. **Borsotti, M., R. L. Patrão, X. Jiang, A. Napoleone, and R. R. Negenborn (2025).** *A deep reinforcement learning framework for O&M planning in offshore wind*. Accepted for publication in the proceedings of the joint 16th International Conference on Computational Logistics (ICCL 2025) and 1st EURO Mini Conference on Maritime Optimization and Logistics (EUROMar 2025), Springer Lecture Notes in Computer Science.
3. **Borsotti, M., X. Jiang, and R. R. Negenborn (2026a).** *Review of deep reinforcement learning for offshore wind farm maintenance planning*. *Wind Energy Science*, 11, pp. 1185–1204. doi: 10.5194/wes-11-1185-2026.
4. **Borsotti, M., X. Jiang, and R. R. Negenborn (2026b).** *Stochastic optimization of predictive maintenance scheduling for offshore wind farms*. *Ocean Engineering*, 357, 125424. doi: 10.1016/j.oceaneng.2026.125424.
5. **Borsotti, M., R. R. Negenborn, and X. Jiang (2026c).** *A review of multi-horizon decision-making for operation and maintenance of fixed-bottom offshore wind farms*. *Renewable and Sustainable Energy Reviews*, 226, 116450. doi: 10.1016/j.rser.2025.116450.

Conference contributions

1. **Wind Energy Science Conference (WESC) 2023** - Glasgow, United Kingdom, 23–26 May 2023.
Optimizing the Operation and Maintenance of Offshore Wind Farms.
Oral presentation.
2. **7th International Conference on Maritime Technology and Engineering (MARTECH) 2024** - Lisbon, Portugal, 14–16 May 2024.
Model Predictive Control Framework for Optimizing Offshore Wind O&M.
Oral presentation.
3. **North American Wind Energy Academy/WindTech Conference (NAWEA) 2024** - New Brunswick, New Jersey, United States, 30 October–1 November 2024.
Stochastic Optimization of Predictive Maintenance Scheduling for Offshore Wind Farms.
Oral presentation.
4. **Wind Energy Science Conference (WESC) 2025** - Nantes, France, 24–27 June 2025.
Deep Reinforcement Learning for Offshore Wind O&M Planning.
Oral presentation.
5. **International Conference on Computational Logistics (ICCL) 2025** - Delft and Rotterdam, the Netherlands, 8–10 September 2025.
A Deep Reinforcement Learning Framework for O&M Planning in Offshore Wind.
Oral presentation.
6. **Wind Energy 2050** - Delft, the Netherlands, 11 November 2025.
A Deep Reinforcement Learning Framework for O&M Planning in Offshore Wind.
Poster presentation.

TRAIL Thesis Series

The following list contains the most recent dissertations in the TRAIL Thesis Series. For a complete overview of more than 400 titles, see the TRAIL website: www.rsTRAIL.nl.

The TRAIL Thesis Series is a series of the Netherlands TRAIL Research School on transport, infrastructure and logistics.

Borsotti, M., *Towards Learning-Based Decision-Making for Maintenance of Offshore Wind Farms*, T2026/26, October 2026, TRAIL Thesis Series, the Netherlands

Momen, S., *Electrified Multimodal Urban Transport Planning and Operations under Uncertainty*, T2026/25, October 2026, TRAIL Thesis Series, the Netherlands

Gagno Azolin Tassarolo, L., *Shared Mobility in Urban Transport Systems: Accessibility and user behaviour under public transport disruptions*, T2026/24, October 2026, TRAIL Thesis Series, the Netherlands

Torabi Kachousangi, F., *Enhancing Urban Mobility: Role of shared mobility on travel behaviour*, T2026/23, September 2026, TRAIL Thesis Series, the Netherlands

Tang, X., *Energy-Aware Operations for Decarbonizing Ports and Shipping under Uncertainty: Learning- and optimization-based approaches*, T2026/22, September 2026, TRAIL Thesis Series, the Netherlands

Yoo, S., *Perceived Accessibility by Air Transportation: Determinants, Heterogeneity, and Climate Policy Implications*, T2026/21, September 2026, TRAIL Thesis Series, the Netherlands

Cunillera Pérez, A., *Robust Train Trajectory Optimization and Train Motion Model Calibration*, T2026/20, September 2026, TRAIL Thesis Series, the Netherlands

Öztürker, M., *On the Bright Side of Vehicle Automation: The effects of service quality improvements and ride experience on users' preferences for automated public transport*, T2026/19, July 2026, TRAIL Thesis Series, the Netherlands

Faber, R.M., *Travel in Transition: A longitudinal exploration of travel behaviour*, T2026/18, June 2026, TRAIL Thesis Series, the Netherlands

Xiong, X., *Hydrodynamics-Aware Model Predictive Formation Control for Energy-Efficient Multi-Vessel Systems*, T2026/17, July 2026, TRAIL Thesis Series, the Netherlands

Louro, T.V., *Electric Bicycles: impacts on job accessibility and spatial equity in Brazilian megacities*, T2026/16, July 2026, TRAIL Thesis Series, the Netherlands

Knaap, R.J.H. van der, *Multi-period Line Planning and Timetabling for Varying Railway Passenger Demand*, T2026/15, June 2026, TRAIL Thesis Series, the Netherlands

Liang, K., *Adaptive Takeover Time Budgets in Conditionally Automated Driving*, T2026/14, June 2026, TRAIL Thesis Series, the Netherlands

Suryana, L.E., *Operationalising Meaningful Human Control for Reason-Responsive Decision-Making in Automated Vehicles*, T2026/13, June 2026, TRAIL Thesis Series, the Netherlands

Yan, H., *Cycling Speed: Variation and Stability within Rides*, T2026/12, May 2026, TRAIL Thesis Series, the Netherlands

Abhishek, D., *Safe Navigation of Autonomous Vessels in Inland Waterways under Uncertain and Abnormal Operational Condition*, T2026/11, June 2026, TRAIL Thesis Series, the Netherlands

Garrido-Valenzuela, F., *Pixels, People, Places: Computer Vision and Image Embeddings for Perception-Aware Urban Analytics*, T2026/10, April 2026, TRAIL Thesis Series, the Netherlands

Spierenburg, L., *Advances in the Analysis of Residential Segregation and Urban Riots*, T2026/9, April 2026, TRAIL Thesis Series, the Netherlands

Boot, M., *Evaluating Experiences with Smart Cycling Technologies: Sensor-based evaluations of outdoor cycling experiences with Smart Cycling Technologies*, T2026/8, March 2026, TRAIL Thesis Series, the Netherlands

Wen, X., *Data-Driven Spatial-Temporal Modeling for Bicycle Traffic Prediction*, T2026/7, March 2026, TRAIL Thesis Series, the Netherlands

Wang, Z., *Optimising Performance of Automatic Train Operation on Railway Networks*, T2026/6, March 2026, TRAIL Thesis Series, the Netherlands

Hadi, A.H., *DEM Modelling of Multi-Component Segregation in the Blast Furnace Charging System*, T2026/5, February 2026, TRAIL Thesis Series, the Netherlands

Farhani, M., *Demand Management Strategies for Operations of Shared Mobility Services*, T2026/4, February 2026, TRAIL Thesis Series, the Netherlands

Yao, X., *Driving Heterogeneity in Traffic Flow Theory: An action-based framework for identification, modelling, and simulation*, T2026/3, January 2026, TRAIL Thesis Series, the Netherlands

Versluis, N.D., *Optimising Railway Traffic Management under Radio-Based Distance-to-Go Signalling*, T2026/2, January 2026, TRAIL Thesis Series, the Netherlands

Jiao, Y., *Proactive Collision Risk Quantification in Multi-directional Traffic Interactions*, T2026/1, January 2026, TRAIL Thesis Series, the Netherlands

Asadi, M., *Accessibility and Road Safety: Integration of road safety in accessibility evaluation*, T2025/20, November 2025, TRAIL Thesis Series, the Netherlands



TRAIL

Summary

Offshore wind maintenance requires decisions under uncertain component health, weather-driven accessibility, logistical constraints, and costly downtime. This thesis develops a framework to investigate how prognostic information can support such decisions. It compares optimization-based maintenance planning using mixed-integer linear programming with learning-based planning using deep reinforcement learning. The results highlight their respective strengths and limitations, supporting the development of more adaptive, reliable, and cost-effective offshore wind O&M decision-support systems.

About the Author

Marco Borsotti graduated with honors in Mechanical Engineering from Sapienza University of Rome in 2021. After two years in renewable-energy asset management, he conducted his PhD research in the Maritime and Transport Technology Department of Delft University of Technology.

TRAIL Research School ISBN 978-90-5584-401-2



Radboud University



rijksuniversiteit
 groningen



UNIVERSITY OF TWENTE.

TU/e

Technische Universiteit
 Eindhoven
 University of Technology