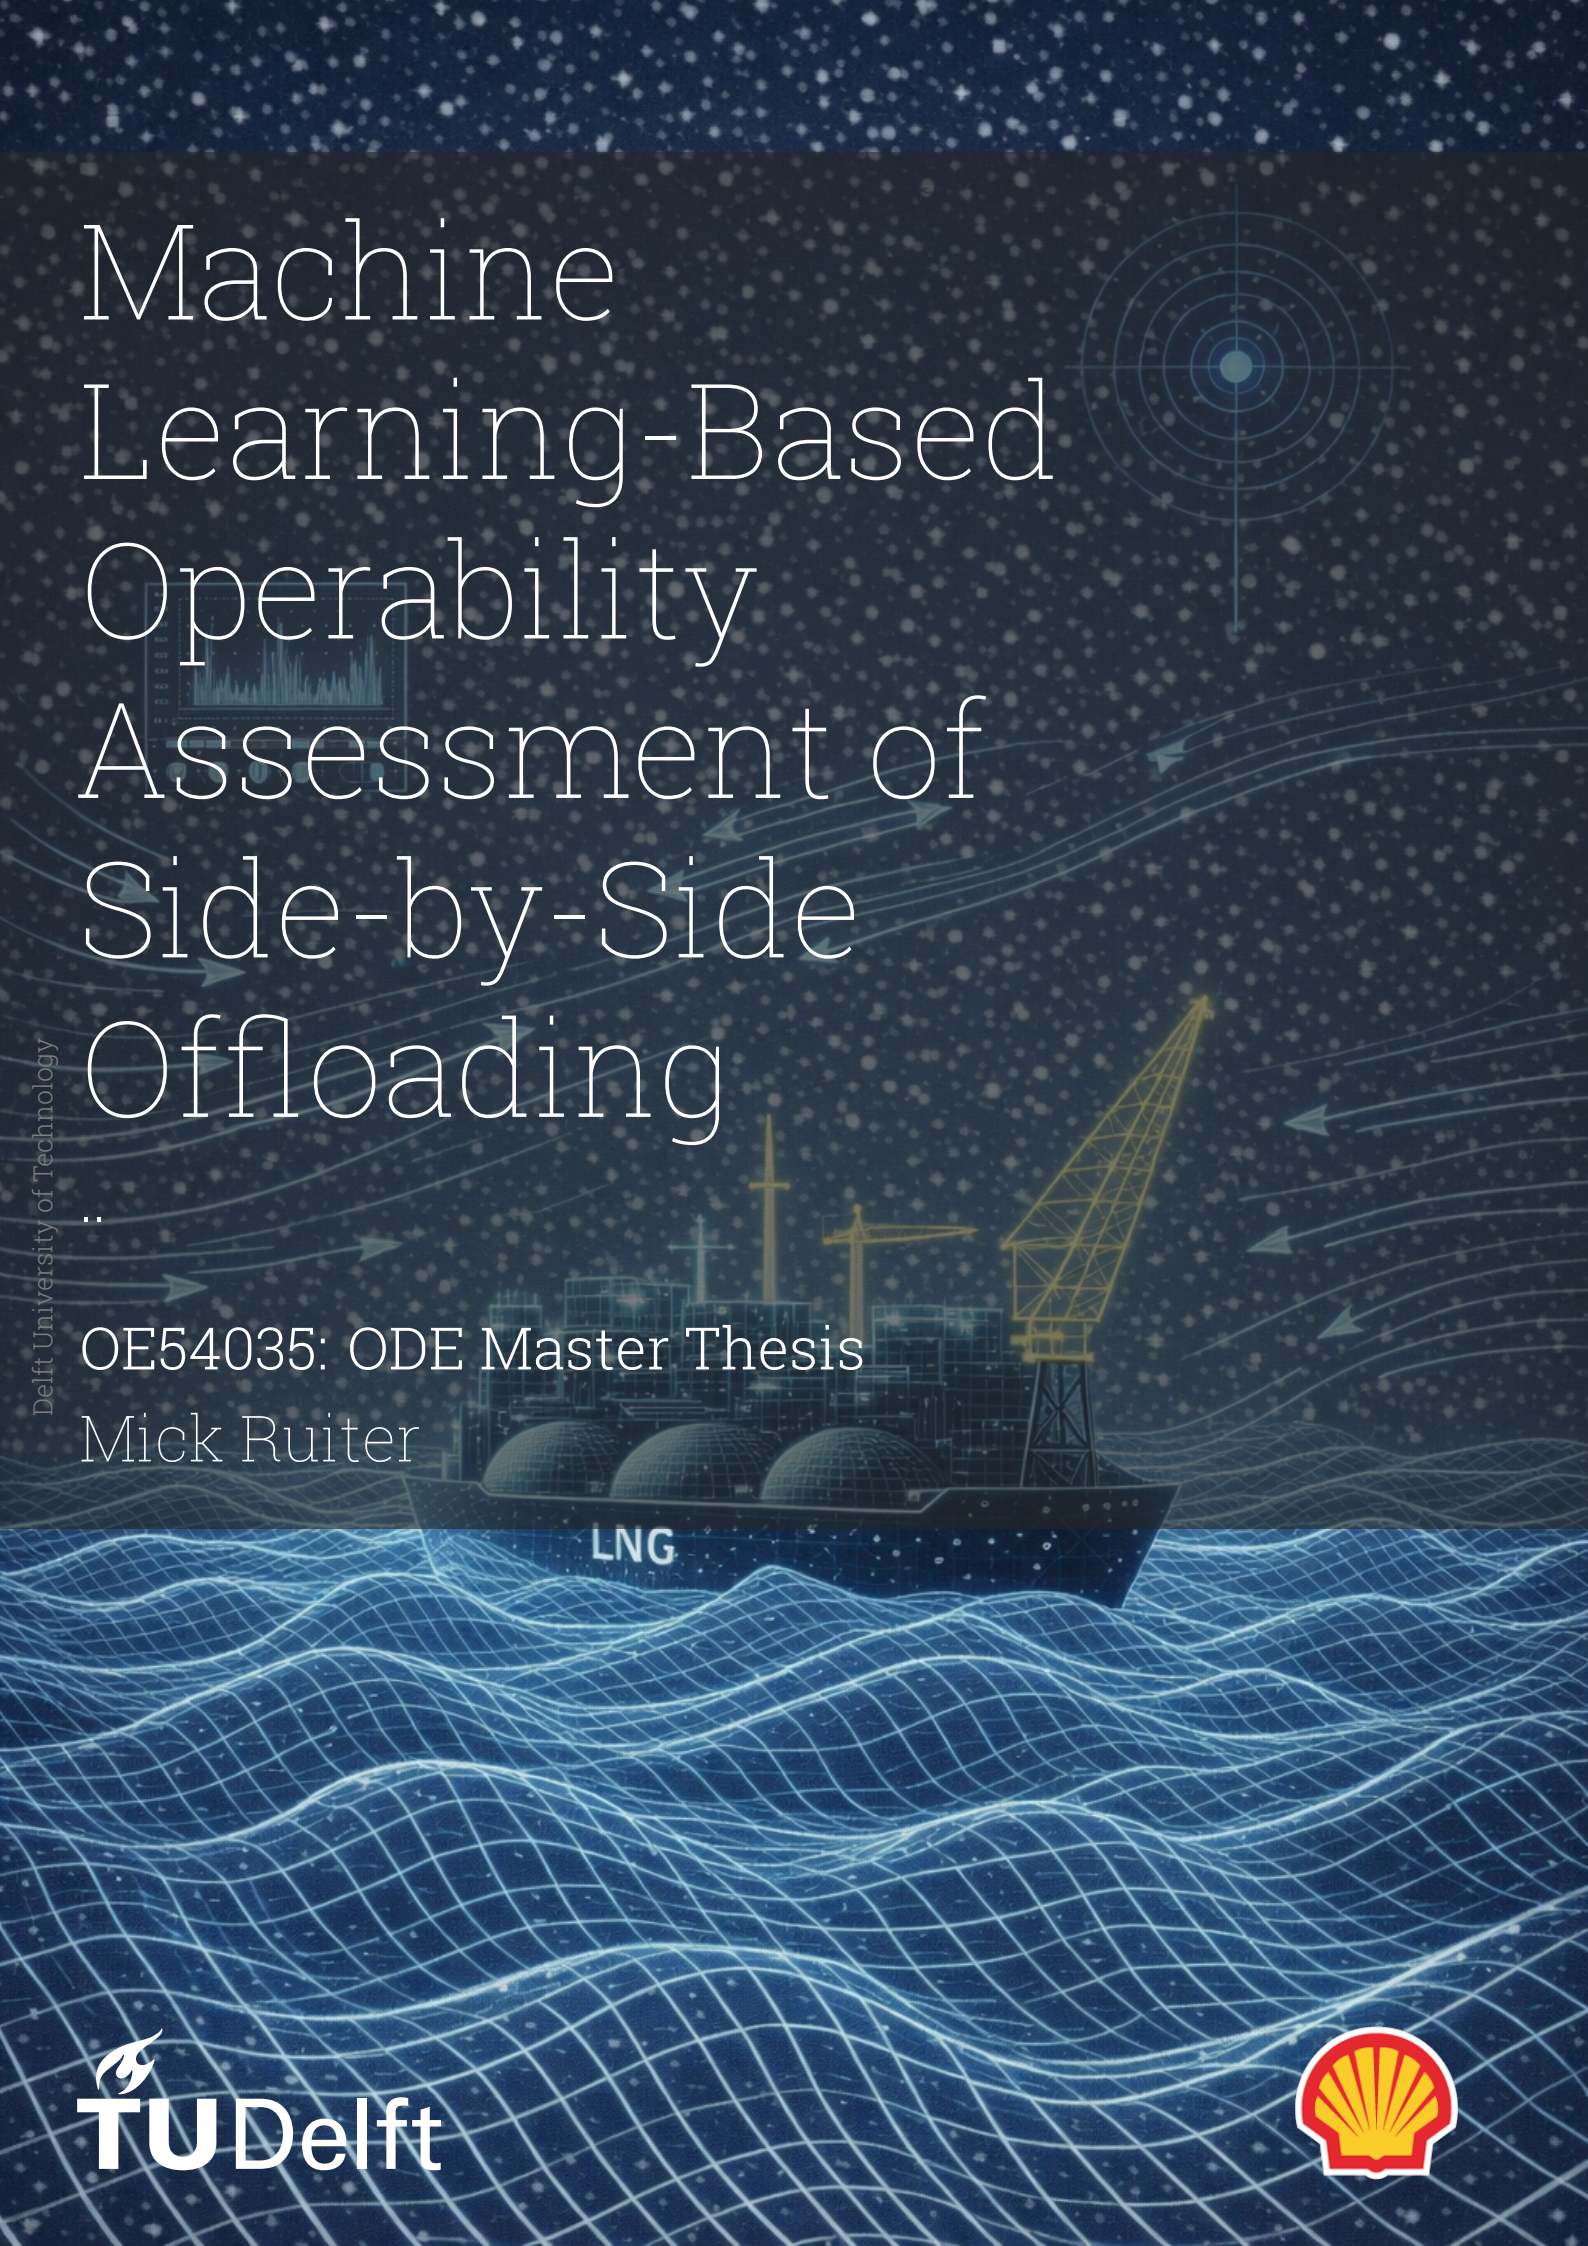
Machine Learning-Based Operability Assessment of Side-by-Side Offloading

..

OE54035: ODE Master Thesis

Mick Ruiter

Delft University of Technology



Machine Learning-Based Operability Assessment of Side-by-Side Offloading

by

Mick Ruiter

Thesis Committee:	Prof.dr. A.V. Metrikine - Committee Chair
	Ir. A.C.M. van der Stap
	Dr. H. Wang
	Dr. R. van Vliet - Company Supervisor
Project Duration:	July, 2025 - April, 2026
Faculty:	Faculty of Mechanical Engineering, Delft

Cover: FLNG & LNGC with environmental influences - created by Microsof Copilot

TU Delft

Preface

This master thesis marks the end of my time as a student at Delft University of Technology. I started with the bachelor in Mechanical Engineering and found my way to the Offshore Engineering master, which I enjoyed enormously. From the technical depth of the courses to the strong connection with industry through lectures, events and projects, this programme taught me a great deal about the world of offshore engineering and the people who work in it.

This thesis was carried out at Shell, and I would like to start by thanking all my colleagues there for welcoming me to the team and for their genuine enthusiasm and interest in the work. There was always someone available for a discussion, and that made a real difference.

A very special thank you goes to Renate, my supervisor at Shell, for her dedication throughout this project. Our discussions ranged beyond the thesis itself and were extremely valuable. Her feedback, which she even sent in the evenings or at weekends, was invaluable, and I am grateful for the time and energy she put in.

I also want to thank Octavio and Jason for their involvement and expertise. The simulation campaign you set up created the foundation for everything in this thesis. Jason, thank you for your enthusiasm about the potential of ML applications in this field. It was contagious.

Hongrui, thank you for your guidance especially in the earlier phases of the thesis, when machine learning was still new territory for me. Your help pushed me further than I would have gotten on my own. Andrei, thank you for your technical insights and for keeping a critical eye on what the models could and could not do. Andre, thank you for the discussions on industry relevance, on how to present results in the right way, and for keeping sight of the bigger picture.

I also want to thank the people at MARIN, and Jaap in particular, for their involvement in the project and for providing the full time series from the simulation campaign. Jaap, your perspectives and your willingness to help me understand the underlying theory were greatly appreciated.

Finally, I want to thank my family, my girlfriend and my friends. You always reminded me that there are a lot more fun and important things to life than writing a thesis.

This marks the end of my time as a student. I look forward to building a career in offshore engineering, and I hope to carry with me the lessons and values from my years at TU Delft.

AI Usage Statement

In completing this thesis, I made use of Microsoft Copilot and Claude to support coding and writing tasks. These tools were used to improve code quality, assist with phrasing and structure, and help identify inconsistencies in the text.

Mick Ruiter
Delft, April 2026

Abstract

Operability assessment of multibody offshore operations is commonly performed using large numbers of high-fidelity time-domain simulations. When environmental conditions are discretised across realistic ranges of wave, wind, and current parameters, simulation campaigns can involve up to 10^5 individual three-hour runs, resulting in substantial computational cost. This challenge is shared across a broad class of coupled floating-body operations, including offshore hydrocarbon offloading, CO_2 ship-to-ship transfer, and floating wind installation. This thesis investigates the use of data-driven prediction models to support simulation-driven operability assessment while retaining response-based decision criteria, using an FLNG–LNGC side-by-side offloading configuration as a case study.

The study is based on a simulation dataset produced by the Maritime Research Institute the Netherlands (MARIN) for Shell, consisting of time-domain simulations of a spread-moored FLNG with an LNG carrier in side-by-side configuration. From this dataset, a sea-wave subset of 8 317 simulations is used for model development and evaluation. Each simulation includes summary response extrema as well as full time-series records of wave elevation, vessel motions, mooring line tensions, and fender loads.

Two complementary machine learning (ML) approaches are considered. The first is a summary-statistics-based learning (SSBL) approach, which predicts operability-limiting response extrema directly from static environmental descriptors using feed-forward regression models. The second is a time-series-based learning (TSBL) approach operating on full response time series. Within this framework, a segment-based forecasting model based on long short-term memory (LSTM) networks and a full time-series model based on transformer architectures with cross-attention heads are employed.

For time-series modelling, response signals are decomposed into wave-frequency and low-frequency components to distinguish first-order wave-induced behaviour from second-order drift-dominated dynamics. At a training fraction of approximately 70%, wave-frequency models achieve mean coefficients of determination R^2 between 0.6 and 0.8, while low-frequency LSTM models attain mean R^2 values of approximately 0.8. To reduce training requirements while preserving operability-critical behaviour, a compact training dataset of 1 500 simulations is derived using a two-stage selection strategy combining stratified sampling over key sea-state parameters with K-means clustering in response space.

For operability classification based on response extrema across the full sea-wave dataset, a multilayer perceptron trained on the reduced dataset achieves an overall accuracy of 98.22% for 6 816 unseen sea states. Time-series-based models achieve an overall classification accuracy of 86.7% and enable reconstruction of complete multibody response histories when the training data sufficiently represent dominant physical regimes, including head-sea, beam, and near-beam wave conditions.

From a computational perspective, the original MARIN sea-only time-domain simulations require approximately 1 428 core-hours to evaluate 8 317 sea states. In contrast, transformer-based time-series models require approximately 18.9 GPU-hours for one-time training on 1 500 sea states and 2.85 GPU-hours for prediction over remaining sea states, while extrema-based model prediction is effectively instantaneous. These results demonstrate that data-driven models can substantially reduce the computational burden of simulation-driven operability assessment for coupled floating-body systems more broadly.

Contents

Preface i

Summary ii

Nomenclature xv

1 Introduction 1

2 General Review of Technology 4
2.1 What is AI? . . . 4
2.2 State-of-the-Art Applications within Engineering . . . 9
2.3 Bottlenecks in Offshore Structures Engineering . . . 12
2.4 AI in Offshore Structures Engineering . . . 15
2.4.1 Generative Design of Offshore Support Structures . . . 15
2.4.2 Operability Assessment of Complex Marine Operations . . . 18
2.4.3 AI-Driven Predictive Maintenance and Digital Twin Systems for Offshore Platforms 21
2.5 Conclusion . . . 23

3 In Depth Review of Technology 24
3.1 Body Motions . . . 25
3.2 Mooring Systems for FLNG Facilities . . . 26
3.3 Offloading Methods . . . 27
3.3.1 Side-by-Side (SBS) Offloading . . . 27
3.3.2 Tandem Offloading . . . 28
3.3.3 Implications of Mooring and Offloading Choices . . . 29
3.4 Side-by-Side Offloading of an spread-moored FLNG . . . 29
3.5 Criteria for Side-by-Side Offloading Operability . . . 29
3.6 Hydromechanical Considerations . . . 30
3.7 Conventional Operability Assessment . . . 32
3.8 Machine Learning in Multibody/Offshore Dynamics . . . 33
3.8.1 Long-Horizon Time-Series Motion Prediction . . . 33
3.8.2 ML for Mooring Line Tension and Load Prediction . . . 34
3.8.3 ML for Operability Assessment and Decision Support . . . 34
3.9 Toward ML-Driven SBS Operability Assessment and Research Gaps . . . 34

4 Case Description and Theoretical Framework 36
4.1 Configuration and Coordinate Systems . . . 37
4.2 Operability Criteria . . . 38
4.3 Simulation Framework . . . 38
4.4 Environmental Inputs . . . 40
4.5 Simulation Outputs . . . 42
4.6 Hydrodynamic Background . . . 44

4.6.1 RAO-based representation for irregular seas . . . 44
4.6.2 Second Order Wave (Drift) Forces . . . 45
4.6.3 Remarks on Hydrodynamic Coupling and Operability . . . 45
4.6.4 Coupled Multi-Body System (12-DOF) . . . 45
4.6.5 Interpretation . . . 46
4.7 Low-Frequency vs. Wave-Frequency Motions . . . 47
4.7.1 Wave-Frequency Motions . . . 47
4.7.2 Low-Frequency Motions . . . 47
4.7.3 Engineering Implications . . . 48
4.8 Discussion . . . 48

5 Methodology 50
5.1 Summary Statistics-Based Learning . . . 51
5.1.1 Random Forest . . . 52
5.1.2 Multi-Layer Perceptron (MLP) . . . 52
5.2 Time Series-Based Learning . . . 54
5.2.1 Response signal decomposition . . . 55
5.2.2 Segment-Based Sequence to Sequence Learning . . . 57
5.2.3 Full Time Series-Based Learning . . . 62
5.3 Evaluation and Methodological Synthesis . . . 66

6 Execution 68
6.1 Summary-Statistics-Based Learning (SSBL) . . . 68
6.1.1 Data and Preprocessing . . . 68
6.1.2 Training, Validation, and Test Strategy . . . 69
6.1.3 Multi-Head Multilayer Perceptron (MLP) Architecture . . . 70
6.2 Time Series-Based Learning (TSBL) . . . 71
6.2.1 Data preprocessing . . . 71
6.2.2 Butterworth Signal Decomposition . . . 72
6.2.3 Segment-Based Sequence-to-Sequence Learning . . . 74
6.2.4 Full Time Series Based Learning . . . 79

7 Preliminary Results Analysis 83
7.1 SSBL Models . . . 83
7.1.1 Random Forest Results . . . 83
7.1.2 MLP Results . . . 84
7.1.3 SSBL Operability Assessment . . . 85
7.1.4 Interpertation of Results . . . 88
7.2 Time Series-Based Models . . . 89
7.2.1 Segment-Based Sequence to Sequence Learning . . . 89
7.2.2 Full Time Series-Based Learning . . . 98
7.2.3 Interpertation of Results . . . 118

8 Final Results 120
8.1 Dataset Selection . . . 121
8.1.1 Selection Methodology . . . 121
8.2 SSBL Operability Assessment . . . 124
8.2.1 False-Safe Environmental Combinations . . . 125
8.2.2 Computational Burden of the SSBL Operability Assessment . . . 126

8.3 TSBL Operability Assessment . . . 127
8.3.1 False-Safe Environmental Combinations . . . 127
8.3.2 Training and Inference Efficiency of the LF and WF Models . . . 128

9 Discussion 130
9.1 Input limitations and implications . . . 130
9.2 Evaluation of the Modelling Approach . . . 131
9.2.1 Summary-Statistics-Based Learning . . . 131
9.2.2 Time-Series-Based Learning: Segment-Based Prediction . . . 132
9.2.3 Time-Series-Based Learning: Full Reconstruction . . . 132
9.2.4 Low-Frequency Model Training and Optimisation . . . 133
9.2.5 Incorporating Extracted Low-Frequency Drift Forcing as Model Input . . . 134
9.2.6 Wind and Current Forcing . . . 135
9.2.7 Practical Meaning of Model Accuracy . . . 136
9.3 Minimal Dataset . . . 136
9.4 Alternative Approach: Linearised Frequency Domain Modelling . . . 137
9.5 Integrated Interpretation . . . 139

10 Conclusions & Recommendations 141
10.1 Conclusion . . . 141
10.1.1 Integrated Interpretation . . . 145
10.2 Recommendations . . . 146
10.2.1 Application with Extended Input Spaces and Operational Data . . . 146
10.2.2 Provide Physically Complete Model Inputs . . . 146
10.2.3 Improve Model Reliability and Data Efficiency . . . 147
10.2.4 Towards a General Operability Screening Model . . . 147
10.2.5 Exploring the Role of Frequency Domain Modelling . . . 147

References 149

A General Review of of Technology 158

B Theoretical Framework aNySIM 160
B.1 Environmental Modelling . . . 161
B.1.1 Linear Wave Modelling . . . 161
B.1.2 Wind and current fields . . . 164
B.2 Hydromechanic Load Models . . . 164
B.2.1 Linear Hydrostatics . . . 164
B.2.2 First Order Wave Loads . . . 166
B.2.3 Radiation Forces and Added Mass . . . 168
B.2.4 Viscous Damping . . . 169
B.2.5 Second-Order (Drift) Wave Loads . . . 169
B.2.6 Current Forces . . . 171
B.2.7 Wind Forces . . . 172
B.3 Mechanical Forces: Mooring Lines and Fenders . . . 172
B.3.1 FLNG Anchoring Mooring System (Linear Model) . . . 172
B.3.2 Side-by-Side Mooring Lines (Nonlinear and Dynamic Model) . . . 173
B.3.3 Fender Mechanics . . . 174
B.4 Rigid-Body Dynamics Framework . . . 175

B.4.1 Six-Degree-of-Freedom Vessel Dynamics . . . 175
B.4.2 Coordinate Transformations and Kinematic Consistency . . . 176
B.4.3 Coupled Multi-Body System (12-DOF) . . . 176
B.4.4 Interpretation . . . 177
B.5 Conclusion . . . 177

C Autocorrelation Function Plots 178
C.1 Wave Frequency ACF's . . . 178
C.2 Low Frequency ACFs . . . 178

D Detailed Architectures of Models 180
D.1 Segment-Based Sequence-to-Sequence Learning: Model Specifications . . . 181
D.1.1 Low-Frequency (LF) Model . . . 181
D.1.2 Wave-Frequency (WF) Model . . . 182
D.2 Full Time Series-Based Learning . . . 184
D.2.1 LF Model Specifics . . . 185
D.2.2 WF Model Specifics . . . 185
D.2.3 Hyperparameter Optimization . . . 185

E Prediction Statistics 188
E.1 SSBL Full Dataset Predictions . . . 188
E.2 TSBL Head Waves Prediction . . . 190

F Final Results Plots & Prediction Statistics 192
F.1 SSBL Insights . . . 192
F.2 False Safe Environmental Pattern Analysis . . . 194
F.2.1 Most Frequent Combinations Across All False Safe Cases . . . 194
F.2.2 False Safe Patterns for Manifold Surge Limit . . . 195
F.2.3 False Safe Patterns for Manifold Sway Limit . . . 195
F.2.4 False Safe Patterns for Line Tension Limit . . . 195
F.2.5 False Safe Patterns for Fender Force Limit . . . 196
F.3 TSBL Environmental Pattern Analysis . . . 196
F.3.1 Combinations of False Safe Classifications . . . 196
F.3.2 Combinations False-Safe Classification . . . 197
F.3.3 Combinations Causing Surge-Limit False Safe Cases . . . 197
F.3.4 Top Combinations Causing Sway-Limit False Safe Cases . . . 197
F.3.5 Combinations Causing Roll-Limit False Safe Cases . . . 198
F.3.6 Top Combinations Causing Line-Tension False Safe Cases . . . 198

List of Figures

1.1 Rendering of an FLNG–LNGC side-by-side offloading arrangement [2]. 1

2.1 Hierarchy of AI subfields. 5
2.2 Building blocks of an AI system. 5
2.3 Machine Learning Overview (Based on [5] & [4]) 6
2.4 Schematic representations of neural network architectures [4]. 7
2.5 Airbus generative design [21] 9
2.6 Structural Health Montioring Overview [34] 13
2.7 Generative design workflow 16
2.8 Workflow of a neural network-based surrogate model for operability assessment. 19
2.9 Schematic representation of digital twin framework [64]. 22

3.1 Schematic representation of an FLNG facility showing gas extraction, turret mooring, liquefaction, and LNG offloading [31]. 25
3.2 Schematic overview of the multi-body motion system. Each vessel responds in 6 DOFs, producing a coupled 12-DOF [67]. 26
3.3 A turret-mooring arrangement example [69] 26
3.4 A spread-mooring arrangement example [70] 27
3.5 Schematic of MLA envelopes: flanging envelope, Emergency Shut Down envelopes, and maximum reach [2]. 28
3.6 Tandem offloading configuration [69] 28

4.1 Sign convention and coordinate system used in the simulations [83]. 37
4.2 SBS mooring layout showing fenders and line leads, plotted as rectangualar for length between perpendicular. 37
4.3 Integrated workflow: DIFFRAC hydrodynamics $\rightarrow$ aNySIM time-domain simulation $\rightarrow$ operability statistics. 39
4.4 Comparison of environmental parameterization between datasets. 42
4.5 Illustration of first- and second-order motions: the rapid first-order response is superimposed on a slowly varying second-order modulation and a mean drift component. 48

5.1 Summary statistics ML model: from sea-state inputs to predicted operability parameters. 52
5.2 High-level overview of TSBL workflow. 55
5.3 Recursive prediction visualization with train/test split and in- and output windows. 57
5.4 High-level overview of the full time series-based learning framework. During training the transformer learns a mapping from wave elevation and environmental conditions to 12 DOF vessel motions and mooring tensions across known sea states ($T_p = 7, 9, 10, 11$ s; $H_s = 2.0$ m). At inference time only the wave signal and environmental scalars are needed; no reference simulation is required. 62

6.1 Overview MLP Architecture 70

6.2 Frequency spectra of FLNG surge (left) and heave (right) with the 0.03 Hz low-frequency cutoff used for LF-WF decomposition. 73
6.3 Example LF–WF decomposition for FLNG surge. Top: LF component (with mean) and zero-mean WF component. Bottom: original signal and reconstructed LF + WF. 73
6.4 Overview of the segment-based sequence-to-sequence architecture showing the WF model (upper branch, $H_c = 50$ s, overlap 150 s) and LF model (lower branch, $H_c = 100$ steps, overlap 900 s), both using dual LSTM encoders with recursive chunking. . . . 74
6.5 Overview of the full time series learning architecture. 79
6.6 Training and validation loss curves for the LF (left) and WF (right) models. The dashed line indicates the epoch with the lowest validation loss for each model. Both curves illustrate the convergence behaviour of the models under the selected optimization configuration. 82

7.1 Random Forest performance: scaled RMSE (left) and R^2 (right) across training-data splits. 84
7.2 MLP performance: scaled RMSE (left) and R^2 (right) across training-data splits. 84
7.3 MLP predictions for minimum and maximum roll for both vessels, including operability limits. 85
7.4 MLP predictions for minimum and maximum relative surge and sway, with operability limits indicated. 85
7.5 MLP predictions for maximum line tension and maximum fender force with operability limits included. 86
7.6 Confusion-matrix summary of uptime/downtime classifications for the test set (62,586 simulations). 87
7.7 Recursive low-frequency predictions for the FLNG. 90
7.8 Recursive low-frequency surge and sway predictions for the LNGC. 90
7.9 Low-frequency ACFs and memory times computed using the $1/e$ envelope method. . . 91
7.10 Recursive wave-frequency predictions for the FLNG (early window, first 250 s). 92
7.11 Recursive wave-frequency predictions for the LNGC (early window, first 250 s). 92
7.12 Recursive wave-frequency heave and roll predictions for the LNGC (late window, 2250–2500 s). 92
7.13 Recursive wave-frequency heave and roll predictions for the FLNG (late window, 2250–2500 s). 93
7.14 R^2 over time for recursive wave-frequency FLNG predictions. 93
7.15 R^2 over time for recursive wave-frequency LNGC predictions. 93
7.16 ACF & memory time for FLNG roll (wave-frequency). 94
7.17 ACF & memory time for LNGC roll (wave-frequency). 94
7.18 Learning-curve analysis for the wave-frequency time-series model using FLNG roll as reference parameter. All models are trained for 100 epochs per training fraction. 95
7.19 Learning-curve analysis for the low-frequency segment-based LSTM model using FLNG surge as reference parameter. Models are trained for 100 epochs at each training fraction. 95
7.20 Wave-elevation time series with envelope and detected wave groups over a three-hour simulation. 96
7.21 Comparions of wave signal for different $H_{\mathrm{s,sea}}$. 98
7.22 FLNG surge and sway predictions vs. reference for unseen $H_{\mathrm{s,sea}}$. 99
7.23 LNGC surge and sway predictions vs. reference for unseen $H_{\mathrm{s,sea}}$. 99
7.24 LF motion cross-correlation for varying $H_{\mathrm{s,sea}}$. 100
7.25 Surge motion (LF) comparison across $H_{\mathrm{s,sea}}$ (FLNG). 100

7.26 Comparison of wave signal for different $H_{\mathrm{s,sea}}$, mean $\pm$ std bars (with min–max ranges). 100
7.27 FLNG first 250 s: heave and roll predictions vs. reference (unseen $H_{\mathrm{s,sea}}$). 101
7.28 LNGC first 250 s: heave and roll predictions vs. reference (unseen $H_{\mathrm{s,sea}}$). 101
7.29 WF motion cross-correlation for varying $H_{\mathrm{s,sea}}$. 101
7.30 Wave-frequency (WF) LNGC roll motion across $H_{\mathrm{s,sea}}$ (FLNG). 101
7.31 Comparison of PSD of different $T_{\mathrm{p,sea}}$ 103
7.32 Comparison of wave signals for different $T_{\mathrm{p,sea}}$. 103
7.33 FLNG surge and sway predictions for unseen peak period $T_{\mathrm{p,sea}}$. 103
7.34 Low-frequency drift-force spectra for all training and test cases. The shaded region indicates the low-frequency range used for LF response modelling. Changes in $T_{\mathrm{p,sea}}$ reshape the drift-force spectrum rather than producing a simple amplitude scaling. . . . 104
7.35 LF motion cross-correlation matrix for varying $T_{\mathrm{p,sea}}$. 104
7.36 Wave-frequency motion correlation matrix for varying $T_{\mathrm{p,sea}}$ (FLNG). 104
7.37 Comparison of LF responses for different $T_{p,sea}$, mean $\pm$ std bars (with min–max ranges).105
7.38 FLNG heave and roll predictions for unseen peak period $T_{\mathrm{p,sea}}$. 105
7.39 LNG heave and roll predictions for unseen peak period $T_{\mathrm{p,sea}}$. 106
7.40 WF motion cross-correlation matrix for varying $T_{\mathrm{p,sea}}$. 106
7.41 Wave-frequency roll motion comparison (FLNG). 106
7.42 Coverage of significant wave height versus peak period ($H_{\mathrm{s,sea}}$–$T_{\mathrm{p,sea}}$). 108
7.43 Coverage of wind velocity versus significant wave height (U–$H_{\mathrm{s,sea}}$). 108
7.44 Calm-condition relative surge response: FLNG–LNGC prediction vs. reference ($R^2 = 0.039$, RMSE = 0.0167 m). The model captures the mean offset but not the detailed LF structure under low-excitation conditions. 109
7.45 Calm-condition relative sway response: FLNG–LNGC prediction vs. reference ($R^2 = 0.326$, RMSE = 0.0167 m). The model reproduces the mean sway offset reasonably well under low-excitation conditions. 109
7.46 Calm-condition roll response: FLNG (top, $R^2 = 0.611$, RMSE = 0.0003 rad) and LNGC (bottom, $R^2 = 0.598$, RMSE = 0.0002 rad) prediction vs. reference. 110
7.47 Calm-condition mooring line tension response for Line 10: prediction vs. reference ($R^2 = -6.002$, RMSE = 6693.6 N). Despite the low R^2, the model correctly captures the maximum tension level. 110
7.48 Critical-condition relative surge response: FLNG–LNGC prediction vs. reference ($R^2 = 0.922$, RMSE = 0.208 m). 111
7.49 Critical-condition relative sway response: FLNG–LNGC prediction vs. reference ($R^2 = 0.602$, RMSE = 0.1296 m). 111
7.50 Critical-condition roll response: FLNG (top, $R^2 = 0.819$, RMSE = 0.0024 rad) and LNGC (bottom, $R^2 = 0.762$, RMSE = 0.0010 rad) prediction vs. reference. 111
7.51 Critical-condition mooring line tension response for Line 10: prediction vs. reference ($R^2 = 0.840$, RMSE = 68,351 N). 111
7.52 Confusion-matrix summary of uptime/downtime classifications for all 456 head-sea states.112
7.53 Comparions of wave signal for different $H_{\mathrm{s,sea}}$. 114
7.54 FLNG surge and sway predictions vs. reference for unseen θ_{sea}. 115
7.55 LNGC surge and sway predictions vs. reference for unseen θ_{sea}. 115
7.56 LF motion cross-correlation matrix for varying θ_{sea}. 116
7.57 LF sway motion comparison across new wave-direction cases (LNGC). 116
7.58 LNGC surge and sway predictions vs. reference for unseen θ_{sea}. 116
7.59 FLNG heave and roll WF predictions vs. reference for unseen wave direction θ_{sea}. . . . 117

7.60 LNGC heave and roll WF predictions vs. reference for unseen wave direction θ_{sea}. . . . 117
7.61 WF motion cross-correlation matrix for varying θ_{sea}. 117
7.62 Wave-frequency roll motion comparison across new wave-direction cases (FLNG). . . . 117

8.1 Coverage of significant wave height V_{wind} (m) in the selected dataset (red) relative to the full dataset (blue). 122
8.2 Displayed coverage of key sea-state parameters in the selected training subset. 123
8.3 Coverage of mean surge vs. mean sway for the FLNG. 123
8.4 Coverage of mean surge vs. mean sway for the LNGC. 123
8.5 Coverage of maximum mooring line tension in the selected subset (red) relative to the full dataset (blue). 124
8.6 Coverage of FLNG roll in the selected subset (red) compared to the full dataset (blue). . 124
8.7 Confusion-matrix summary of SSBL operability classification for all 6816 cases (Accuracy: 98.22%; False Safe: 1.81%; False Unsafe: 1.73%). 124
8.8 Confusion-matrix summary of TSBL uptime/downtime classification for 6,816 simulations (Accuracy: 86.7%; False Safe: 19.9%; False Unsafe: 5.5%). 127

9.1 LF drift forcing extracted from the wave elevation. The upper panel shows $\eta(t)$ and its group envelope. The lower panel shows the corresponding LF drift force consisting of a mean component and a slowly varying oscillatory contribution. 135
9.2 Comparison between RAO based reconstruction and the wave frequency filtered MARIN simulation for roll motion of the FLNG (top) and LNGC (bottom). The dashed coloured lines show the linear reconstruction and the solid black lines show the MARIN wave frequency response. 138

B.1 Mean JONSWAP spectrum [109] 162
B.2 Wave record analysis and regeneration: transformation between frequency-domain spectrum and time-domain realisation [98]. 163
B.3 Two complementary representations of linear wave-to-motion behaviour [98] 163
B.4 Shift of the centre of buoyancy due to heel: formation of emerged and immersed wedges, movement of B to B_ϕ, and definition of the metacentre N_ϕ [98]. 165
B.5 Linear potential-flow boundary conditions for two interacting floating bodies. 166
B.6 Time-domain representation of an irregular wave elevation (bottom) and the associated second-order drift force (top). The drift force consists of a constant and a slowly varying part that follows wave groupiness. 170

C.1 ACF & memory time for FLNG Heave (WF). Memory time: 17 s; Dominant period: 8.31 s.178
C.2 ACF & memory time for LNGC Heave (WF). Memory time: 11 s; Dominant period: 8.54 s.178
C.3 ACF & memory time for FLNG Pitch (WF). Memory time: 20 s; Dominant period: 8.37 s. 178
C.4 ACF & memory time for LNGC Pitch (WF). Memory time: 15 s; Dominant period: 8.62 s. 178
C.5 ACF & memory time for FLNG Surge (LF). Memory time: 585 s; Zero up-cross: 134.91 s.179
C.6 ACF & memory time for FLNG Sway (LF). Memory time: 698 s; Zero up-cross: 111.98 s. 179
C.7 ACF & memory time for FLNG Yaw (LF). Memory time: 442 s; Zero up-cross: 115.06 s. 179
C.8 ACF & memory time for LNGC Surge (LF). Memory time: 583 s; Zero up-cross: 134.91 s.179
C.9 ACF & memory time for LNGC Surge (LF). Memory time: 738 s; Zero up-cross: 120.91 s.179
C.10 ACF & memory time for LNGC Yaw (LF). Memory time: 424 s; Zero up-cross: 111.48 s. 179

F.1 Roll predictions (max/min) for FLNG and LNGC with operability limits. 192

F.2 Relative surge and sway predictions (max/min) with operability limits. 193
F.3 Predictions for maximum line tension and maximum fender force with operability limits. 193

List of Tables

2.1 Examples of AI Applications and Their Industry Disruptions 9
2.2 Overview of Generative Design for Offshore Support Structures 17
2.3 Overview of AI-Enhanced Operability Assessment for Offshore Operations 20

3.1 Overview of vessel motions in six degrees of freedom (6-DOF). 25
3.2 Overview of Typical SBS Offloading Operability Criteria (after [31]). 30

4.1 Operability criteria used to evaluate each simulation (adapted from MARIN Summary Report). 38
4.2 Generic matrix (sea-only): parameterization used in the study. 40
4.3 Wind and current combinations applied to each sea-only wave case. 41
4.4 Batch matrix for swell-only and combined sea–swell conditions (adapted from MARIN dataset). 41
4.5 Overview of individual simulation output parameters. 43
4.6 Comparison of wave-frequency and low-frequency motions. 48

6.1 Model output channels: vessel motions (FLNG and LNGC) and mooring line tensions. . 72
6.2 Environmental conditions corresponding to the frequency spectra in Figure 6.2. 72
6.3 Low-frequency model training hyperparameters (encoder–decoder with recursive rollout). 75
6.4 WF hyperparameters. 77

7.1 Representative environmental regimes producing False Safe (FS) cases across roll and manifold responses. 87
7.2 Environmental condition set including significant wave height. 89
7.3 Summary of dominant period, memory, and cycle counts (for a 3 h record: 70% training = 7 560 s; 30% test = 3 240 s). 97
7.4 Significant wave heights and constant environmental parameters across training and predicted cases. 98
7.5 Typical periods and constant environmental parameters across training and predicted cases. 102
7.6 Dominant wave–period ranges for each peak period $T_{\mathrm{p,sea}}$. 102
7.7 Calm-condition environmental parameters for test `36490_01XMF_02_000_038_01`. 109
7.8 Critical-condition environmental parameters for test. 111
7.9 Typical periods and constant environmental parameters across training and predicted cases. 114

8.1 Parameter weighting scheme for dataset selection. 122
8.2 Coverage of selected vs. full range for critical parameters 123
8.3 Per-criterion violation and false-safe rates on the 6,816 test cases. System-level false-safe rate is 1.81 % (67 cases), all involving line tension only. 125
8.4 False-safe regimes identified in the SSBL test set. All cases have $V_{cur} = 0.24$ m/s. . . . 125

8.5 Essential CPU specifications used for SSBL operability model training and inference . . 126
8.6 Most frequent V_{wind}–$H_{s,sea}$–$T_{p,sea}$ combinations in the 736 TSBL false-safe cases. All entries have $V_{cur} = 0.24$ m/s. 128
8.7 Essential GPU specifications used for LF and WF model training and inference 129

A.1 Detailed Overview of AI Methods by Subfield 159

D.1 LF motion encoder specification 181
D.2 LF wave CNN specification 182
D.3 LF motion decoder specification 182
D.4 WF encoder specification 182
D.5 WF wave CNN specification 182
D.6 WF motion decoder specification 183
D.7 LF wave CNN (1-channel elevation) 184
D.8 WF wave CNN (4 channels: elevation + derivatives) 184
D.9 Transformer encoder parameters (LF vs. WF) 185
D.10 RF baseline (LF mean only) 185
D.11 LF oscillation decoders 185
D.12 WF decoder architecture 185
D.13 Search spaces for the LF and WF models. 186
D.14 LF baseline configuration retained for deployment. 186
D.15 Optimised WF configuration retained for deployment. 187
D.16 Optimised Random Forest hyperparameters for position and tension prediction. 187

E.1 Top-5 most frequent 3-parameter combinations in False Safe cases for FLNG roll (min/-max). 188
E.2 Top-5 most frequent 3-parameter combinations in False Safe cases for LNGC roll (min/-max). 188
E.3 Top-5 most frequent 3-parameter combinations in False Safe cases for Manifold Distance X. 188
E.4 Top-5 most frequent 3-parameter combinations in False Safe cases for Manifold Distance Y. 189
E.5 Line Tension violation summary. 189
E.6 Fender Force violation summary. 189
E.7 Most difficult environmental condition ranges for each predicted parameter, ranked by False Safe (FS) rate. 190
E.8 Most difficult environmental conditions per tension line (Lines 1–8 left, 9–16 right). FS Count column widened for readability. 191

F.1 Performance metrics (R^2, RMSE, and scaled RMSE) for the summary-based MLP model. 194
F.2 Top 5 environmental combinations in all False Safe cases (swell parameters removed). 194
F.3 Top 5 environmental combinations for False Safe surge-limit violations (swell removed). 195
F.4 Top 5 environmental combinations for False Safe sway-limit violations (swell removed). 195
F.5 Top 5 environmental combinations for False Safe line tension violations (swell removed). 195
F.6 Top 5 environmental combinations for False Safe fender force violations (swell removed). 196
F.7 Top 5 most frequent 3□parameter combinations in False Safe cases. 196
F.8 Top 5 most frequent 4-parameter combinations in False Safe cases. 197
F.9 Top 5 environmental combinations leading to Relative Surge limit violations. 197

F.10 Top 5 environmental combinations leading to Relative Sway limit violations. 197
F.11 Top 5 environmental combinations leading to Roll limit violations. 198
F.12 Top 5 environmental combinations associated with tension-limit violations. 198

Nomenclature

Abbreviations

Abbreviation	Description
Machine Learning	
AI	Artificial Intelligence
CNN	Convolutional Neural Network
GAN	Generative Adversarial Network
LSTM	Long Short-Term Memory (network)
ML	Machine Learning
MLP	Multi-Layer Perceptron
RBM	Restricted Boltzmann Machine
RF	Random Forest
RNN	Recurrent Neural Network
SSBL	Summary Statistics-Based Learning
TSBL	Time Series-Based Learning
VAE	Variational Autoencoder
Performance Metrics and Loss Functions	
ACF	Autocorrelation Function
MSE	Mean Squared Error
RMSE	Root Mean Squared Error
RMS	Root Mean Square
Optimisation	
TPE	Tree-structured Parzen Estimator
Offshore Engineering	
DOF	Degrees of Freedom
FLNG	Floating Liquefied Natural Gas (facility)
FPSO	Floating Production Storage and Offloading
LNG	Liquefied Natural Gas
LNGC	LNG Carrier
MBL	Minimum Breaking Load
MLA	Marine Loading Arm
SBS	Side-by-Side (offloading configuration)
SHM	Structural Health Monitoring
Hydrodynamics and Wave Theory	
FFK	Froude–Krylov Force
JONSWAP	Joint North Sea Wave Project (wave spectrum)
LF	Low Frequency

Abbreviation	Description
QTF	Quadratic Transfer Function
RAO	Response Amplitude Operator
WF	Wave Frequency
Software and Simulation Tools	
DIFFRAC	MARIN frequency-domain diffraction/radiation panel code
XMF	Extensible Modelling Framework (aNySIM)
Computing	
CPU	Central Processing Unit
GB	Gigabyte
GPU	Graphics Processing Unit
HPC	High-Performance Computing
Organisations and Standards	
DNV	Det Norske Veritas
MARIN	Maritime Research Institute Netherlands

1

Introduction

Determining whether an offshore operation can be executed safely under given environmental conditions often relies on computationally expensive time-domain numerical simulations. In these simulations, system responses such as vessel motions, mooring line tensions and fender loads are calculated at every time step for a duration of typically 3 hours per environmental condition. For offshore operations involving floating systems, acceptable operating conditions are therefore not defined solely by environmental thresholds, but by response-based criteria such as vessel motions, relative displacements, mooring line tensions and fender loads.

This response-based approach is also critical for operations involving multiple floating bodies in close proximity, such as FLNG side-by-side offloading. In such configurations, nonlinear hydrodynamic interaction between vessels, low-frequency drift motions and mechanical coupling through moorings and fenders dominate system behaviour and govern operability limits. As described in DNV-RP-H103, these effects cannot be captured adequately using linear frequency-domain methods alone, necessitating fully coupled time-domain simulations to assess operability-limiting responses in multibody offshore systems [1]. A FLNG–LNGC side-by-side offloading arrangement is shown in Figure 1.1.

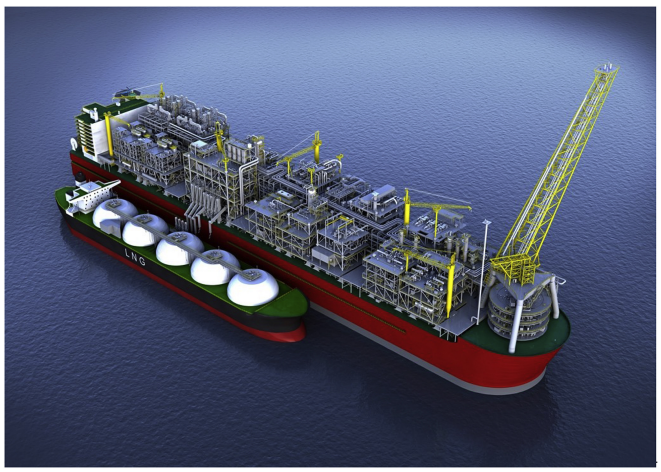


Figure 1.1: Rendering of an FLNG–LNGC side-by-side offloading arrangement [2].

Time-domain simulations explicitly resolve wave-frequency and low-frequency motions, nonlinear mooring and fender forces and transient interaction effects over extended simulation durations. However, this physical fidelity comes at significant computational cost. As a consequence, large operability assessments require extensive batches of long-duration time-domain simulations, typically 3 hours per environmental condition. For complete environmental datasets, this can result in up to $100{,}000$+ individual simulation runs [2].

In conventional operability assessment workflows, the detailed time histories generated by time-domain simulations are not used directly in the final decision-making process. Instead, each simulation is reduced to a small set of summary statistics, such as mean values, standard deviations and extreme maxima or minima, which are checked against predefined operability criteria. Once these summary metrics are obtained, the underlying time series are typically discarded, even though they are computationally expensive to generate. This observation raises two related questions. First, whether the full time-series data contain additional information that can be used more effectively, allowing the same insight to be obtained from fewer simulations. Second, whether operability assessment as a whole must rely so heavily on large numbers of time-domain simulations, or whether part of this dependence can be reduced by alternative or complementary modelling approaches while maintaining sufficient accuracy.

ML provides a potential framework to explore this question. By training ML models on high-fidelity time-domain simulation data, operability-limiting responses may be predicted efficiently while exploiting structure present in the response time series that is not retained in conventional summary-based assessments. However, a clear gap exists in the literature: there is no systematic comparison of summary-statistics-based and time-series-based ML models for a fully coupled multibody offshore system, and no end-to-end framework capable of reconstructing complete multibody response time series directly from environmental inputs.

This leads to the central research question of this thesis:

> *How can machine learning-based models support and accelerate operability assessment for side-by-side offloading?*

This thesis addresses this challenge by developing and evaluating ML-based models for operability assessment of an FLNG–LNGC side-by-side offloading configuration. High-fidelity time-domain simulation data are used to train surrogate models that predict key operability metrics, including vessel motions, relative manifold displacements, mooring line tensions and fender compression loads under varying environmental conditions.

Two classes of ML models are investigated. Summary-statistics-based models use static environmental descriptors and sea-state parameters to predict peak responses directly, offering simplicity and computational efficiency. In contrast, time-series-based models exploit the full three-hour response histories obtained from time-domain simulations and employ learning architectures capable of capturing temporal dynamics. Within this category, both full-length time-series learning and segment-based learning approaches are developed, where shorter response segments are used as training inputs. By training and evaluating these model types on the same dataset, the study systematically compares their predictive accuracy, robustness and potential to reduce reliance on extensive time-domain simulation effort.

Model performance is assessed on unseen sea states and benchmarked against reference time-domain simulations to evaluate the extent to which ML models can reproduce operability-limiting responses, reduce the computational burden of conventional operability assessment, and provide sufficiently accurate classification of operable and non-operable conditions.

The thesis also tries to take a helicopter view on how ML methods could be applied to operability assessment of multibody offshore operations. Rather than optimising a solution for one specific project, it uses an FLNG–LNGC side-by-side offloading configuration as a case study to systematically investigate what different ML paradigms can and cannot offer for this class of problems. The FLNG–LNGC configuration is representative of a broader class of coupled floating-body operations, including offshore hydrocarbon offloading, CO_2 ship-to-ship transfer, and floating wind installation, all of which share the same defining characteristics: multibody response-based operability limits and the dependence on costly time-domain simulations. The methods and findings of this thesis are therefore not specific to side-by-side offloading but applicable wherever these characteristics are present.

This thesis is structured as follows. Chapter 2 provides a general review of artificial intelligence and its applications within engineering disciplines, establishing the broader methodological context for this work. Chapter 3 presents a focused literature review on side-by-side offloading operability assessment and numerical simulation practice. Chapter 4 describes the FLNG–LNGC case study and the time-domain simulation dataset. Chapter 5 introduces the methodology of this thesis, while Chapter 6 details the implementation of these models. Chapters 7 and 8 present and analyse the model results. Chapter 9 discusses the findings in the context of operability assessment practice and simulation dependence. Finally, Chapter 10 summarises the conclusions and provides recommendations for future work.

2

General Review of Technology

As part of this thesis, a literature review was conducted to identify potential disruptive applications of Artificial Intelligence (AI) within Offshore Structures Engineering.

To do so, this chapter provides a general overview of AI, its applications within engineering, and its specific relevance to offshore structures engineering. The chapter begins by defining AI and outlining its broader use cases across engineering disciplines. It then discusses key bottlenecks currently faced in offshore structures engineering and explores how AI techniques may help address these challenges or introduce new opportunities for innovation.

Note: This chapter reflects the state of AI and its applications in September 2025. While it aims to give a comprehensive overview of AI developments and potential offshore engineering applications at that time, both AI and offshore engineering are rapidly evolving fields. Consequently, this chapter cannot capture every relevant AI method, emerging research direction, or industry challenge. Its purpose is to provide a structured foundation for understanding where AI may meaningfully contribute within the context of offshore structures engineering and establish the foundation for this thesis' direction.

2.1. What is AI?

AI is a multidisciplinary field focused on creating intelligent agents capable of executing tasks that would typically require human intelligence. It aims to mimic human cognitive abilities, enabling machines to perceive, analyze, learn, and make decisions. The term "Artificial Intelligence" was formally introduced at the Dartmouth Conference in 1956, where the idea of building machines capable of simulating human intelligence was proposed [3]. AI has evolved from simple automation to sophisticated technologies such as machine learning, deep learning, and natural language processing, driving innovation across industries like healthcare, finance, and transportation. As AI continues to advance, it holds the potential to transform how we live and work, presenting both incredible opportunities and complex challenges that shape our future.

To identify a promising application of AI within Offshore Structures Engineering, it is important to gain anunderstanding of what AI is, what its subfields are, and what these are capable of.

Although the classification of AI subfields is debated and the boundaries between them are often fluid, this section aims to present an overview that is as comprehensive and accurate as possible.

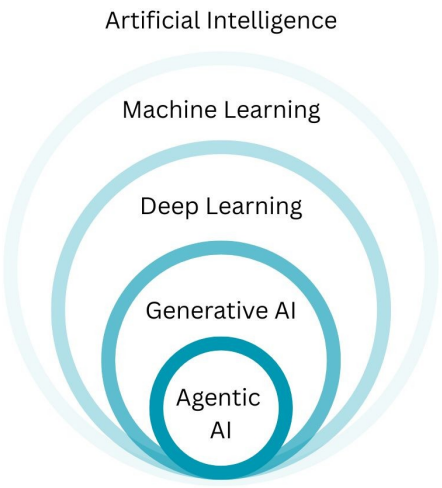


Figure 2.1: Hierarchy of AI subfields.

Figure 2.1 presents a layered conceptualization of AI, illustrating its progression from broad computational intelligence to highly autonomous systems. The layered circles represent increasing specialization and capability. At the outermost layer, Artificial Intelligence encompasses all systems designed to replicate or augment human cognitive functions. All AI systems are composed of three foundational elements: data, algorithms, and features.



Figure 2.2: Building blocks of an AI system.

Data: Data represents the raw input, whether from sensors, user interactions, or digital records, that the system uses to learn from.

Algorithms: Algorithms are the mathematical and logical procedures that process this data. They define how the system learns patterns, makes predictions, or identifies anomalies. An algorithm can be as simple as a decision tree or as complex as a deep neural network (explained later in this chapter). It acts as the engine that transforms input data into meaningful insights or actions.

Features: Features, in this broader context, refer to the capabilities or actions the AI system performs based on the information it processes. These include tasks such as classifying objects, generating recommendations, or triggering alerts. Features are what make the system useful. They are the visible outcomes of the algorithm's internal logic applied to data.

Together, these components enable AI systems to operate intelligently across different domains by turning raw data into actionable decisions.

Machine Learning (ML): Machine Learning is a branch of artificial intelligence focused on algorithms that can automatically discover patterns and relationships within data. ML models learn how input parameters relate to observed outputs and use these learned patterns to make predictions on new,

unseen data. As noted by Sarker (2021), machine learning enables computer systems to improve their predictive ability by extracting structure from data [4].

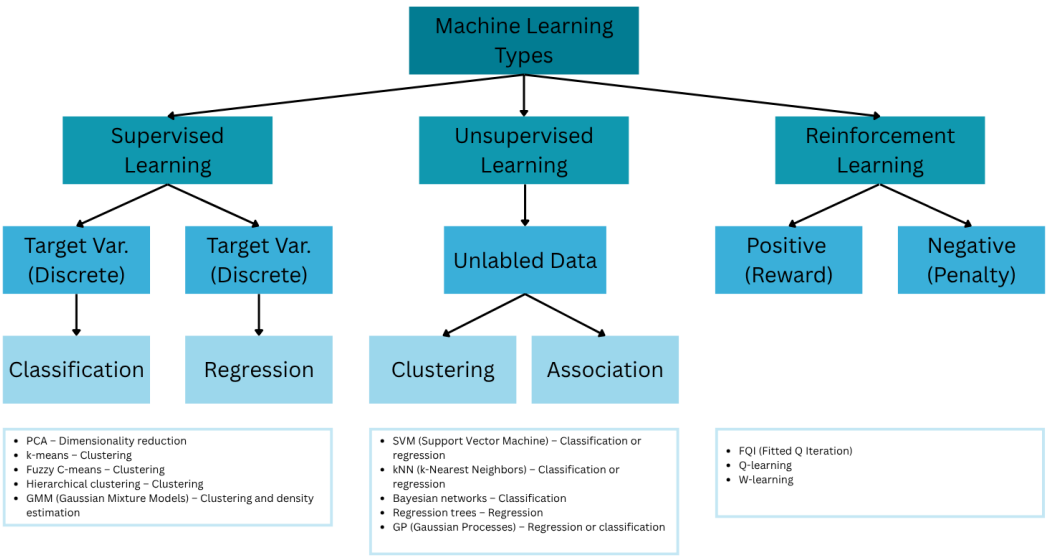


Figure 2.3: Machine Learning Overview (Based on [5] & [4])

ML algorithms are designed to identify patterns in data and use these patterns to make predictions or decisions. These algorithms can be broadly categorized into three main types, also depicted in Figure 2.3:

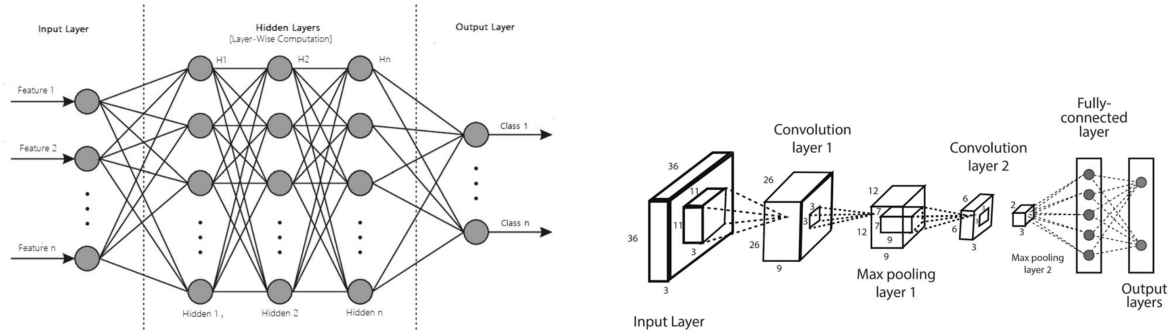
- **Supervised Learning:** Supervised learning involves training a model on a labeled dataset, where each input is paired with the correct output. The model learns to map inputs to outputs and is commonly used for tasks like classification and regression. Regression refers to the task of predicting continuous numerical values based on input features, such as estimating future stock prices from historical market data. Classification, on the other hand, involves assigning inputs to discrete categories, such as determining whether a stock's performance will be labeled as "rising" or "falling." It is particularly effective when the goal is well-defined and labeled data is available [6].
- **Unsupervised Learning:** Unsupervised learning deals with unlabeled data. The algorithm tries to find hidden patterns or intrinsic structures in the input data. Common applications include clustering, dimensionality reduction, and anomaly detection [6].
- **Reinforcement Learning:** Reinforcement learning is based on an agent interacting with an environment to learn optimal actions through trial and error. The agent receives rewards or penalties based on its actions and aims to maximize cumulative rewards. This approach is widely used in robotics, game playing, and autonomous systems [7].

The effectiveness of a machine learning model depends on both the quality of the data and the suitability of the learning algorithm. Selecting a proper learning algorithm that is suitable for the target application is challenging [4].

Deep Learning, a subset of machine learning, utilizes multi-layered neural networks to model complex, high-dimensional relationships, making it particularly effective for tasks such as image recognition, natural language processing and time series prediction. Deep learning is a powerful subset of machine learning that focuses on learning hierarchical representations of data through multiple layers

of abstraction. Unlike traditional models that rely heavily on feature engineering -the manual selection or construction of input variables from raw data- deep learning algorithms can process raw data and automatically discover the features necessary for tasks such as classification, regression and control [5].

Deep learning encompasses a variety of architectures, each suited to different data structures and learning tasks. Feedforward neural networks (FNNs) are the simplest form, where data flows in one direction through layers; they are effective for general function approximation [8]. In a simple neural network, as displayed in Figure 2.4a, each node (or neuron) in one layer is connected to every node in the next layer through weighted links. During training, the network adjusts these weights so that its output matches the desired target. This adjustment is done using ”backpropagation”, a procedure that calculates how much each weight contributed to the prediction error and updates the weights accordingly. In this way, the network gradually learns which connections are important for producing accurate outputs.



(a) Feedforward Neural network layout

(b) Convolution Neural Network

Figure 2.4: Schematic representations of neural network architectures [4].

Convolutional neural networks (CNNs) specialise in extracting patterns from grid-like data such as images or time-series [9]. They use small filters that slide across the input to detect local features, followed by pooling layers that reduce the data while keeping the most important information. Fully connected layers at the end combine these extracted features into a prediction, as illustrated in Figure 2.4b.

Recurrent neural networks (RNNs) are designed for sequential data and maintain an internal state that allows the model to use information from previous time steps [10]. This makes them suitable for signals where the current output depends on what happened earlier. Long Short-Term Memory (LSTM) networks [11] extend RNNs by adding gating mechanisms that control how information is stored, forgotten, or passed forward. These gates allow LSTMs to retain important information over much longer time spans, which standard RNNs struggle to do.

Deep learning models are powerful but come with challenges: they often require large datasets, significant computational resources, and their internal workings are difficult to interpret. Techniques such as dropout, data augmentation, and momentum-based optimisation are used to reduce overfitting and improve generalisation. Despite these limitations, deep learning is well suited for engineering applications where relationships between inputs and outputs are complex, nonlinear, and high-dimensional.

Moving further inward in the AI layers, **Generative AI** refers to models that learn the structure of data and can generate new, synthetic samples that resemble the training distribution. These models underpin much of the current progress in AI, enabling capabilities such as content generation, simulation

enhancement, and data augmentation [12]. Their use in offshore engineering is still emerging, but they offer promising opportunities for generating synthetic simulations or exploring design alternatives at scale.

At the core lies **Agentic AI**, which represents systems capable of autonomous, goal-directed behaviour. These systems combine perception, planning, memory, and decision-making, allowing them to act within physical or virtual environments [13, 14]. Agentic AI moves beyond passive prediction to active problem solving and is viewed as the next step toward autonomous engineering systems.

Table A.1 summarises the main AI methods across the different layers, together with their key capabilities and limitations.

2.2. State-of-the-Art Applications within Engineering

Recent advances in data-driven methods are reshaping scientific and industrial fields, enabling new approaches to modeling, optimization, and decision-making. In engineering, AI enables the automation of complex tasks, the discovery of patterns in high-dimensional data, and the development of predictive models that complement or even replace traditional physics-based approaches.

To define a meaningful and innovative research direction for this thesis, it is essential to consider developments in other engineering disciplines where AI has already demonstrated significant impact. Fields such as aerospace, civil, and mechanical engineering have embraced AI for tasks ranging from predictive maintenance and structural health monitoring to design optimization and autonomous control [15], [16].

AI Application	**Industry**	**Key Players**	**Specific Disruption**
Predictive Maintenance	Manufacturing	GE	Reduced downtime and maintenance costs by predicting equipment failures before they occur [17].
Autonomous Vehicles	Automotive & Aerospace	Tesla, Waymo	Revolutionized transportation with self-driving capabilities, increasing safety and efficiency [18].
Generative Design	Aerospace & Automotive	Airbus, GM	Optimized designs for lighter, stronger, and more efficient products, reducing material usage [19].
Energy Optimization in Data Centers	Technology/Data Centers	Google DeepMind	Improved energy efficiency by reducing power consumption in data centers through AI optimization [20].

Table 2.1: Examples of AI Applications and Their Industry Disruptions

Table 2.1 outlines key disruptions and successes of AI across engineering domains. It shows how AI has led to measurable improvements in areas like design, maintenance, data analysis, and automation. Reviewing these examples helps clarify the practical impact of AI and provides a foundation for understanding its potential in more specialized fields, such as offshore engineering. Three of those are discussed below.

Airbus is using Autodesk's generative design tools to create lighter, more efficient aircraft components, most famously the bionic partition—a cabin wall redesigned using algorithms that mimic natural cellular structures. The resulting component is 45% lighter while maintaining required strength. By using cloud-based computation to explore thousands of structural variations, Airbus demonstrates how generative design can unlock weight savings and sustainability gains.

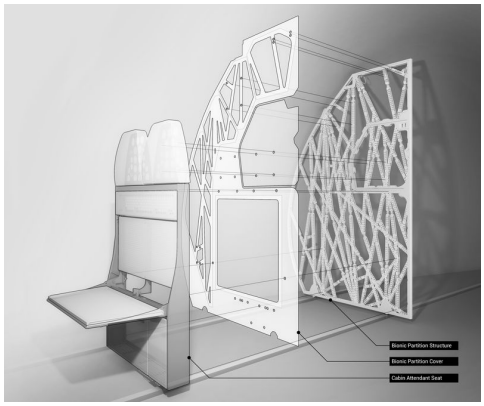


Figure 2.5: Airbus generative design [21]

GE Vernova uses its SmartSignal platform to deliver AI-based predictive maintenance, applying digital-

twin models to spot early anomalies in pumps, compressors, and other critical equipment. By comparing live sensor data with expected machine behavior, SmartSignal flags deviations—such as temperature spikes, pressure drops, or vibration changes—and issues diagnostics that indicate likely failure modes. This early warning capability allows operators to act well before breakdowns occur, reducing downtime and preventing costly failures [22].

The offshore industry also presents an opportunity for disruptive innovation through digitization and intelligent systems. The sector is characterized by complex, dynamic environments and high operational costs, making it an ideal candidate for AI-driven solutions that enhance safety, efficiency and decision-making.

By leveraging data from simulations, sensors, and operational logs, AI could support real-time monitoring, predictive modeling and optimization of offshore oil and gas operations, including drilling, production, and facility inspections [23].

The following areas represent the most mature and widely adopted applications of AI across engineering disciplines. These areas have advanced due to a combination of data availability, computational capability and clear performance benefits.

Predictive Maintenance & Strucutural Health Monitoring

- Used in: Mechanical, aerospace, civil and industrial engineering.
- Techniques: Time-series forecasting, anomaly detection, supervised learning.

AI-driven predictive maintenance systems analyze historical and real-time sensor data such as vibration, temperature, pressure, and acoustic signals. These inputs are used to detect patterns and anomalies that precede equipment failure. The goal is to predict when maintenance should be performed to prevent unexpected breakdowns, thereby reducing downtime, extending asset life and optimizing maintenance schedules [24].

Generative and AI-Assisted Design

- Used in: Aerospace, and product design engineering.
- Techniques: Evolutionary algorithms, topology optimization, neural networks.

AI-assisted design tools take design constraints, performance requirements, and material properties as input. They use optimization algorithms to explore a vast design space and generate innovative geometries that meet or exceed performance targets. The outcome is often a lightweight, structurally efficient design that would be difficult to achieve through manual methods [12].

Smart Manufacturing and Process Automation

- Used in: Manufacturing, systems and industrial engineering.
- Techniques: Computer vision, reinforcement learning, digital twins.

In smart manufacturing, AI systems process data from cameras, sensors, and control systems to monitor and optimize production lines. Inputs include visual inspection data, machine parameters, and environmental conditions. AI enables adaptive control, fault detection, and autonomous decision-making, resulting in higher product quality, reduced waste, and improved operational efficiency[25].

Digital Twins and Real-Time Simulation

- Used in: Civil, mechanical, electrical, and systems engineering.
- Techniques: Physics-informed machine learning, hybrid modeling, data assimilation.

Digital twins are virtual replicas of physical systems that continuously receive data from sensors embedded in the real system. Inputs include structural loads, environmental conditions, and operational parameters. AI models simulate system behavior in real time, enabling predictive analytics, performance monitoring, and scenario testing. This supports better decision-making throughout the lifecycle of the asset [26].

Materials Discovery and Optimization

- Used in: Materials science, chemical, and mechanical engineering.
- Techniques: Deep learning, graph neural networks, Bayesian optimization.

AI accelerates materials discovery by learning from large databases of known crystal structures and first-principles simulation results. Inputs typically include chemical compositions and crystal structures, which AI models use to predict properties such as formation energy, stability, or ionic transport behavior. These predictions guide researchers toward promising candidate materials before committing to expensive simulations or laboratory synthesis. The outcome is a dramatically faster identification of new, stable materials with desirable characteristics [27].

2.3. Bottlenecks in Offshore Structures Engineering

To find potential for disruptive implementations of AI within offshore structures engineering it can be helpful to look at the full process or lifecycle within this subject, to define where the bottlenecks are or where improvements would be necessary.

Design Phase

The design phase in offshore structures engineering is characterized by several bottlenecks that hinder efficiency and innovation. One major challenge is the reliance on conservative design codes, which prioritize safety but often lead to over-engineered solutions and limited flexibility in concept development. These codes, originally developed for oil and gas platforms, are not always well-suited for emerging applications such as offshore renewables or hydrogen infrastructure [28].

Another bottleneck is the non-integrated nature of design tools and workflows. Engineers often rely on a mix of software packages for structural, hydrodynamic, and mooring analyses, which lack seamless interoperability. This disconnection adds complexity to the design process and slows down iteration cycles. Moreover, early-stage design typically involves manual trial-and-error approaches, which are both time-consuming and susceptible to bias.

Amiri et al. (2021) highlight that offshore platforms face unique design challenges due to harsh environmental conditions, dynamic loading, and the need for long-term reliability. These factors demand robust modeling approaches, yet current practices often fall short in capturing nonlinear interactions and uncertainty. The lack of integrated, data-driven design frameworks further limits the ability to optimize structures holistically [29].

Installation & Operation

Current operability assessments in marine operations often rely on empirical sea state thresholds without systematic links to structural responses, leading to unclear and potentially conservative limits. They typically focus on isolated phases, neglecting operational continuity, and lack consistent treatment of uncertainties such as forecast errors and modeling limitations. Numerical tools also struggle with nonlinear, non-stationary dynamics, reducing reliability in complex offshore conditions [30].

With regard to the operability assessment of FLNG side-by-side offloading, Zhao et al. (2018) [31] highlight the challenges faced by operators due to the lack of standardized guidelines. Operability is affected by complex interactions between environmental conditions such as wave height, swell direction, and wind, and vessel-specific characteristics. These factors contribute to nonlinear behaviors, particularly in roll motion, sloshing, and gap resonance, which are not adequately captured by conventional modeling tools. The authors advocate for improved models and data to support more reliable, physics-informed assessments tailored to operational realities.

The Prelude FLNG study shows that traditional numerical simulations can become a bottleneck. To understand side-by-side offloading, the team had to run more than 100,000 time-domain simulations, because there were so many variables: different ships, loading conditions, mooring setups, headings, and sea states [2]. The authors also explain that the problem was so complex that it required multiple research projects and new modelling techniques just to handle all the combinations. This clearly shows how resource-heavy operability assessments can be. It also highlights why data-driven methods are needed; they can learn from a smaller number of simulations and avoid this computational workload.

Maintenence & Monitoring

A persistent bottleneck in offshore engineering is the continued operation of legacy platforms well beyond their 20–30 year design lives. Over time, corrosion wastage, fatigue cracking, hydrogen□assisted

embrittlement and environmental cracking interact and progressively reduce residual strength and toughness. The problem is made worse by the fact that many older platforms were built with materials and welds that do not meet today's toughness and fatigue standards. Limited inspection coverage in inaccessible areas and incomplete historical records further increase uncertainty in life□extension decisions and required safety margins [32].

Moan (2018) published a discussion of structural bottlenecks in offshore integrity management. A primary constraint is the continued reliance on generic inspection regimes, which often do not reflect the actual uncertainty or degradation patterns of individual offshore structures. This mismatch leads to insufficient coverage of critical areas where deterioration is most likely to occur. Fatigue in welded joints and mooring components represents another key bottleneck: small fabrication-induced defects commonly escape early detection, yet can propagate rapidly under cyclic loading, outpacing the assumptions embedded in conventional SN-curve-based design procedures [33].

A further limitation is the difficulty of incorporating real uncertainty into maintenance planning, including uncertainty related to fabrication variability, inspection limitations, environmental loading, and human or organizational error. Because these uncertainties are typically addressed through conservative assumptions rather than continuous updating with in-service observations, inspection and repair strategies remain largely reactive rather than adaptive. This insufficient integration of operational data restricts the shift from traditional maintenance management to true integrity management, thereby constraining the industry's ability to anticipate fatigue progression and allocate resources effectively [33].

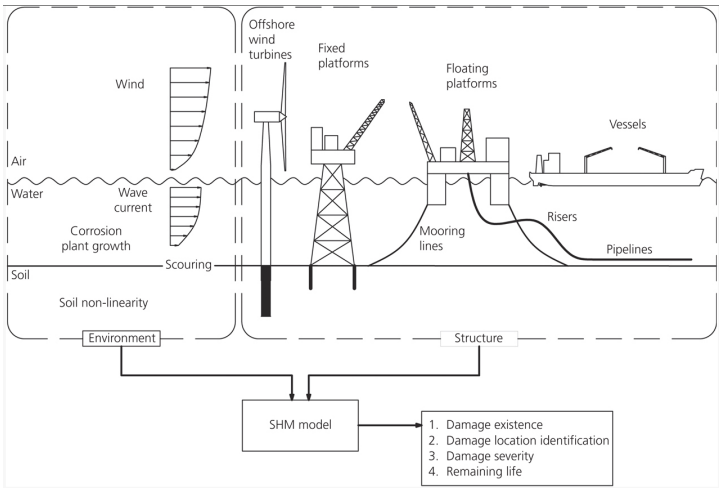


Figure 2.6: Structural Health Montioring Overview [34]

When it comes to Structural Health Monitoring (SHM) systems, a major challenge is the environmental uncertainties that offshore structures are subjected to, such as unpredictable wind, waves, and marine currents. These stochastic forces complicate the accurate modeling and prediction of structural behavior, especially in fatigue life assessment [34]. Finally, real-time monitoring in such dynamic environments is another major hurdle. The complexity of monitoring the continuously changing offshore conditions and ensuring that real-time data is accurately captured and analyzed requires sophisticated and often computationally expensive methods, such as digital twins, to predict long-term structural performance. An overview of such a sytstem of offshore structures is displayed in Figure 2.6.

Decommissioning

A large wave of decommissioning is happening: about 2,000 offshore oil and gas platforms are expected to be retired by 2040, with especially high activity in the North Sea and the U.S. Gulf of Mexico. Rules and permits can be outdated or inconsistent across countries. Engineering is hard because many platforms are old, corroded, and difficult to cut and lift safely. Logistics are tight, with limited heavy-lift vessels, long sailing distances, and narrow weather windows. Costs are uncertain because detailed cost data are rarely shared and estimates are often optimistic. Environmental knowledge still has gaps, for example on long-term ecological effects, invasive species, and how to count emissions from decommissioning work. Stakeholders such as fishers, conservation groups and operators often have conflicting goals. Finally, health, safety, and environmental (HSE) risks remain high during lifts, cutting, diving, and waste handling and must be managed carefully.

Decommissioning offshore platforms involves complex environmental and regulatory challenges. Complete removal often leads to significant loss of marine biodiversity, especially in areas lacking natural reef habitats [35].Despite the ecological benefits of partial removal or reefing, such options are often limited by regulatory constraints and inconsistent environmental assessment frameworks across regions [36]. Removal operations also pose environmental risks, including contaminant dispersal and high energy and carbon costs [37]. A multidisciplinary, ecosystem□based approach is therefore recommended, with opportunities to repurpose infrastructure for artificial reefs, renewable energy, or research platforms [35, 38].

2.4. AI in Offshore Structures Engineering

Based on the previous sections, three promising research directions are defined, which will be discussed below. These directions aim to address bottlenecks in offshore structural engineering, such as prolonged design cycles, limited operability assessments and high maintenance costs. The following sections will explore how emerging technologies can enhance efficiency and reduce conservatism in offshore infrastructure and industry.

2.4.1. Generative Design of Offshore Support Structures

The application of generative design in the context of offshore support structures represents a promising advancement in structural engineering. Generative design refers to computational methods that autonomously explore a wide range of design alternatives based on predefined performance criteria and constraints. In offshore engineering, this approach offers the potential to overcome traditional bottlenecks such as conservative structural sizing, limited geometric exploration, and time-intensive manual iterations.

As defined by Shea et al. [39], generative design aims to produce geometrically novel, efficient, and buildable structures by leveraging modern computational and manufacturing capabilities. Singh and Gu [40] emphasize its role in expanding the design space, enabling the discovery of innovative configurations that may not emerge through conventional design processes. Krish [41] further characterizes generative design as a designer-driven, parametrically constrained exploration process, typically implemented within parametric CAD environments.

In the offshore domain, generative design has yet to be widely adopted, but early efforts suggest its transformative potential. Karimi et al. [42] developed a multi-objective optimization framework for floating offshore wind turbine support structures, balancing cost and performance using a frequency-domain model. Their parametric study identified four-float semi-submersibles and TLPs as optimal, expanding the design space for floating offshore wind turbines (FOWTs), though without employing generative design methods.

Jang et al. [43] developed a reinforcement-learning approach for structural design that uses a neural network to predict design performance and an reinforcement learning agent to explore new design options. Applied to automotive wheels, their method quickly generated a wide range of feasible and realistic designs, showing how reinforcement learning can support automated and diverse structural design exploration.

Shen et al. [44] introduced a metamodel-based generative design framework for onshore wind turbine foundations, employing a multi-output Random Forest model optimized via NSGA-II. Their approach achieved a reduction in computational time, highlighting the feasibility of surrogate-assisted design exploration in offshore contexts.

Saeed et al. (2024) [45] applied a density-based topology optimization approach to achieve a lightweight yet structurally robust design under extreme environmental conditions. The authors further translated the TO results into a manufacturable truss-based configuration and proposed reinforcement strategies to mitigate local stress concentrations.

Sun and Ma [46] showed that combining reinforcement learning with topology optimisation can help generate new structural designs. They tested several ways for the algorithm to explore design options. They were able to produce a good variety of usable 2D and 3D designs.

Although topology optimization has seen increasing use in offshore structural design, generative de-

sign remains underexplored in this domain. Compared to other fields, offshore engineering has yet to fully leverage autonomous design generation and diversity-driven exploration. Integrating generative design, particularly reinforcement learning and multi-objective optimization, offers opportunities to unlock unseen, high-performance support structures. Future progress lies in developing offshore-specific generative frameworks that balance manufacturability, environmental resilience and computational efficiency.

2.4.1.1. Generative Design Workflow

Generative design for offshore support structures leverages machine learning to autonomously create new configurations that meet structural and environmental requirements. As illustrated in Figure 2.7, the process begins with historical design data, such as geometry and performance metrics, which is used to train a generative model. A common approach is the Variational Autoencoder (VAE), where an encoder maps designs into a latent space and a decoder reconstructs feasible variants from sampled points.

Alternatively, Generative Adversarial Networks (GANs) can be used to create new design options. In a GAN, one network generates designs and another checks whether they look realistic compared to known examples. The generator starts from simple random inputs, which act like random "seeds" that influence the shape of the design. Once trained, the generator can produce new support structure designs, which are then converted into geometric form and evaluated through engineering analysis.

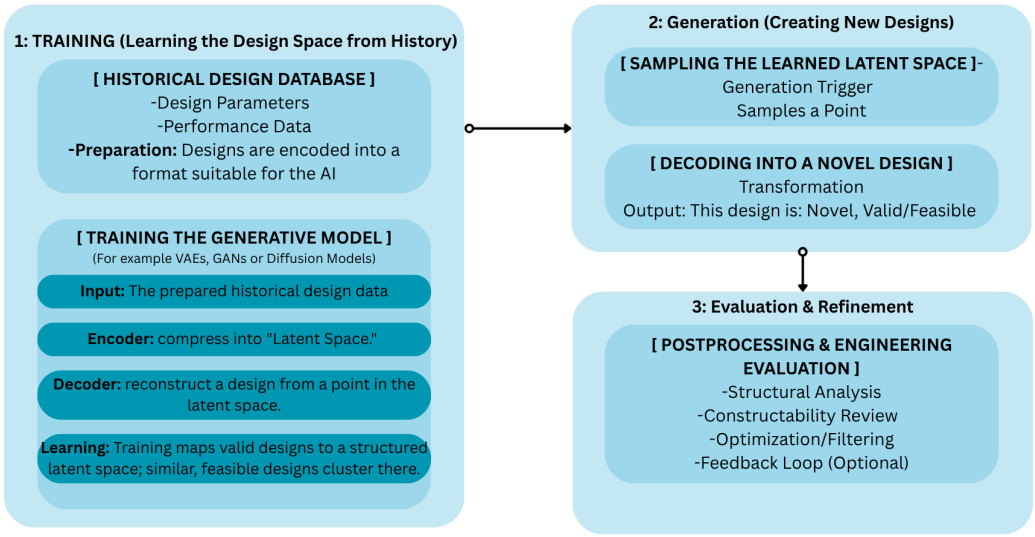


Figure 2.7: Generative design workflow

Table 2.2 summarizes the key aspects of generative design for offshore support structures. It outlines the purpose and functioning of generative algorithms, such as GANs and VAEs, in automating the creation of novel structural configurations. The table also highlights key input requirements, including environmental loads, material properties, and fabrication constraints. It emphasizes how these methods can contribute to improving design efficiency and encouraging innovation. Importantly, it draws attention to critical concerns such as physical feasibility, regulatory compliance, and the need for high-quality training data, which must be addressed to ensure successful application in offshore engineering.

Table 2.2: Overview of Generative Design for Offshore Support Structures

Topic	**Description**
Purpose	Automate the exploration of structural configurations and generate new foundation structures using generative algorithms.
How it Works	Learns patterns from existing data using models such as GANs and VAEs, then samples or creates new variations that follow those patterns, even if they were never seen before.
Inputs Required	Design domain, environmental loads (e.g., wave, wind, current), material properties, boundary conditions, fabrication constraints and a dataset of historical designs.
Potential Impact	Enables discovery of innovative support structures, reduces overdesign, improves structural efficiency, shortens design cycles and supports disruptive innovation in offshore design methodologies.
Key Concerns	Ensuring physical feasibility, interpretability of AI-generated designs and availability of sufficient high-quality training data.

2.4.2. Operability Assessment of Complex Marine Operations

Operability assessment is a critical component in the planning and execution of marine operations, particularly for complex offshore operations. It refers to the evaluation of the available time windows during which operations can be safely performed, based on environmental conditions and system limitations. Accurate operability predictions are essential for minimizing downtime, ensuring safety and optimizing logistics.

Current industry standards, such as those by Det Norske Veritas (DNV, 2011) [47] recommend defining operational limits in terms of environmental parameters like significant wave height (H_s), wind speed, and vessel responses. However, these limits are often based on experiences or simplified assumptions, lacking systematic derivation from structural response analyses [48].

Several studies have addressed this gap. Clauss and Riekert [49, 50] proposed motion-based criteria for floating crane operations. De Wilde et al. [51] contribute to operability analysis by demonstrating a direct time-domain approach for LNG offloading using high-performance computing. Their method enables long-term simulation of full operational sequences under various environmental conditions, offering detailed insight into downtime events and providing a valuable benchmark for faster surrogate models.

More recently, Guachamin Acero et al. [30] and Li et al. [52] introduced systematic approaches to identify critical events and derive response-based operational limits using time-domain simulations.

Despite recent progress, a fully integrated methodology that connects operational limits with weather window analysis and fast data-driven decision-making is still not fully established.

AI could offer potential improvements by enabling automatic pattern recognition from simulation or sensor data, enhancing short-term sea state forecast integration and supporting operability assessment by reducing the dependence on numerical simulations.

Recent developments in machine learning have advanced the prediction of nonlinear ship motions. RNNs and LSTM networks are widely used in time series modeling because they can learn patterns over time and handle complex, nonlinear behavior. LSTM models have been successfully applied to ship motion prediction tasks, such as estimating deterministic responses under wave excitation [53] and inferring wave direction from onboard measurements [54]. Guo et al. [55] used LSTM networks trained on wave basin data to forecast heave and surge motions of a semi-submersible, based on wave and motion time series, including noisy and motion-only input scenarios.

A key opportunity lies in developing machine learning models that reduce the need for repeated full-scale simulations by generalizing across scenarios and offering deeper insight into the governing dynamics.

2.4.2.1. Surrogate Modelling Workflow

To reduce reliance on computationally expensive time-domain simulations, deep learning models can be employed as surrogate models. These models learn from simulation-derived time series data to approximate vessel and mooring responses under varying environmental conditions, enabling efficient operability prediction.

Figure 2.8 illustrates the surrogate-modelling workflow. The process begins with input from numerical simulations or field measurements, including vessel motions, mooring loads, and environmental parameters. These time series are preprocessed and passed to a neural network with an input layer, several hidden layers, and an output layer. The model learns the temporal and nonlinear relationships needed to predict operability-limiting metrics such as excessive motions or line tensions.

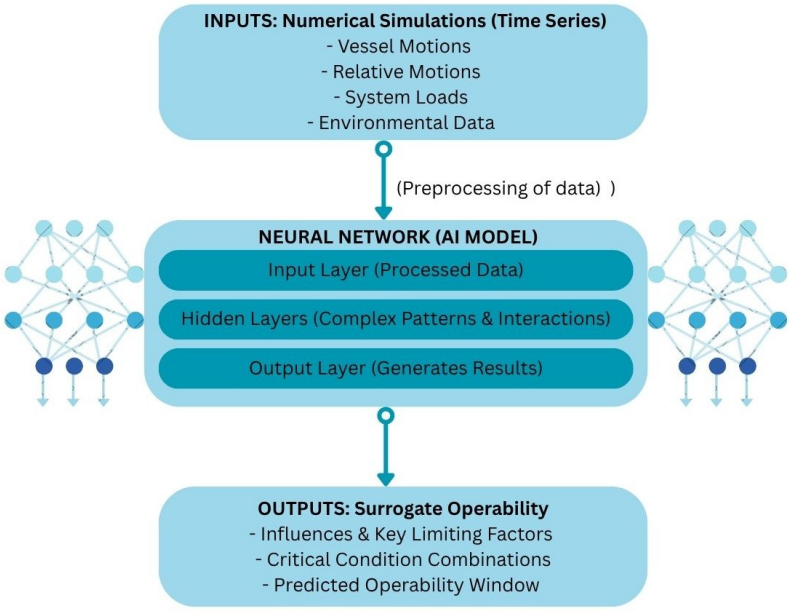


Figure 2.8: Workflow of a neural network-based surrogate model for operability assessment.

The model provides the following outputs:

- Predicted operability windows over the planning horizon
- Combinations of conditions that are most critical
- The key factors that limit operations
- An adaptive assessmen of feasibility

In this context, adaptive assessment means that the evaluation updates itself as new data become available. Rather than relying on fixed thresholds alone, the model revises its operability estimates when it encounters changing sea states, vessel responses, or updated measurements. This enables a more accurate and situation-specific judgement of whether offshore operations can proceed safely.

Together, these outputs support an efficient and scalable assessment of complex offshore operations, while allowing the decision process to improve as more data are collected.

Table 2.3: Overview of AI-Enhanced Operability Assessment for Offshore Operations

Topic	Description
Purpose	Assess safe operational windows for offshore activities based on environmental and system constraints, aiming to reduce downtime and improve safety.
Approach	Use time-domain simulations and AI models (e.g., LSTM networks) trained on time series data to predict operability under varying conditions.
Inputs	Vessel motions, relative motions, mooring loads, and environmental parameters from simulations or sensor data.
Impact	Reduce reliance on extensive numerical simulations during planning phases and enable fast, data-driven operability assessments to support decision-making for offshore operations.
Challenges	Ensuring physical realism, interpretability of AI outputs, and availability of high-quality training data.

2.4.3. AI-Driven Predictive Maintenance and Digital Twin Systems for Offshore Platforms

Offshore platforms operate in harsh environments where structural degradation and operational risks are constant concerns. Traditional maintenance strategies often fall short in addressing these challenges efficiently. As offshore assets age and complexity increases, there is a growing need for intelligent systems that can monitor conditions in real time and predict failures before they occur. Integrating digital twins with AI-driven predictive maintenance offers a proactive approach to improving safety, reliability and lifecycle performance.

Offshore platforms are complex and high-risk structures that require strong maintenance strategies to ensure safety, reliability and cost-efficiency. Traditional approaches such as reactive and condition-based maintenance are increasingly being replaced by predictive maintenance, which uses real-time data to anticipate failures before they happen [56].

Reactive maintenance can compromise safety, while preventive strategies may lead to unnecessary servicing. To optimize both safety and efficiency, maintenance planning must be informed by accurate, data-driven condition assessments.

A promising advancement in this domain is the integration of digital twins with AI-driven structural health monitoring (SHM) systems. A digital twin is a data-integrated model that mirrors the real-time structural state of a physical asset, including configuration, material properties, and degradation [57, 58]. By continuously updating the model with sensor data, digital twins reduce modeling assumptions and uncertainty, enabling more accurate diagnostics and prognostics.

Sha et al. [59] proposed a digital twin system tailored to offshore platforms. Their framework integrates modules for physical entity monitoring, digital modeling, predictive analytics, and visualization. It supports damage detection, residual life estimation, and maintenance planning, offering a scalable solution for intelligent predictive maintenance.

Carvalho et al. [60] conducted a systematic review of machine learning methods applied to predictive maintenance. Their work highlights the importance of selecting appropriate algorithms and preprocessing techniques to ensure reliable diagnostics and prognostics in industrial settings.

Chen et al. [61] emphasized the role of digital twins in risk-based inspection planning. Their study demonstrated how real-time condition data can optimize inspection schedules and resource allocation, improving both safety and cost-efficiency.

Semeraro et al. [62] provided a broad literature review on the digital twin paradigm, identifying key challenges in system integration, data synchronization, and lifecycle modeling. Their work underscores the need for standardized frameworks and cross-disciplinary collaboration in marine engineering applications.

A review by Imran et al. [63] highlights how mature and feasible modern corrosion-detection approaches have become, showing that deep-learning models consistently achieve accuracies above 90-95% across ship hulls, offshore platforms, and subsea pipelines, often surpassing traditional inspection methods. This wide range of studies demonstrates not only that automated computer-vision systems are already technically viable, but also that the field has progressed far enough to support integration into digital-twin and predictive-maintenance frameworks for real-world offshore applications.

In line with these developments, Sindi et al. [64] introduced a practical digital healthcare engineering framework for aging offshore structures, proposing a streamlined workflow for smart monitoring and digital-twin–driven maintenance.

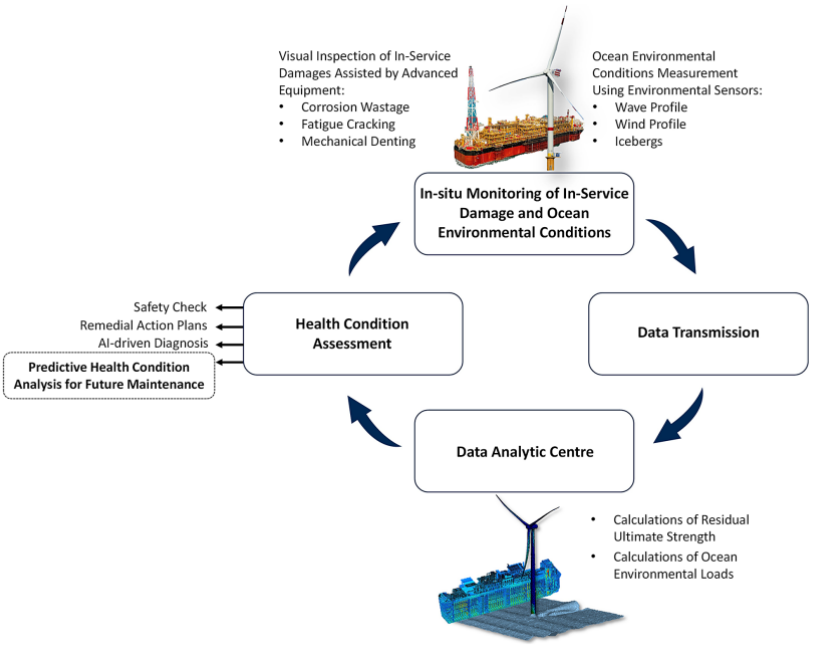


Figure 2.9: Schematic representation of digital twin framework [64].

As illustrated in Figure 2.9, their system combines periodic visual inspection of in-service damage with real-time environmental sensing, followed by satellite-based data transmission to a central analytics hub. There, CFD and nonlinear FEM are used to compute environmental loads and residual strength, enabling continuous health assessment, AI-based diagnosis, and predictive maintenance planning. This modular architecture reinforces the value of integrating sensing, simulation, and machine learning into a unified digital twin to support reliable, data-driven asset management.

Although existing studies contribute important building blocks—such as digital-twin architectures [59], machine-learning methods for predictive maintenance [60], risk-based inspection planning [61], and automated corrosion detection [63]—there remain clear gaps in how these components are integrated into a unified, offshore-specific maintenance ecosystem. The following research directions address these gaps and extend beyond current work:

- **Modular system integration:** Most existing frameworks focus on isolated elements (e.g., sensing, modelling, analytics). A coherent modular architecture that links monitoring, simulation, and decision-making remains underdeveloped. Research should define how these modules interact in practice and how new sensors or models can be added seamlessly.
- **Hybrid physics–ML modelling:** Prior work either emphasises physics-based modelling or machine learning, but few approaches combine them systematically. A hybrid methodology would improve accuracy by using physics as a baseline while allowing ML to capture residual behaviour and degradation patterns learned from data.
- **Real-time decision support:** While risk-based inspection is well studied, real-time adaptive maintenance planning, updated continuously as new condition data arrive is still limited. Developing such adaptive decision tools would directly support offshore operational scheduling and resource allocation.
- **Lifecycle prognosis under actual operating conditions:** Existing studies treat fatigue, corrosion, and degradation models largely in isolation. A consolidated lifecycle module that feeds real

motion and load data into damage models would provide more realistic estimates of remaining useful life.

Together, these directions extend beyond the current literature by focusing not only on individual technologies but on their coordinated integration into a practical digital twin workflow tailored for offshore platforms. By continuously learning from sensor data and updating structural models in real time, digital twins enable proactive decision-making, reduce uncertainty, and extend asset life. Their ability to simulate future scenarios and optimize inspection schedules transforms maintenance from a reactive task into a strategic, data-driven process. As research advances, these systems hold the potential to redefine offshore asset management, making it safer, more efficient and even more economically sustainable.

2.5. Conclusion

This chapter has reviewed the current state of artificial intelligence (AI) and its transformative potential within offshore structures engineering.

Across engineering, AI has already proven its value in real industrial settings. Companies such as Airbus use generative design to create lighter components, cutting material use and improving efficiency. GE applies AI-based predictive maintenance to detect equipment issues early and avoid costly downtime. Similar successes in sectors like automotive, aerospace, and manufacturing show how AI can improve safety, reduce costs, and support smarter decision-making. These examples demonstrate that AI is not only technically feasible but already delivering practical benefits—highlighting its strong potential for offshore engineering as well

First, generative design provides a route toward automated, data-driven structural innovation. By using reinforcement learning, surrogate models and generative algorithms, engineers can explore a broader design space with far fewer manual iterations, offering the potential for lighter and more efficient offshore support structures. Second, AI-based operability assessment reduces dependence on computationally expensive time-domain simulations. Deep learning models trained on environmental and response time-series data can approximate vessel and mooring behavior rapidly, enabling faster identification of safe weather windows and more adaptive operational planning. Third, predictive maintenance and digital twins are emerging as powerful tools for lifecycle asset management. By combining sensor data or camera data, environmental modeling and machine-learning diagnostics, digital twins can improve inspection planning, support early fault detection and reduce maintenance uncertainty.

In this thesis, the focus is placed on applying machine-learning methods to operability assessment, with side-by-side offloading as the central case study.

A unique opportunity for a machine learning-based operability assessment study is provided by an available simulation campaign conducted by MARIN for Shell. In this study the operation side-by-side offloading is simulated presenting a real-world scenario with significant operational complexity. MARIN has conducted extensive numerical simulations for this case, resulting in a dataset of time series responses under varied sea states. This dataset forms the foundation for developing and testing machine learning models that can predict operability outcomes and support the next generation of data-driven assessment tools.

In summary, the integration of AI-driven surrogate models, supported by high-quality simulation data from real-world case studies, represents a promising and innovative direction for advancing operability assessment in offshore engineering.

3

In Depth Review of Technology

After defining the topic of this thesis in the previous chapter, namely the use of surrogate modelling for the operability assessment of side-by-side (SBS) offloading, this chapter presents a focused literature review. While the case study later in this thesis concerns LNG transfer from an FLNG to an LNGC, SBS offloading is a multi-body problem that also appears in similar form in emerging CO_2 ship-to-ship and ship-to-platform transfer chains, as well as floating-to-floating wind installation [65]. CO_2 shipping exhibits the same essential multi-body characteristics because offshore offloading typically requires a CO_2 carrier to interact dynamically with a buoy or platform, creating coupled motions and operability constraints [66]. Many of the hydrodynamic mechanisms, operability limits, and modelling challenges reviewed here therefore apply beyond LNG and FLNG-specific configurations, meaning that the methods and insights developed in this thesis are relevant to a broader class of offshore multi-body operations.

As Floating Liquefied Natural Gas (FLNG) side-by-side (SBS) forms the case study on which this thesis is focused, this chapter gives an overview of the operational principles of FLNG SBS offloading, including the different mooring configurations and offloading methods used in industry practice. It then examines spread-moored SBS offloading in more detail and discusses the implications of this arrangement for vessel motions, mooring loads, and overall operability. Furthermore, the chapter reviews both conventional operability-assessment approaches and recent developments in machine learning-based prediction of ship motions, mooring tensions, and operable conditions. Finally, the chapter identifies the remaining literature gap and formulates the research question addressed in this thesis.

FLNG facilities are offshore units that integrate gas liquefaction, storage, and offloading systems into a single floating structure. By bringing part of the LNG value chain offshore, FLNG technology enables the development of remote or deep-water gas fields without reliance on onshore liquefaction plants, long subsea pipelines, or coastal export terminals. A key operational activity of these facilities is the transfer of LNG to conventional carriers.

In this thesis, the focus lies specifically on SBS offloading of a spread-moored FLNG. This configuration introduces distinctive hydrodynamic and operability challenges due to the close proximity of the two large vessels and the environmental forces acting upon them. Because the feasibility and safety of SBS operations depend strongly on the station-keeping performance of the FLNG, various mooring arrangements are briefly discussed. The mooring system determines the unit's ability to maintain

heading, limit low-frequency motions, and manage wind, wave, and current loads during offloading. Understanding the interaction between mooring configuration, vessel motions, and offloading operability is therefore fundamental to assessing the performance and limitations of SBS operations.

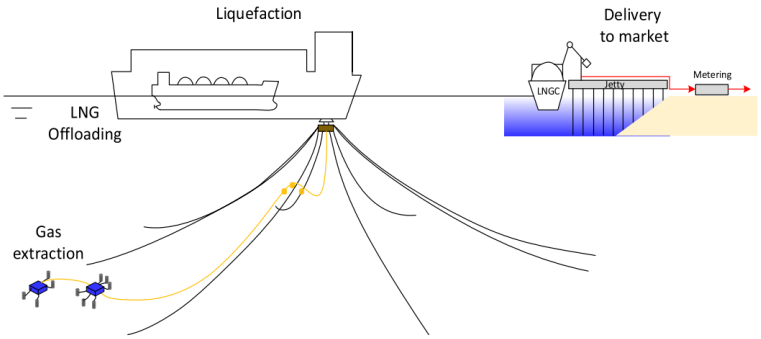


Figure 3.1: Schematic representation of an FLNG facility showing gas extraction, turret mooring, liquefaction, and LNG offloading [31].

Figure 3.1 illustrates a typical turret-moored FLNG facility that extracts and liquefies natural gas before exporting LNG to a carrier.

3.1. Body Motions

Floating bodies respond to waves, wind, and current through six degrees of freedom (6-DOF). These motions describe how a floating structure translates and rotates in space, and they form the basis for analysing hydrodynamic behaviour in both single- and multi-body systems. For side-by-side offloading, understanding these motions is essential, as the FLNG and LNG carrier move independently while also influencing each other through hydrodynamic interaction, fender contact, and mooring-line forces.

Table 3.1: Overview of vessel motions in six degrees of freedom (6-DOF).

Category	Motion	Symbol	Description
Translational	Surge	ξ_1	Fore–aft displacement along the longitudinal (X) axis.
	Sway	ξ_2	Side-to-side displacement along the transverse (Y) axis.
	Heave	ξ_3	Vertical displacement along the vertical (Z) axis.
Rotational	Roll	ξ_4	Rotation about the longitudinal (X) axis.
	Pitch	ξ_5	Rotation about the transverse (Y) axis.
	Yaw	ξ_6	Rotation about the vertical (Z) axis.

The motions within a multi-body system are displayed in Figure 3.2.

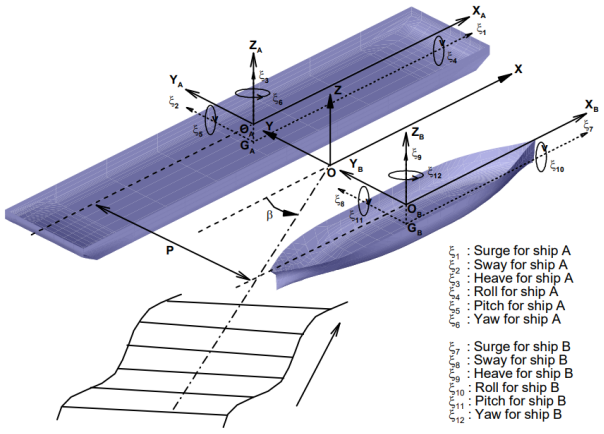


Figure 3.2: Schematic overview of the multi-body motion system. Each vessel responds in 6 DOFs, producing a coupled 12-DOF [67].

3.2. Mooring Systems for FLNG Facilities

For the purpose of assessing the implications of a chosen mooring system, three common types of FLNG mooring systems are discussed below as reference.

Turret Mooring Systems: Turret moorded FLNG can freely weathervane under the action of waves, wind, and current. This capability reduces environmental loading and allows the FLNG to naturally align with the sea state. Such freedom of rotation improves the hydrodynamic behavior during offloading, especially in swell-dominated environments [31]. Turret mooring is a defining feature of major FLNG developments, including Shell's Prelude FLNG [2].

Studies by MARIN for Shell demonstrate that turret-moored vessels exhibit notably reduced sway and roll responses when aligned with incident waves, improving relative motion performance for SBS configurations. Furthermore, fender loads and line tensions remain within acceptable operational limits across a broader range of conditions when compared with fixed-heading systems [68].

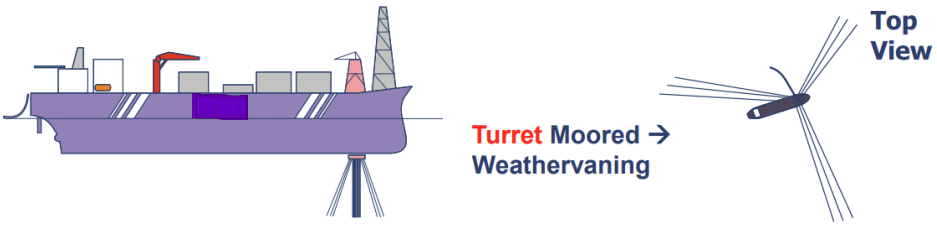


Figure 3.3: A turret-mooring arrangement example [69]

Spread Mooring Systems: In a spread mooring system, the FLNG is held in a fixed heading through multiple mooring lines anchored to the seabed. While spread mooring provides high positional stability, the inability to weathervane can expose the vessel to beam-sea conditions. Such conditions increase vessel motions and hydrodynamic loads during offloading and may significantly reduce operability, particularly for SBS arrangements that are sensitive to roll and lateral motions [2].

Simulation results for spread-mooring show that spread-moored floating units experience substantially

higher roll motions in beam seas, more often exceeding typical LNG-loading-arm limits of approximately $\pm 3^{\circ}$ [70]. Natural periods of surge and sway for spread-moored systems are typically above 100 seconds due to the high static stiffness of shallow-water mooring spreads, placing them outside wave-frequency excitation bands. However, because heading is fixed, beam-sea excitation produces large wave-frequency roll and sway amplitudes that dominate operability. Additionally, spread-moored systems have limited allowable lateral offsets before mooring lines become taut, especially in shallow water, which poses a risk of overstressing the anchoring mooring system [70].

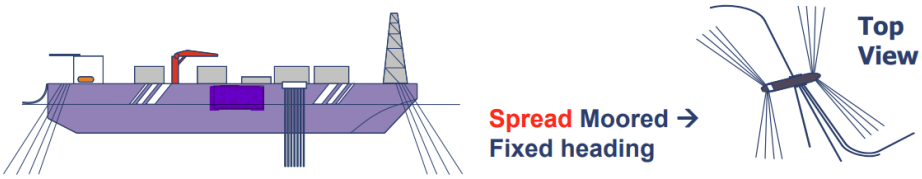


Figure 3.4: A spread-mooring arrangement example [70]

Jetty-Moored FLNG: In a jetty-moored arrangement, the FLNG is restrained by conventional shore (or dolphin) moorings and fenders. Heading is effectively fixed by the berth axis, so the system behaves similarly to a spread-moored floating unit from a heading-control standpoint, but line geometry and fender characteristics are tuned by jetty layout. Simulations show that in bow/quartering seas operability is generally acceptable, while beam-sea exposure increases lateral offsets, roll, and fender compression; nonetheless, with appropriate breast and spring line layouts and suitably sized cell or pneumatic fenders, line tensions remain within limits for many operational cases considered [68, 70]. In practice, jetty sites are selected in situations where a temporary solution is preferred, rather than constructing a complete onshore LNG terminal and the associated infrastructure.

3.3. Offloading Methods

FLNG–LNGC offloading can be achieved through a limited number of established configurations, and understanding these methods is essential for assessing operability and system design. Two of which are discussed in the following subsections.

3.3.1. Side-by-Side (SBS) Offloading

Side-by-side offloading is the most widely applied method for LNG export from FLNG facilities. In SBS configuration, the LNG carrier (LNGC) berths parallel to the FLNG with pneumatic fenders providing hull separation and mooring lines controlling relative position [31]. This proximity allows the use of marine loading arms (MLAs), which provide high-flow cryogenic transfer and are adapted from onshore LNG terminals for offshore use. On Prelude FLNG, MLA motion limits form critical operability constraints, and exceedances trigger emergency shutdown procedures [2].

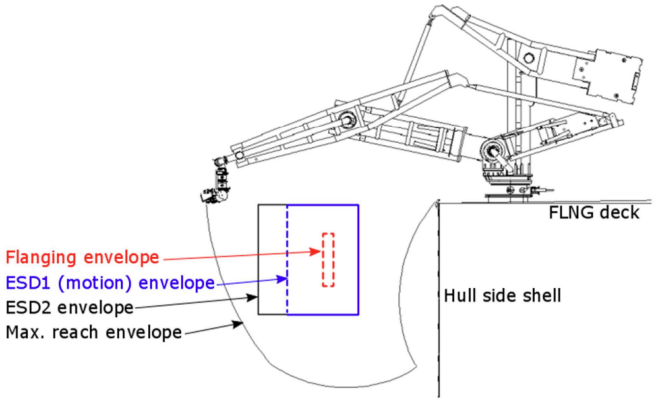


Figure 3.5: Schematic of MLA envelopes: flanging envelope, Emergency Shut Down envelopes, and maximum reach [2].

Figure 3.5 shows the MLA motion envelopes that govern operability limits during SBS offloading.

SBS offloading in open-sea conditions introduces significant challenges. The narrow gap between hulls leads to complex hydrodynamic interactions including gap resonance, LNGC roll motions, and fender and mooring loads with nonlinear properties [31]. Roll motion, in particular, is often the limiting criterion for safe SBS operation. Accurate prediction of roll damping, relative motions, and multi-body dynamics is therefore essential for operability assessment.

For turret-moored configurations, fender loads and line tensions remain mostly within allowable limits even under moderate-severity sea states [68]. Conversely, in spread-moored SBS cases, roll motions and relative displacements at manifold locations increase significantly in beam seas, producing higher fender compression and mooring line tensions [70].

3.3.2. Tandem Offloading

Tandem offloading positions the LNGC astern of the FLNG, which is very common in Floating Production Storage and Offloading (FPSO)–shuttle tanker operations. Hydrodynamically, this configuration benefits from vessel alignment with incident waves, which eliminates gap resonance effects and significantly reduces lateral motions [71]. As a result, tandem offloading generally exhibits more favorable motion characteristics compared with SBS, especially in harsher environments.

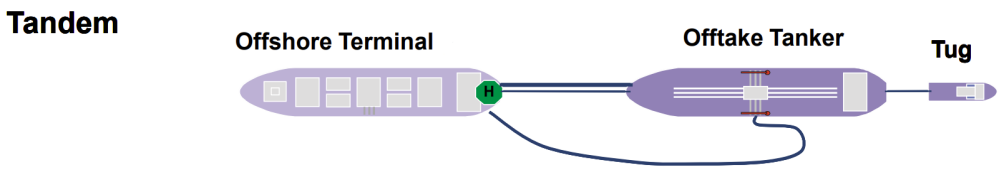


Figure 3.6: Tandem offloading configuration [69]

However, LNG tandem transfer requires flexible cryogenic hoses. Unlike oil transfer hoses, LNG hoses must withstand extremely low temperatures and dynamic loading [31]. Additionally, most LNG carriers are not equipped with bow loading manifolds required for tandem LNG transfer. Hydrodynamic studies and model tests show that although the FLNG remains relatively stable, the LNGC may exhibit significant fish-tailing behavior depending on hawser configuration, environmental forces, and vessel separation [71].

3.3.3. Implications of Mooring and Offloading Choices

The choice of mooring system directly influences the feasibility of SBS or tandem offloading. Turret-moored FLNGs benefit from natural weathervaning, reducing relative motions and improving MLA operability, thereby making SBS the preferred option [2]. Spread-moored FLNGs, constrained in heading, may encounter unfavorable beam-sea conditions that significantly reduce SBS operability. Time-domain simulations show that for spread-moored SBS configurations, roll angles frequently exceed the allowable limits for marine loading arms, manifold-relative motions become large in beam waves, and mooring-line tensions may approach operability limits in moderate sea states [70].

Tandem offloading becomes more attractive for spread-moored systems because the astern alignment eliminates beam-wave excitation and reduces lateral motions and fender loads. However, this option is currently limited by the availability and qualification of flexible cryogenic hose systems. In contrast, turret-moored FLNGs retain adequate SBS operability over a broader range of environmental headings due to weathervaning, reducing the need for tandem arrangements except in extreme conditions [68].

3.4. Side-by-Side Offloading of an spread-moored FLNG

The side-by-side (SBS) offloading operation between an FLNG and an LNG carrier consists of three main phases: (1) berthing and mooring, (2) LNG transfer, and (3) departure and unmooring. Although this thesis focuses on the LNG transfer phase, a brief overview of the full sequence is provided below.

1) Berthing and Mooring: The LNGC approaches and aligns alongside the FLNG with tug support. Safe berthing requires low relative motions to ensure controlled fender engagement and mooring-line connection. Prelude experience shows that excessive roll or sway can quickly compromise this phase [2], while hydrodynamic studies highlight that beam-sea exposure on spread-moored FLNGs increases drift forces and LNGC roll during approach [31]. Typical duration: 2–3 hours.

2) LNG Transfer: Once moored, operability is governed by the marine loading arm (MLA) motion limits. Roll is typically the most restrictive motion due to multi-body interaction and gap resonance [31]. LNG transfer is the most important phase of the offloading operation because it involves the export of the facility's primary product and also represents the longest part of the offloading sequence. Transfer operations typically last around 24 hours, depending on LNGC size and loading rate, which makes it essential that a safe and stable operability window is reliably determined to ensure uninterrupted and secure transfer.

3) Departure and Unmooring: Departure reverses the berthing sequence: mooring lines are released in a controlled order while tugs maintain separation. As during approach, stable conditions are required to avoid excessive relative motion. Prelude experience indicates that uncontrolled roll or lateral movement in marginal seas can lead to secondary contact [2]. Departure typically requires 1 hour.

3.5. Criteria for Side-by-Side Offloading Operability

SBS offloading operability describes the ability of an LNG carrier (LNGC) to safely berth, remain moored, transfer LNG, and depart alongside an FLNG facility under prevailing metocean conditions. Operability is achieved only when environmental limits, mooring-system capacities and vessel-motion constraints are simultaneously satisfied throughout the full offloading window. The following provides a concise overview of the key criteria described by Zhao et al. [31].

Mooring System Criteria Safe offloading requires that mooring line tensions remain below a specified fraction of their breaking load and that fender deflections stay within allowable limits [31]. These limits ensure the mechanical integrity of the connection between the vessels under dynamic loading.

Vessel Motion Criteria Allowable relative motions in surge, sway, and heave are governed by the operating limits of the marine loading arms. A minimum separation distance must be maintained to avoid hull contact. LNGC roll is typically the most restrictive criterion, as excessive roll reduces clearance, increases relative motions, and interacts with sloshing and gap-resonance effects [31].

Operability is evaluated by checking each criterion over successive 3-hour sea states within the total offloading window. If any limit is exceeded, the entire window is deemed inoperable [31]. This method captures seasonal variability, multi-year metocean fluctuations, and coupled hydrodynamic interactions between the vessels.

Table 3.2: Overview of Typical SBS Offloading Operability Criteria (after [31]).

Criterion Category	Description
Environmental limits	Combined $H_{s,\text{total}}$; sensitivity to long-period swell; maximum wind speed; tug operability constraints.
Mooring system limits	Maximum allowable mooring line tensions; maximum fender deflection to prevent damage.
Vessel relative motions	Limits on $X_{\max}$, $Y_{\max}$, $Z_{\max}$ based on marine loading arm envelopes; minimum required hull clearance.
LNGC roll limit	Most restrictive criterion; governs collision risk, MLA operability, and interaction with sloshing and gap resonance.
Overall operability check	All limits must be satisfied over the full offloading window.

3.6. Hydromechanical Considerations

Side-by-side (SBS) offloading brings two large ship-shaped hulls into close proximity, creating a tightly coupled hydrodynamic system. The resulting motion behaviour differs fundamentally from that of isolated vessels because waves are scattered, radiated, trapped and amplified in the narrow gap between the hulls. A hydromechanical understanding of these effects is therefore essential for assessing the feasibility and operability of SBS offloading.

Roll Motion and Viscous Effects. Roll motion of the LNG carrier is widely recognised as one of the most influential degrees of freedom in SBS configurations. Ship-shaped hulls exhibit relatively low roll damping in deep water, with damping governed predominantly by viscous flow separation, vortex shedding at sharp bilges, and possible contributions from bilge keels [72]. When two vessels are positioned closely together, the roll behaviour is further modified by proximity effects: radiated waves from one hull interact with the other, altering the effective hydrodynamic damping and shifting natural roll periods [67]. This interaction can generate coupled roll responses even in wave headings such as head seas where an isolated vessel would roll only marginally.

Gap Hydrodynamics and Resonance. The narrow gap between the FLNG and the LNG carrier forms a semi-enclosed fluid region in which free-surface oscillations can develop. Depending on gap

width, hull geometry and incident wave direction, these oscillations can become resonant, resulting in amplification of water surface elevations and velocities between the vessels. This so-called *gap resonance* leads to strong wave-induced pressures and contributes to increased sway, roll and heave responses of the LNG carrier [73, 74]. Hydrodynamic studies show that these resonant modes arise from the constructive interference of incident, diffracted and radiated waves, and may occur even in moderate sea states due to the geometric confinement of the fluid [67].

Hydrodynamic Interaction Between Hulls. Beyond gap resonance, the two vessels interact through first-order wave radiation and diffraction. Close proximity leads to scattering of waves between the hulls, modifying added mass, damping and exciting forces compared with single-body cases [75]. Huijsmans, Pinkster and De Wildeb show that mean drift forces—particularly in sway—are significantly altered when the hydrodynamic interaction is included, due to momentum exchange and fluid acceleration in the gap region [75]. Similarly, Buchner et al. demonstrate that multi-body interaction generates substantial additional low-frequency loads, even in head seas, driven by wave reflection from the upstream bow and subsequent trapping of low-frequency wave components between the hulls [76]. These interaction effects contribute to unexpected sway and yaw responses that would not arise for isolated vessels.

Coupled Multibody Motions. The proximity of the vessels causes their motion responses to become dynamically coupled. As shown by Kim et al. [67], relative sway, heave and roll amplitudes may be dominated by coupled resonances that occur when the motions of the vessels become out of phase. Even small differences in natural periods can cause large relative motions, as the combination of radiated and diffracted waves creates additional excitation pathways. Naciri and Waals [77] describe low-frequency surge, sway and yaw modes that arise from the collective mass, hydrodynamic coefficients and geometric arrangement of the two vessels. These modes are weakly damped and can become energetically significant when excited by second-order wave drift forces.

Relative Motions at Offloading Points. From a hydromechanics viewpoint, relative motions at the manifold or loading-arm location result from a superposition of global 6-DOF vessel responses and local hydrodynamic interaction. Kim et al. [67] show that:

- Longitudinal relative motions are often largest in head and bow-quartering seas, influenced by the phase difference between surge and pitch motions of the two vessels.
- Lateral relative motions can arise even in head seas due to interaction-induced sway and yaw, despite nominally symmetric wave excitation.
- Vertical relative motions are highest in beam seas, where roll and heave are jointly amplified by gap dynamics and wave incidence.

These relative motions are driven by hydrodynamic phase differences between the vessels and by the amplification of responses associated with multibody interaction.

Low-Frequency Hydrodynamic Loads. Second-order wave drift forces are important in SBS configurations. When the two vessels are close together, their wave-induced pressure fields interact, which can either increase or decrease the low-frequency forces acting on them depending on the wave direction [75].

Naciri, Waals and De Wilde [77] further identify characteristic slow-drift modes of the coupled system, some of which involve out-of-phase sway and yaw motions with low damping. These LF motions can

dominate the overall relative displacement envelope, especially in environments with long-period swell or oblique waves. Overall, the hydromechanics of SBS offloading are governed by a combination of viscous roll effects, free-surface dynamics in the narrow gap, multi-body wave radiation and diffraction, and low-frequency drift forces.

3.7. Conventional Operability Assessment

Following the hydromechanical discussion in the previous section, it is important to examine how these physical mechanisms have traditionally been represented in operability assessments. Conventional offshore LNG offloading analyses rely heavily on simplified, threshold-based approaches that approximate the system behaviour rather than resolving the full complexity of multibody hydrodynamics. Typical tools include environmental scatter diagrams, linear Response Amplitude Operators (RAOs), and look-up tables or operability matrices derived from frequency-domain results. These methods are computationally efficient and remain widely used in early-phase engineering, but they struggle to reflect the full dynamics arising from vessel–vessel interaction.

Early work by Huijsmans et al. (2001) [75] and Fournier et al. (2006) [78] showed that directly applying linear potential-flow solvers to side-by-side configurations can lead to unphysical amplification of gap waves. This occurs because the free-surface boundary condition in potential flow does not represent the viscous dissipation present in narrow gaps. To mitigate this, Pauw et al. (2007) [79] introduced the epsilon-damping lid method, which suppresses unrealistic free-surface oscillations and yields hydrodynamic coefficients suitable for time-domain simulation.

Conventional operability assessment has also historically treated berthing and approach manoeuvres separately from the offloading operation. Van Doorn and ten Hove (2002) [80] demonstrated that fast-time, autopilot-based simulations can provide consistent estimates of tug forces and approach trajectories, although full-mission bridge simulations remain necessary to evaluate human-factor effects.

In parallel, simplified downtime assessments traditionally relied on limiting environmental criteria—such as significant wave height (H_s), peak period (T_p), or maximum wind speed—applied directly to long-term metocean databases. However, such scalar thresholds do not capture important hydromechanical phenomena: multibody coupling, gap resonance, low-frequency drift forces, and nonlinear mooring and fender behaviour.

To partly address these limitations within a simplified framework, several authors proposed using time-domain simulations to populate operability matrices. Grin et al. (2005) [81] and Poldervaart et al. (2006) [82] demonstrated that simulation-based look-up tables can incorporate nonlinear effects—including mooring elasticity, low-frequency motions, and hydrodynamic interaction—provided that a sufficiently dense matrix of simulations is performed.

A major advancement was introduced by de Wilde et al. (2009) [51], who proposed a direct time-domain approach in which each environmental condition is simulated for 3.5 hours, with the final 3 hours used to assess downtime. This differs from simulation-based look-up tables because almost every sea state present is simulated directly rather than approximated by interpolation from a precomputed matrix, allowing nonlinear multibody effects to be captured for each individual condition. Their results showed that traditional threshold-based approaches can substantially underpredict downtime, particularly in beam-sea or quartering swell conditions where roll amplification and slow-drift forces dominate.

Building on this methodology, MARIN's aNySIM time-domain software (MARIN, 2025) [83] has become a widely used tool for detailed operability analysis. aNySIM computes the six-degree-of-freedom motions of each vessel, including full multibody interaction, by combining retardation-function-based

hydrodynamics (derived from DIFFRAC) with nonlinear mechanical models for moorings, fenders, and loading arms. The software captures both wave-frequency and low-frequency behaviour, drift forces, gap-induced excitation, and mooring line dynamics, enabling high-fidelity simulations of complete offloading operations. However, this level of realism comes at a computational cost: evaluating an entire environmental matrix may require hundreds of thousands of simulations.

3.8. Machine Learning in Multibody/Offshore Dynamics

Machine learning (ML) has emerged as a powerful tool for modeling complex physical systems in the offshore and marine industries. In this domain, researchers have applied various models to predict the dynamic responses of vessels and structures. For instance, Guo et al. (2021) [84] used a Long Short-Term Memory (LSTM) network to predict the heave and surge motions of a semi-submersible platform. Their work is highly relevant as it demonstrates that an LSTM can learn complex dynamics from physical model test data, achieving high accuracy for predictions up to 46.5 seconds into the future, a significant horizon for operational decisions. Similarly, Zhang et al. (2021) [85] developed a multiscale attention-based LSTM to predict ship motions. Their key contribution was decomposing motion signals into different frequency scales before feeding them to the model, which improved accuracy by allowing the network to learn the distinct characteristics of wave-frequency and low-frequency effects. These applications underscore the utility of LSTMs, which, as established by Hochreiter & Schmidhuber (1997) [86], use a gating mechanism to capture long-term dependencies in time-series data.

To further enhance predictive power, hybrid architectures have been explored. Lee et al. (2023) [87] developed an integrated system combining a Convolutional Neural Network (CNN) with an LSTM encoder-decoder to predict wave-induced ship motions. Their model's relevance lies in its use of the CNN to process spatiotemporal wave-field data, which the LSTM then uses to forecast motion, and in the recursive modelling technique they apply to extend predictions over longer horizons.

More advanced architectures, such as the Transformer model introduced by Vaswani et al. (2017) [88], have also been adapted for these tasks. Zhang et al. (2023) [89] employed a Transformer network to predict the 6-DOF motions of ships in real operational conditions, showing that its attention mechanism can capture the influence of factors like vessel speed and heading on dynamics, a crucial capability for modeling realistic operational scenarios.

3.8.1. Long-Horizon Time-Series Motion Prediction

Long-horizon forecasting is critical for operational planning but poses a challenge due to error accumulation in sequential prediction. To address this, Zhou et al. (2021) [90] developed the Informer model, a specialized Transformer architecture for Long Sequence Forecasting (LSF). By using a self-attention mechanism and a generative-style decoder, Informer can efficiently predict long time-series sequences in a single forward operation, mitigating cumulative errors. This is particularly relevant for predicting entire 3-hour simulations. Building on this, Shen et al. (2025) [91] proposed a pre-trained Informer model for multi-step ship motion prediction. Their mechanism-data dual-driven framework used a physics-based model to generate virtual data for pre-training, which improved the model's accuracy and reduced data dependency when fine-tuned on real experimental data, offering a path to train models even with limited SBS data.

The performance of these models also depends on the input data representation. Liu et al. (2020) [92] investigated input vector optimization for LSTM models, finding that using the Auto-Correlation Function (ACF) to determine the input dimension improved both training efficiency and prediction accuracy for ship motions. This finding is directly relevant for efficiently training the time-series models in this thesis.

3.8.2. ML for Mooring Line Tension and Load Prediction

Accurate prediction of mooring line tensions is crucial for structural integrity. Much of the existing research focuses on the anchoring systems of single floating structures rather than the inter-vessel lines used in side-by-side configurations. For example, Yang et al. (2023) [93] used a Reservoir Computing (RC) model to predict the mooring tension of a barge and a KVLCC2 ship. By decomposing the tension signal into low-frequency and wave-frequency components, their model achieved high accuracy for predictions up to 80 seconds into the future. This decomposition strategy is a key technique that could be applied to the more complex SBS problem. For the even more complex system of a floating offshore wind turbine (FOWT), Zhong et al. (2025) [94] used a Multi-input Deep Operator Network with Attention (MIONet-ATT) to predict mooring tensions from 6-DOF platform motions, demonstrating superior generalization over classical neural networks. These studies highlight that while ML is effective for station-keeping systems, its application to the highly dynamic, coupled mooring lines between two vessels remains less explored.

3.8.3. ML for Operability Assessment and Decision Support

Ultimately, the goal of motion and load prediction is to support operational decisions. Auburtin et al. (2019) [95] developed a model for the Coral South FLNG project using a Radial Basis Function (RBF) Artificial Neural Network (ANN). This model directly mapped sea state parameters (wind, waves, and current) to a binary ”operable” or ”not operable” classification, providing a rapid go/no-go decision tool without running new simulations. This represents a ”summary-statistics” approach, where the model learns a meta-relationship from a database of pre-run simulations, directly informing the first part of this thesis's methodology. In a different approach aimed at real-time support, Nielsen et al. (2018) [96] demonstrated that the observed autocorrelation function of a vessel's recent motion history could be used to make deterministic predictions for the near future without needing any wave input data, which is valuable when environmental data is uncertain.

3.9. Toward ML-Driven SBS Operability Assessment and Research Gaps

The literature reveals two promising research trajectories for applying ML to SBS offloading operability. The first uses simplified, computationally efficient models to map environmental summary statistics directly to a peak response or an operability classification. The work by Auburtin et al. (2019) [95] is a prime example, showing that such models are promising for rapid operational assessments.

The second trajectory involves time-series models that predict the full dynamics of a system. For forecasting the complex, coupled dynamics of a multibody system over short to medium horizons, an LSTM-based forecasting model, such as that used by Lee et al. [87], can be applied. For the even more ambitious task of predicting an entire multi-hour simulation from another, Transformer-based models like Informer offer a promising solution due to their efficiency with very long sequences [90, 91]. In fact, many of the most recent innovations in AI applications stem from Transformer architectures, and applying such models to operability assessment in offshore operations would represent a highly exciting and potentially disruptive development.

However, a clear research gap emerges at the intersection of these approaches and the specific application of SBS offloading. There is an absence of studies applying these modern ML models to a complete, coupled multibody system like an FLNG and LNGC. Specifically, the literature lacks:

- A systematic comparison of summary-statistics-based models versus full time-series models for

predicting the operability of an SBS configuration or any offshore configuration.

- The application of advanced forecasting models like LSTMs or Transformers to predict the coupled 6-DOF motions and inter-vessel loads for a full SBS system.
- An end-to-end prediction framework capable of generating a full 3-hour time-series of motions and loads for an SBS operation directly from environmental conditions.

Addressing these gaps would be particularly impactful because successful ML models could significantly reduce the offshore sector's strong dependence on computationally expensive numerical simulations. A surrogate that reliably replaces or enhances time-domain simulations would not only accelerate operability assessment but could also enable real-time or near-real-time decision support.

Moreover, an especially powerful direction would be to combine simulation or measurement data from multiple SBS configurations and locations. This would make it possible to build a general operability model that works across many different side-by-side setups-learning the common motion patterns of different hull shapes, mooring systems, and environmental conditions. Such a model would be a big step toward an operability tool that is easy to reuse, scale, and apply across the offshore industry.

This deficiency in the literature leads directly to the central question of this thesis:

How can machine learning-based models support and accelerate operability assessment for side by side offloading?

By developing and comparing both summary-statistics-based and time-series-based models for this specific application, this work aims to fill this critical gap and provide a framework for machine learning-based operability assessment of SBS Offloading.

To address this main question, the following sub-questions structure the investigation:

- To what extent can operability-relevant responses be predicted using only environmental summary statistics?
- Can multibody response behaviours be inferred from partial time histories?
- Under which environmental variations can full three-hour multibody responses be reconstructed from wave elevation and environmental inputs when predicting unseen wave conditions?
- How can we construct a minimal dataset which captures all conditions and system dynamics?

4

Case Description and Theoretical Framework

For this thesis, Shell and Maritime Research Institute the Netherlands (MARIN) provided access to simulation data from an offshore Liquified Naturag Gas (LNG) development study. This chapter presents the case study for this thesis on side-by-side (SBS) offloading operability for a a spread-moored Floating LNG (FLNG) facility with an LNG carrier (LNGC) moored alongside. The analysis is based on a simulation campaign conducted by MARIN for an offshore FLNG development. The dataset includes a long-term hindcast-driven batch of simulations and a large generic environmental matrix. These data serve as ground truth for the models developed in this thesis and provide a representative benchmark for industrial offshore operability assessments [83].

A spread-moored FLNG maintains a fixed heading and cannot change its orientation with environmental forcing, making it sensitive to waves, swell, wind, and current. The relative motion between both vessels, and the corresponding loads on moorings and fenders, become critical. High offloading operability is essential to prevent production stops due to tank-top limitations and to maintain continuous gas export [31]. Poor operability translates into costly downtime and deferred production hence the need for a rigorous, data-driven evaluation. It motivates the detailed operability assessment that follows and explains why the next sections focus on the specific criteria, modelling workflow, and environmental inputs for this spread-moored SBS system.

The remainder of this chapter is structured as follows. Section 4.1 describes the FLNG–LNGC setup and coordinate conventions. Section 4.2 introduces the operability criteria used in the simulations. Section 4.3 summarises the numerical workflow based on DIFFRAC and aNySIM. Section 4.4 presents the environmental inputs from the hindcast dataset and the generic matrix. Section 4.5 describes the two datasets generated by the simulations (summary statistics and full time-series records). Section 4.6 provides the hydrodynamic background required to interpret the simulation results. Section 4.7 discusses the distinction between low-frequency and wave-frequency vessel motions. Section 4.8 outlines modelling considerations and limitations relevant to interpreting the aNySIM dataset. Finally, Appendix B provides a deeper explanation of the underlying simulation theory.

4.1. Configuration and Coordinate Systems

The configuration consists of a spread-moored 170 000 m^3-class FLNG, modelled in ballast condition, and a 140 000 m^3-class LNG carrier in loaded condition, where the numerical designations refer to the LNG storage capacities of the vessels. The FLNG's heading is fixed at $112.5°$ clockwise from North (illustrated in Figure 4.2 and the LNGC is moored port-side with a 4.2 m fender gap. The underlying hydrodynamic model was originally developed for a starboard-side arrangement; therefore, a validated heading-mirroring procedure was applied to obtain the port-side configuration used in this study [83].

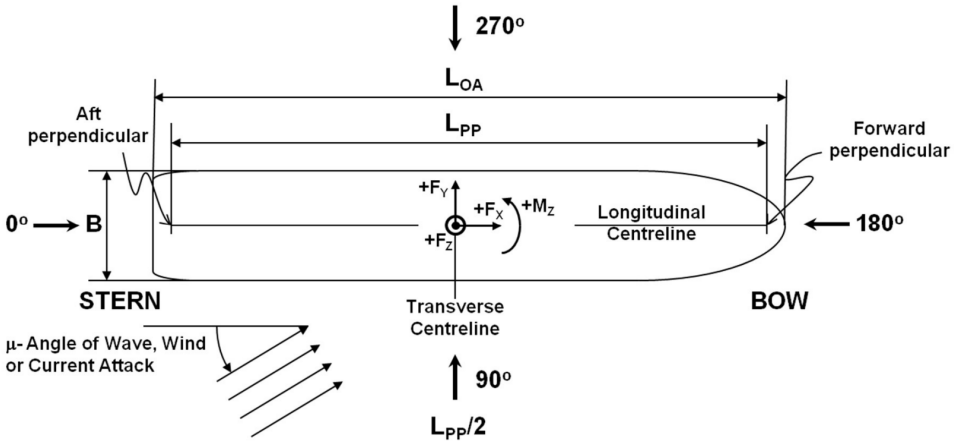


Figure 4.1: Sign convention and coordinate system used in the simulations [83].

Motions and forces are reported in an OCIMF local coordinate system (surge, sway, heave; roll, pitch, yaw) and in a global east–north–up system. Environmental directions are specified as "going to" in the global system, with zero degrees towards the vessels stern (Figure 4.1).

Sixteen SBS mooring lines are modelled with pretension of approximately 100 kN and non-linear restoring curves. Five floating fenders (diameter 4.5 m) define the hull-to-hull clearance. A still-water equilibrium leads to approximately 0.3 m fender compression [83].

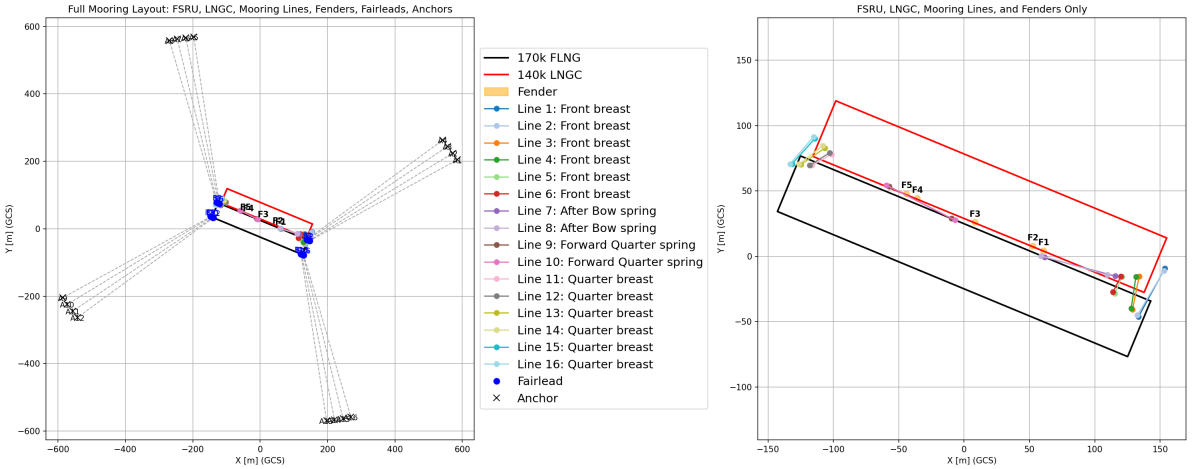


Figure 4.2: SBS mooring layout showing fenders and line leads, plotted as rectangular for length between perpendicular.

4.2. Operability Criteria

The study evaluates each simulation against a set of operability criteria consistent with engineering practice and hardware limits [83]. These criteria, originally defined in the numerical model documentation, are summarised in Table 4.1. They cover mooring line loads, fender compression, relative manifold motions, and roll motions of both vessels. Each parameter is checked for exceedance of its threshold using maximum and minimum values from the time-domain statistics.

Table 4.1: Operability criteria used to evaluate each simulation (adapted from MARIN Summary Report).

Criterion	Description	Limit
Mooring line load	Maximum tension in any of the 16 SBS mooring lines.	$\leq 50\%$ of MBL (672 kN)
Fender compression load	Reaction force corresponding to compression of floating fenders.	$\leq 50\%$ deflection ($\approx$2997 kN)
Relative manifold motions (surge/sway)	Relative motion between LNGC manifold and FLNG loading-arm pedestal.	$\leq \pm 2.0$ m
FLNG roll motion	Roll angle of the FLNG.	$\leq \pm 3^\circ$
LNGC roll motion	Roll angle of the LNGC.	$\leq \pm 3^\circ$

These limits reflect both hardware constraints—loading arms, mooring lines, and fenders—and operational judgement. For example, the roll limits reduce the risk of approaching the loading-arm motion envelope, a practical threshold also reported in simulator-based studies [2].

The MARIN Summary Report provides further detail:

- **Mooring line tension:** The 50% MBL limit follows OCIMF practice for synthetic fibre lines; for the lines used, the governing MBL is 1344 kN.
- **Fender compression:** Reaction forces are taken from the fender compression–force curve. A 50% compression corresponds to 2997 kN.
- **Manifold motions:** Surge and sway must remain within a ± 2 m envelope, representative of typical loading-arm operating limits.
- **Roll motions:** Both vessels must remain within $\pm 3^\circ$.

Together, these thresholds form the pass/fail basis used for both the hindcast and the generic environmental batches. A simulation is deemed operable only when all criteria are satisfied throughout the evaluated period, consistent with standard practice in offshore LNG operations.

4.3. Simulation Framework

The operability assessment by MARIN used in this study is based on two complementary numerical tools: a frequency-domain solver for hydrodynamic coefficients and a time-domain multibody simulator for non-linear system response.

Frequency-domain hydrodynamics (DIFFRAC)

Linear hydrodynamic coefficients are computed with DIFFRAC, MARIN's diffraction–radiation solver. DIFFRAC evaluates the potential-flow boundary value problem using a panel method to obtain added

mass, radiation damping, first-order excitation forces, and second-order drift loads for the coupled FLNG–LNGC system [83]. The analysis is performed for the site water depth and for a full matrix of frequencies and headings, producing the hydrodynamic database used by the time-domain model. The two-body treatment ensures that interaction effects such as gap resonance and shielding are captured consistently.

Time-domain simulation (aNySIM XMF)

Vessel motions and mechanical interactions are computed using aNySIM XMF, MARIN's multibody time-domain simulation framework. aNySIM integrates 6-DOF rigid-body equations of motion including frequency-dependent added mass and damping (via retardation functions), hydrostatics, wave excitation, wind and current loads, and non-linear mooring and fender forces [83]. It uses explicit Runge-Kutta solvers with adaptive time stepping, supported by MARIN's XMF architecture, which provides a modular basis for hydrodynamic and mechanical components. For the FLNG–LNGC configuration, aNySIM models:

- coupled two-body motions using DIFFRAC hydrodynamic input,
- SBS mooring line and fender loads, with pretension and non-linear stiffness curves and force–deflection behaviour,
- relative motions at the manifold [83].

Integrated workflow

The integrated workflow of the different parts combined is illustrated below, in Figure 4.3.

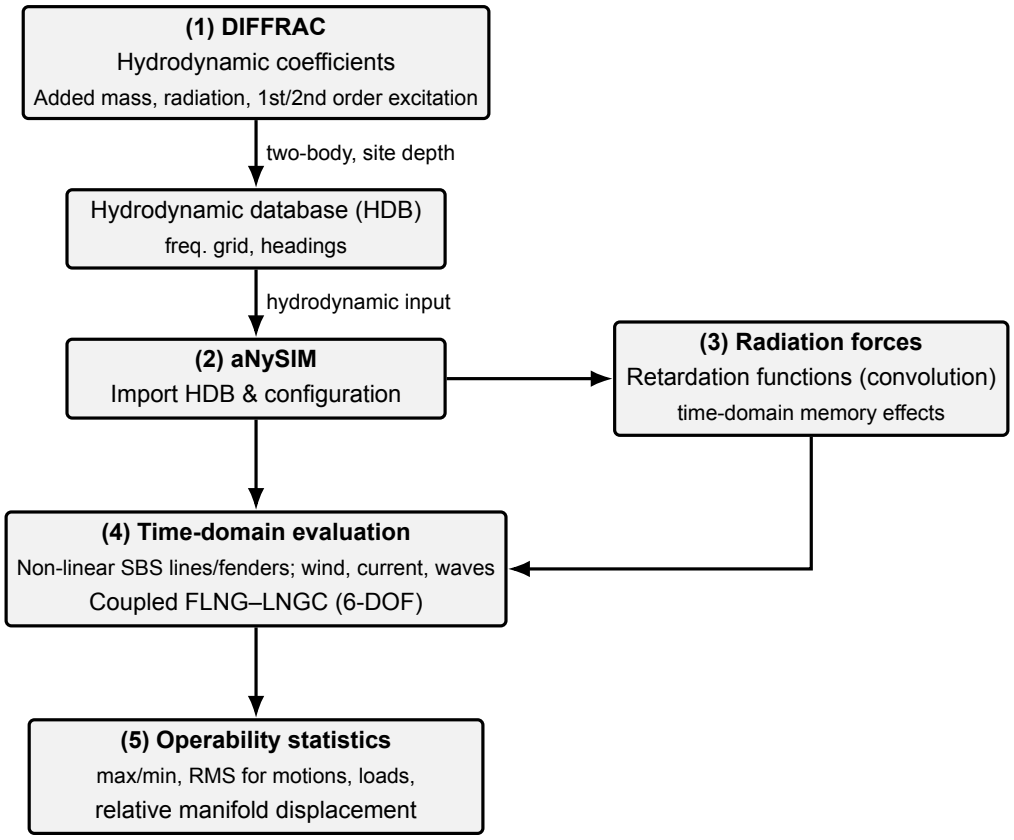


Figure 4.3: Integrated workflow: DIFFRAC hydrodynamics → aNySIM time-domain simulation → operability statistics.

Both DIFFRAC and aNySIM are long-standing MARIN developments. aNySIM builds on earlier simulation tools (TERMSIM, LIFSIM, DPSIM, DREDSIM, DYNFLOAT) and is supported by the Extensible Modelling Framework (XMF), a modular simulation toolkit focused on Newtonian dynamics, object interoperability and extensibility [97].

For the study used in this thesis, a pre-existing multi-body aNySIM model, originally developed for a similar LNG side-by-side offloading assessment, was adapted. A new two-body linear diffraction analysis was generated for the representative water depth. The baseline model assumed the LNGC moored on the starboard side of the FLNG; however, the configuration of interest required a port-side arrangement. To accommodate this, environmental inputs were mirrored about the FLNG centerline, following MARIN's documented and validated mirroring procedure [83]. The resulting outputs are treated as accurate ground truth for the purpose of model development.

4.4. Environmental Inputs

The environmental conditions used in the simulations consist of a generic environmental matrix designed to span a wide range of sea states. This matrix enables parametric exploration across combinations of significant wave height, peak period, and wave direction, paired with predefined wind and current settings [83]. In the MARIN simulations, the wave signal is the *only* environmental input that varies in time; all other environmental parameters, remain constant during each three-hour simulation.

Although the full generic simulation campaign includes both sea and swell wave systems, the present thesis considers only the sea-wave subset. This reduces the number of simulations required for processing and ensures consistency across the dataset, without affecting comparability between different environmental conditions. The conditions included in the generic set are listed in Table 4.2 and Table 4.3.

Table 4.2: Generic matrix (sea-only): parameterization used in the study.

Parameter	**Symbol**	**Unit**	**Values**
Significant wave height	$H_{s,\mathrm{sea}}$	m	0.75, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5
Peak period	$T_{p,\mathrm{sea}}$	s	6, 7, 8, 9, 10, 11
Wave direction	θ_{sea}	deg	90, 105, 120, 140, 160, 180, 200, 220, 240, 255, 270
Total wave cases			$7 \times 6 \times 11 = 462$

Table 4.3: Wind and current combinations applied to each sea-only wave case.

Case	V_{wind}	θ_{wind}	V_{cur}	θ_{cur}	Case	V_{wind}	θ_{wind}	V_{cur}	θ_{cur}
1	1	90	0.12	90	10	4	180	0.24	180
2	1	270	0.12	270	11	4	225	0.24	225
3	4	90	0.24	90	12	4	315	0.24	315
4	4	270	0.24	270	13	10	0	0.24	0
5	10	90	0.24	90	14	10	45	0.24	45
6	10	270	0.24	270	15	10	135	0.24	135
7	4	0	0.24	0	16	10	180	0.24	180
8	4	45	0.24	45	17	10	225	0.24	225
9	4	135	0.24	135	18	10	315	0.24	315

Combining the 462 wave conditions with the 18 wind–current cases yields a total of 8316 sea-wave simulations in the generic dataset.

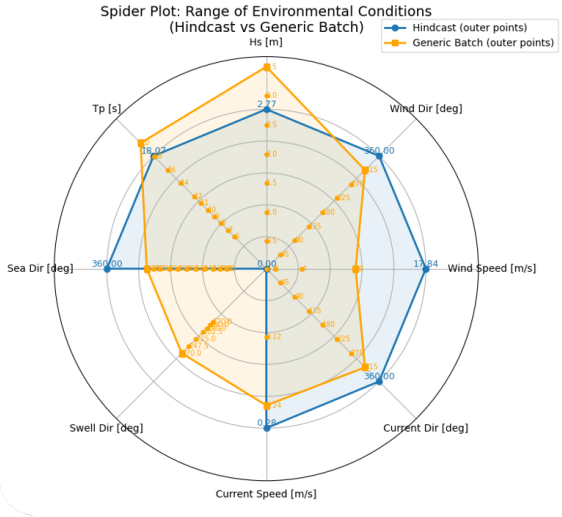
Table 4.4: Batch matrix for swell-only and combined sea–swell conditions (adapted from MARIN dataset).

Parameter	Symbol	Unit	Values
Significant wave height	$H_{s,\text{swell}}$	m	0.5, 0.75, 1.0, 1.5, 2.0, 2.5
Peak period	$T_{p,\text{swell}}$	s	12, 14, 16, 18, 20
Wave direction	θ_{swell}	deg	90, 112.5, 135, 157.5, 170, 180, 190, 202.5, 225, 247.5, 270
Total swell cases			$6 \times 5 \times 11 = 660$
Significant wave height (combined)	$H_{s,\text{sea+swell}}$	m	Sea: 1.0, 1.5, 2.0, 2.5 Swell: 0.5, 1.0, 1.5
Peak period (combined)	$T_{p,\text{sea+swell}}$	s	Sea: 6, 8, 10 Swell: 12, 16, 20
Wave direction (combined)	$\theta_{\text{sea+swell}}$	deg	Sea: 90, 120, 140, 160, 180, 200, 220, 240, 270 Swell: 170, 180, 190
Total combined cases			Sea: $4 \times 3 \times 9 = 108$ Swell: $3 \times 3 \times 2 \times 3 = 54$ Combined: $108 \times 54 = 5832$

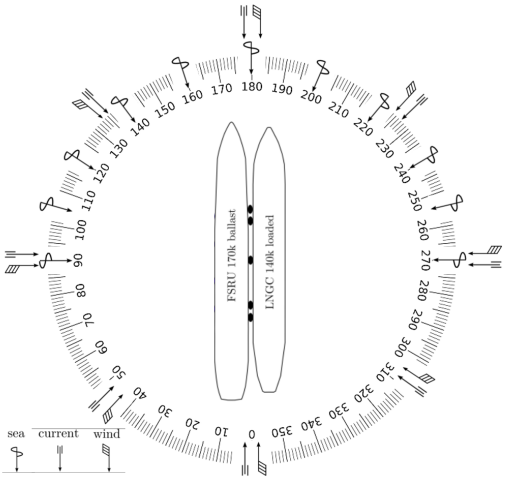
In addition to the 8316 sea-wave simulations, the full MARIN campaign also included swell-only and combined sea–swell conditions. The swell-only (660) and combined sea–swell (5832) cases are each also combined with the same 18 wind–current conditions, resulting in 11,880 and 104,976 cases respectively, giving a total of 125,172 simulations.

Figure 4.4 compares the environmental ranges represented in the generic matrix with those of the

hindcast dataset at the reference FLNG location and illustrates the distribution of incoming wave, wind, and current directions. Although the hindcast set is not used further in the ML analysis, it is included here to show that the generic matrix spans a wider range of wave conditions than those observed on site, ensuring that the model development covers an extended environmental envelope.



(a) Hindcast vs. generic matrix



(b) Incoming wave, wind, and current directions in the generic set

Figure 4.4: Comparison of environmental parameterization between datasets.

Figure 4.4 compares the generic environmental matrix with the hindcast dataset at the reference FLNG location. Although the hindcast set is not used in the ML analysis, it is shown here to illustrate how the generic matrix spans a much wider range of wave, wind, and current conditions than those observed on site. This broader coverage ensures that the generic set governs an extended environmental envelope suitable for training and evaluating machine-learning models beyond the conditions captured in the hindcast record.

4.5. Simulation Outputs

The aNySIM simulations generate two complementary datasets: (i) a *summary dataset* containing per-run statistics used for operability assessment, and (ii) *full time-series records* of all inputs and system responses.

Summary dataset

For every simulated sea state, MARIN provides a summary dataset containing the minimum and maximum values of the output parameters (Table 4.5). These extrema are derived from the full aNySIM time-domain simulations but retain only the key response quantities relevant for operability assessment.

This summary dataset forms the basis of MARIN's operability study for Shell. Each simulated condition is classified as operable or non-operable depending on whether any of the output extrema exceed predefined acceptance limits. The resulting operability classification is reported in the MARIN spreadsheets for both the hindcast dataset and the generic environmental matrix.

Full time-series dataset

In addition to the statistical summaries, MARIN supplied the complete time-series records for all simulations for use in this thesis. These records include the input wave elevation signal as well as the response time series of vessel motions, mooring line tensions, fender forces, and relative motions.

Whereas the summary dataset condenses each simulation into a limited set of extrema, the full time-series data resolve the complete dynamic behaviour of the coupled FLNG–LNGC system. These data capture low-frequency drift, wave-frequency oscillations, transient load peaks, and nonlinear interaction effects that cannot be inferred from summary statistics alone.

The output parameters included in both datasets are listed in Table 4.5.

Table 4.5: Overview of individual simulation output parameters.

Category	Individual parameters
6-DOF motions (COG)	FLNG: surge (x), sway (y), heave (z); roll, pitch, yaw.
	LNGC: surge (x), sway (y), heave (z); roll, pitch, yaw.
Manifold motions (global)	FLNG manifold: x, y, z positions; roll, pitch, yaw.
	LNGC manifold: x, y, z positions; roll, pitch, yaw.
Relative motions (FLNG–LNGC)	Relative manifold motions: surge (x), sway (y), heave (z); relative manifold distance; relative roll angle.
SBS mooring line tensions	Line-by-line tensions: line 1, line 2, …, line 16.
Fender compression forces	Fender 1, fender 2, fender 3, fender 4, fender 5.

Computational Cost of Time-Domain Simulations

Each aNySIM sea-wave simulation covers approximately 3.5 hours of physical time and requires around 10 minutes of wall-clock time to complete. In the MARIN batch execution setup, 100 simulations are assigned sequentially to a single CPU core, leading to a per-core execution time of

$$100 \times 10 \text{ min} = 1000 \text{ min} \approx 17 \text{ h}.$$

For the 8,317 sea-wave simulations analysed in this thesis, this corresponds to

$$8{,}317/100 \approx 84 \text{ CPU cores}.$$

Each MARIN compute node consists of two Intel Xeon 6126 CPUs with 12 physical cores each, providing 24 cores per node. The total number of required compute nodes is therefore

$$84/24 \approx 3.5 \Rightarrow 4 \text{ nodes}.$$

The total computational effort of the original aNySIM workflow for the full sea-wave dataset is thus

$$\text{Total CPU usage} = 84 \text{ cores} \times 17 \text{ h} = 1{,}428 \text{ core-hours}.$$

This figure provides a quantitative reference against which the computational efficiency of the data-driven models developed in this thesis can be evaluated.

4.6. Hydrodynamic Background

The simulations used in this thesis resolve a combination of wave-frequency and low-frequency responses, driven respectively by first-order and second-order wave loads, as well as mooring, fender, wind, and current interactions. While the detailed theory is presented in Appendix B, this section provides a concise overview of the key physical ingredients, first-order RAO-based excitation, second-order drift forces, and the coupled 12-DOF equations of motion-that shape the motion and load characteristics of the system. These concepts are essential for interpreting the MARIN simulation outputs and for understanding the inputs and targets used in the machine-learning models developed in later chapters.

The theoretical background summarized here follows MARIN's aNySIM Theory Documentation [97] and the formulations of Journée and Massie [98]; to avoid repetition, these sources are not cited again in the following subsections and additional citations are introduced only when new theory or assumptions beyond these baselines are used. Unless stated otherwise, the flow is assumed inviscid, incompressible, and irrotational, enabling potential-flow modelling of wave-body interactions together with linear wave theory.

4.6.1. RAO-based representation for irregular seas

Hydrodynamic solvers supply the first-order wave forces through Response Amplitude Operators (RAOs), which express the steady-state response per unit wave amplitude for each frequency and direction. MARIN uses their diffraction analysis in DIFFRAC to obtain these RAOs in the frequency domain. They are transfer functions used only as inputs to time-domain load reconstruction; they are not solved in the time domain directly.

For an irregular wave field represented by n frequency components, the linear excitation in degree of freedom k becomes

$$q_k(t) = \sum_{i=1}^{n} \mathrm{RAO}_k(\omega_i, \mu_i)\, \bar{a}_i, \tag{4.1}$$

where ω_i and μ_i are the frequency and direction of component i, and $\bar{a}_i$ is its complex amplitude. This formulation is fully consistent with the potential-field decomposition and the first-order pressure integration as described in subsection B.2.2.

Application to side-by-side configurations. For two bodies moored side-by-side, such as an FLNG and an LNG carrier, the hydrodynamic behaviour of each vessel is strongly modified by the presence of the other. Reflections, gap resonances, and hydrodynamic shielding introduce additional interaction effects that alter both the incoming wave field and the vessels' radiation and diffraction behaviour. When the hydrodynamic coefficients are obtained from a multi-body panel model, these interaction effects are automatically embedded in the motion responses used by aNySIM. As a result, the predicted first-order loads include the relevant wave-body and body-body interaction mechanisms..

Practical implication. This partition reflects how the separation databases and aNySIM are used: the diffraction set provides first-order excitation (via RAOs), whereas the radiation set provides added mass and damping (in the frequency domain) or retardation kernels (in the time domain). It also clarifies the modelling scope: Froude–Krylov and diffraction describe incoming-wave loading, while radiation accounts for motion-induced wave-making and energy dissipation (see Appendix B and in particular subsection B.2.2). For side-by-side vessels, both sets include cross terms, which are essential for predicting gap resonance, hydrodynamic shielding, and coupled surge–sway–yaw behaviour.

4.6.2. Second Order Wave (Drift) Forces

Second-order (drift) wave loads are the main cause of slow surge, sway, and yaw motions in moored or side-by-side floating structures. These slow-drift forces come from the nonlinear interaction of pairs of regular wave components, which creates difference-frequency loads that align with the natural low-frequency motions of the system. Because mooring systems usually have low stiffness and only limited damping at these frequencies, these slowly changing forces tend to dominate the low-frequency behaviour.

Quadratic Transfer Function (QTF) Formulation: The second-order drift forces can be compactly represented in the frequency domain using Quadratic Transfer Functions (QTFs). The second-order force in degree of freedom k is written as

$$F_k^{(2)} = \sum_{m=1}^{M} \sum_{n=1}^{M} QTF_k(\omega_m, \omega_n, \mu_m, \mu_n)\, \bar{a}_m \bar{a}_n, \tag{4.2}$$

where $\bar{a}_m$ and $\bar{a}_n$ are the complex wave amplitudes and QTF_k is the quadratic transfer function obtained from the second-order boundary-value problem. The QTF contains both sum- and difference-frequency contributions:

$$\omega_+ = \omega_m + \omega_n, \tag{4.3}$$

$$\omega_- = |\omega_m - \omega_n|. \tag{4.4}$$

4.6.3. Remarks on Hydrodynamic Coupling and Operability

When two vessels float close to each other, the water between and around them causes their motions and loads to become strongly connected. The main effects can be summarised as follows:

1. Waves generated by one vessel affect the other: When either hull moves, it radiates waves. These radiated waves act on the neighbouring hull and appear in the hydrodynamic model as *off-diagonal* added-mass and radiation-damping terms. Keeping these cross-coupling terms is essential for predicting correct wave-frequency heave, pitch, and coupled vertical-plane responses [76].
2. The presence of one vessel changes the incoming wave field on the other: Waves scatter between the hulls, and the up-wave vessel partly shields the down-wave vessel. This modifies the first-order excitation loads on both hulls, especially in narrow-gap configurations [77].
3. Low-frequency forces become stronger when the vessels are close: Difference-frequency drift forces in surge, sway, and yaw increase due to proximity effects. These low-frequency motions dominate the relative horizontal movement between the two vessels and therefore strongly influence SBS mooring-line loads and overall operability [77, 51].

Practical implication: Because SBS lines are short and stiff, even small inaccuracies in modelling these interaction effects can lead to large errors in predicted line loads. Reliable operability assessment therefore requires accurate low-frequency forcing, consistent damping, and proper treatment of gap hydrodynamics [77].

4.6.4. Coupled Multi-Body System (12-DOF)

When two vessels are moored side-by-side, their motions are coupled through both hydrodynamics and the mechanical connection (mooring lines, fenders). The global configuration and velocity vectors

collect the 6-DOF states of both vessels and are therefore of dimension 12×1:

$$\mathbf{X} = \begin{bmatrix} \boldsymbol{\eta}_1 \\ \boldsymbol{\eta}_2 \end{bmatrix}, \qquad \mathbf{V} = \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix},$$

with each subvector $\boldsymbol{\eta}_i$, $\mathbf{v}_i \in \mathbb{R}^6$.

Global inertia and hydrodynamic coupling: The global inertia matrix has size 12×12 and is block-diagonal, since each vessel contributes its own 6×6 total inertia matrix $\mathbf{M}_{\infty,i}$:

$$\mathbf{M}_{\infty}^{\text{glob}} = \begin{bmatrix} \mathbf{M}_{\infty,1} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}_{\infty,2} \end{bmatrix}.$$

In contrast, the global radiation kernel $\mathbf{R}^{\text{glob}}(t)$ is also a 12×12 matrix but, it contains non-zero 6×6 off-diagonal blocks. These off-diagonal terms describe hydrodynamic interaction: motions of vessel 1 generate radiated waves that act on vessel 2, and vice versa, and are essential for representing the coupled surge–sway–yaw dynamics in narrow-gap configurations.

Mechanical coupling: Side-by-side mooring lines, fenders and the FLNG anchoring system also introduce cross-body forces. These appear in the mechanical load term $\mathbf{F}^{\text{glob}}(t)$ and result in additional off-diagonal contributions.

Global equation of motion: The coupled 12-DOF system is written as:

$$\mathbf{M}_{\infty}^{\text{glob}}\, \dot{\mathbf{V}}(t) + \int_0^t \mathbf{R}^{\text{glob}}(t-\tau)\, \mathbf{V}(\tau)\, \mathrm{d}\tau + \mathbf{K}^{\text{glob}}\, \mathbf{X}(t) = \mathbf{F}_{\text{glob}}(t), \tag{4.5}$$

where $\mathbf{K}^{\text{glob}}$ contains the hydrostatic restoring terms for each vessel, and $\mathbf{F}_{\text{glob}}$ collects all hydrodynamic, mechanical and environmental loads.

Equation (4.5) forms the dynamic backbone of the side-by-side simulation. It ensures that (i) wave radiation from one hull acts on the other, (ii) drift forces and gap dynamics influence the relative surge–sway–yaw behaviour, and (iii) nonlinear SBS mooring-line and fender loads are applied consistently at each time step.

4.6.5. Interpretation

The structure of the global equations highlights the physical mechanisms governing side-by-side operations:

- The inertia matrix reflects both rigid-body and hydrodynamic inertia.
- The convolution term captures memory effects from radiated waves.
- The stiffness term controls static stability and small-amplitude restoring behaviour.
- The right-hand side collects all excitation, damping, environmental and mechanical loads that drive relative motions and mooring-line forces.

This framework provides a description of the coupled dynamics of the FLNG and LNGC, and is consistent with both the hydrodynamic data generated by DIFFRAC and the nonlinear mechanical models used in aNySIM developed by MARIN.

4.7. Low-Frequency vs. Wave-Frequency Motions

Floating structures exhibit a combination of wave-frequency (WF) motions, driven directly by the incident waves, and low-frequency (LF) motions, excited by the slowly varying second-order forces and environmental loads. This distinction is fundamental for understanding the behaviour of moored vessels, where horizontal-plane responses are dominated by LF effects while vertical-plane motions are controlled by WF excitations.

4.7.1. Wave-Frequency Motions

Wave-frequency (WF) motions originate from the linear, first-order excitation forces acting on the hull. These forces oscillate at the same frequencies as the incident waves and are directly proportional to wave amplitude. As a result, WF motions dominate the vertical and rotational degrees of freedom: heave, roll, and pitch. Their magnitudes follow from the linear hydrodynamic coefficients, hydrostatic stiffness, and response amplitude operators (RAOs).

For a moored tanker in head seas, the WF content of surge is typically small because the horizontal-plane natural periods of the mooring system lie far below the wave-frequency range. By contrast, heave and pitch respond strongly at WF, with amplitudes proportional to wave elevation and wave slope, respectively, as is commonly observed in motion records.

4.7.2. Low-Frequency Motions

Low-frequency (LF) motions arise mainly in the horizontal degrees of freedom—surge, sway, and yaw—and are governed by the interplay between second-order drift forces, wind, current, and the restoring forces from the mooring system or fenders. Because the horizontal restoring stiffness is small compared to hydrostatic restoring in the vertical plane, the natural periods of these modes lie far below the wave-frequency range, often between 30 s and several minutes. The LF natural period can be approximated by

$$T_{\mathrm{LF}} = 2\pi\sqrt{\frac{m + A_{\mathbf{LF}}}{k_{\mathrm{moor}}}}, \tag{4.6}$$

where m is the vessel mass, A_{LF} the added mass at low frequency, and k_{moor} the effective horizontal restoring stiffness. Because both radiation damping and viscous damping are small at these low frequencies, LF responses can reach large amplitudes when the excitation frequency approaches the natural frequency. Accounting for the natural low-frequency period is essential in a side-by-side system because it determines how strongly the two vessels will resonate under slow-drift excitation. When the low-frequency content of the second-order drift forces aligns with the horizontal-plane natural frequency of the FLNG-mooring system, small excitation can lead to large relative excursions between the FLNG and the LNGC. Such resonant amplification directly affects manifold alignment, fender compression, and mooring-line loading.

The key contribution for moored ships arises from the difference-frequency components,

$$\omega_- = |\omega_i - \omega_j|,$$

which fall within the low-frequency range of surge, sway, and yaw. These forces include both a slowly varying part and a non-zero mean component, the latter responsible for steady wave drift. Because they scale with the square of the wave amplitude, the resulting LF motions grow approximately with the square of the significant wave height.

A typical illustration of the relation between first- and second-order motions is shown in Fig. 4.5, where the rapid first-order oscillation is modulated by a slow second-order envelope and a mean offset.

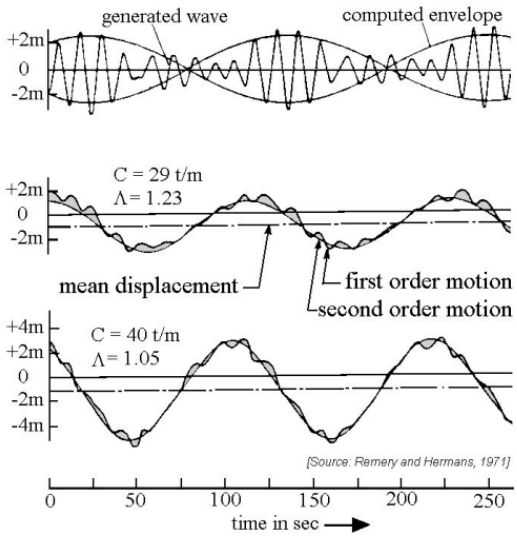


Figure 4.5: Illustration of first- and second-order motions: the rapid first-order response is superimposed on a slowly varying second-order modulation and a mean drift component.

In irregular seas, slowly varying drift forces arise naturally from the wave groups present in a random sea state. These groups introduce low-frequency modulations in the wave envelope, and the corresponding second-order drift forces vary at frequencies close to the difference frequencies of the wave components. When these encounter frequencies lie near the natural low-frequency period of the moored vessel, the structure undergoes large slow-drift excursions due to resonance. The resulting horizontal motions typically grow in proportion to the square of the significant wave height, reflecting the quadratic nature of the underlying second-order forcing.

An overview of wave-frequency and low-frequency motion is characteristics given in Table 4.6.

Motion Type	Origin	Dominant DOFs	Restoring Mechanism
Wave-Frequency (WF)	First-order, linear wave excitation	Heave, Roll, Pitch	Hydrostatics
Low-Frequency (LF)	Second-order drift, wind, current	Surge, Sway, Yaw	Mooring lines, fenders

Table 4.6: Comparison of wave-frequency and low-frequency motions.

4.7.3. Engineering Implications

The combined WF and LF behaviour has important implications for mooring system design and operability assessment. WF motions determine vertical accelerations, slamming risk, and relative motions, while LF motions govern the maximum excursions and cyclic mooring loads. Because the LF restoring and damping are small, slow-drift motions often dominate horizontal response and must be included explicitly in design calculations using second-order hydrodynamic loads. Accurate representation of both first- and second-order forces is therefore essential for reliable prediction of vessel motions in waves.

4.8. Discussion

This section explains what the hydrodynamic modelling choices imply for machine-learning models that aim to predict operability from the MARIN simulations, and how the results in this thesis should be

interpreted.

Uncertainty near operability limits. Time-domain models can agree on mean motions yet differ on peak loads because of details such as wave-force interpolation, heading treatment, solver strategy, or fender friction. Small changes in relative surge or pitch may push line tensions or fender forces just above or below a limit. For this thesis, this means labels close to thresholds (e.g., 50% MBL, 50% fender compression) are naturally less stable in the dataset; prediction errors around these limits should therefore be expected and read with caution [77, 83].

Headings and low-frequency regimes. In close proximity, gap effects and slow-drift sway–yaw introduce extra uncertainty, especially in beam and quartering seas. For this thesis, some headings and period ranges are intrinsically harder to learn from the data and will tend to show larger prediction scatter than others [76, 51].

Sensitivity of wave-frequency responses. Wave-frequency heave and pitch in the MARIN simulations are generated from hydrodynamic coefficients that are first computed in the frequency domain and then used by aNySIM to build the time-domain loads through RAOs and retardation functions [76, 83]. This means that the time-domain motion statistics used to train the ML models inherit any sensitivity present in the underlying frequency-domain hydrodynamic data. Small variations in added mass, radiation damping or first-order excitation therefore translate into shifts in wave-frequency response levels, and the ML predictions reflect this built-in sensitivity of the simulation model.

Nonlinear mechanics and configuration dependence. Mooring and fender behaviour is nonlinear and depends on pretension, damping, and friction. Two sea states that look similar can lead to different peaks if these settings differ slightly. In this thesis, variance in load-related targets may persist even when average motions look alike; such differences do not necessarily indicate modelling errors [77, 51].

Dataset coverage and splits. The applicability of the models outside the training domain depends strongly on the generic dataset, as this set spans a wide range of environmental conditions and headings beyond what occurs in the hindcast. Its broad envelope defines the limits within which the models can reasonably be expected to generalise, while its dense environmental grids may place very similar cases next to each other. Consequently, performance should be interpreted with respect to both overall behaviour and behaviour within specific regimes [83].

Scope of interpretation. The simulations done by MARIN represent high-quality, industry-standard assessments, yet residual uncertainty remains in areas such as gap damping, viscous effects, and extremes. In this thesis, the machine-learning results are interpreted as fast surrogates of the aNySIM assessment rather than as full-scale truth. Their role is to support rapid screening and highlight cases that demand high-fidelity analysis[76, 51, 83].

5

Methodology

The methodology of this thesis focuses on developing machine learning models to predict the operability of side-by-side (SBS) offshore offloading systems under varying sea states. Traditional operability assessments rely on extensive numerical simulations, which are computationally expensive and time-consuming. This work explores data-driven alternatives that aim to reduce computational effort while maintaining prediction accuracy.

The modeling pipeline is designed to leverage the available simulation dataset in two distinct ways. The first approach simplifies the problem by using summary statistics of vessel responses, while the second approach utilizes full three-hour simulation time series data to capture dynamic interactions. These complementary strategies allow for a systematic comparison between simplified and detailed representations of vessel behavior.

1. **Summary Statistics-Based Learning (SSBL):** This approach uses aggregated response values, the maximum and minimum values of critical parameters from three-hour simulations. It investigates whether these condensed representations can provide reliable operability predictions for unseen sea states with reduced computational complexity.
2. **Time Series-Based Learning (TSBL):** This approach incorporates complete wave and response time series, enabling the model to learn temporal dependencies and system dynamics. It aims to improve prediction fidelity and obtain understanding of the vessel–environment interactions.

Both approaches are evaluated for their ability to generalise across sea states and their potential applicability in operational decision-making. The comparison highlights trade-offs between model complexity, accuracy, and scalability providing insights into the suitability of machine learning techniques for offshore operability analysis.

The combined use of SSBL and TSBL provides a comprehensive framework for exploring machine learning applications in offshore operability assessment. By contrasting simplified representations of vessel responses with full dynamic time series, this methodology demonstrates the trade-offs between computational efficiency and predictive fidelity. Together, these approaches enable a deeper understanding of how data-driven models can replicate or approximate physics-based simulations, offering insights into their potential for reducing the reliance on a large number of numerical analyses. This dual strategy not only aims to highlight the feasibility of machine learning in predicting operability under

varying sea states but also lays the foundation for developing a structured framework to guide future assessments.

5.1. Summary Statistics-Based Learning

In the SSBL part of the research, the different machine learning models are used to predict the maximum response of the operability limiting parameters in a three-hour simulation of a given sea state. No further context, just the environmental input parameters are given and the maximum response is predicted. The models are trained on different percentages of the total training set to assess the impact of the amount of data on these kinds of predictions. What will we learn from this? Is it possible to predict limiting parameters for operations like SBS offloading, which involves a multibody system, simplified to a context where only the environmental parameters are known?

The objective is to predict the maximum operability responses in a three-hour sea state from its environmental condition vector $\mathbf{c} \in \mathbb{R}^d$, which collects the wave, wind, and current parameters. The targets are the corresponding maxima $\mathbf{y}_{\max} \in \mathbb{R}^K$ of the operability-limiting responses.

In this work, the output dimension K reflects the structure of the multibody system. Only the motions that directly govern SBS operability are considered: the roll motion of the FLNG, the roll motion of the LNGC, and the relative manifold distance in surge and in sway (four motion-related parameters in total). For the mooring line tensions, a single parameter is used to represent the maximum tension across all sixteen mooring lines. For the fenders, a single parameter represents the maximum compression load among the five fenders. This results in

$$K = 4 + 1 + 1 = 6,$$

which are predicted simultaneously as the operability-limiting response maxima.

The supervised learning task is expressed as the mapping

$$f_\theta : \mathbf{c} \longmapsto \hat{\mathbf{y}}_{\max}, \tag{5.1}$$

where f_θ is a machine learning model trained to approximate $\mathbf{y}_{\max}$ (maximum response) from the static environmental inputs alone:

$$\hat{\mathbf{y}}_{\max} = f_\theta(\mathbf{c}) \approx \mathbf{y}_{\max}.$$

This setup tests how well we can predict the peak responses of the multibody offshore system from the environmental parameters of a sea state, where each peak is taken from a three□hour time□domain simulation driven by a linear JONSWAP spectrum, without relying on the underlying time□series data itself. The model is outlined in Figure 5.1. The three different machine learning models are shortly discussed and explained in the following subsections.

The importance of this first part of the research lies in identifying which predicted parameters are graspable for a simple setup model as in SSBL and which parameters are more difficult to understand and why. And how does this translate into the implementation of the models applied in the second phase of the research.

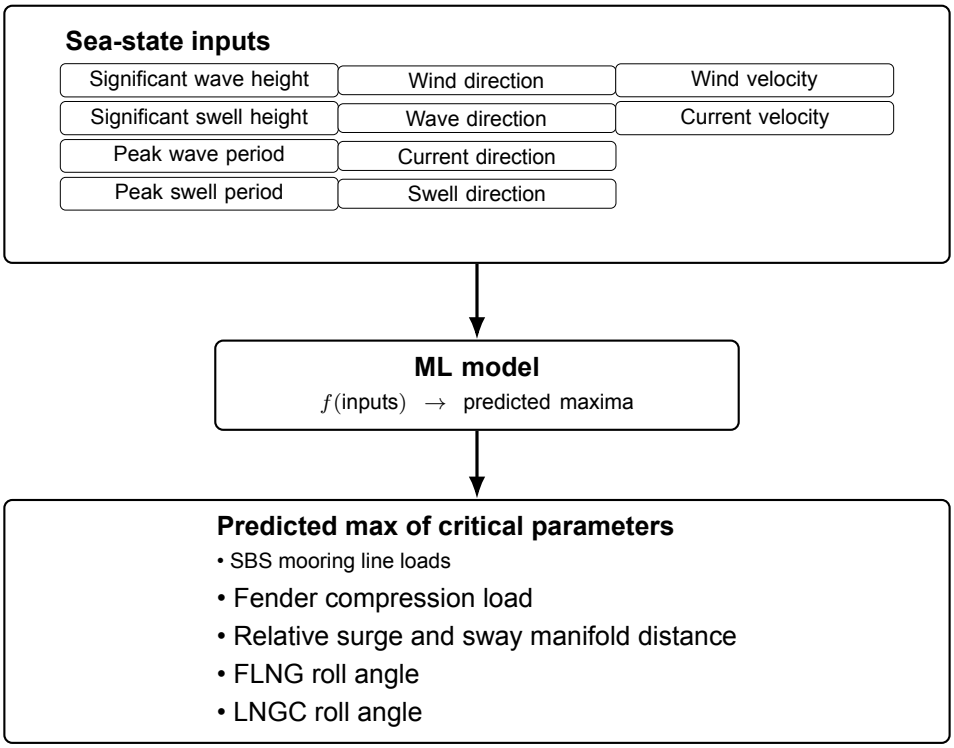


Figure 5.1: Summary statistics ML model: from sea-state inputs to predicted operability parameters.

5.1.1. Random Forest

Random Forests were first introduced by Ho (1995) [99] as “random decision forests,” combining multiple decision trees built from randomly selected subsets of data and features. Breiman (2001) [100] later formalised the algorithm and extended it to regression, showing that aggregating predictions from many trees trained on bootstrap samples improves accuracy and robustness.

A Decision Tree predicts outcomes by recursively splitting data based on feature values until a stopping criterion is met, such as maximum depth or minimum samples per leaf. For regression, each leaf outputs the average target value of its samples, enabling the model to capture nonlinear relationships while remaining interpretable [101].

Random Forest Regression enhances this approach through:

- Bootstrap Sampling: Each tree is trained on a random subset of the data.
- Feature Subsampling: At each split, only a random subset of features is considered.

The final prediction is the average of all tree outputs, reducing variance and mitigating overfitting while maintaining strong predictive performance.

5.1.2. Multi-Layer Perceptron (MLP)

The multilayer perceptron (MLP) is a feedforward neural network that learns a non linear mapping between inputs and outputs using stacked layers of interconnected neurons. Each neuron computes a weighted sum of its inputs followed by a non linear activation function, which allows the network to represent complex relationships that cannot be captured by linear models [102]. The MLP used in this thesis consists of an input layer, one or more hidden layers, and an output layer. Information flows from input to output without recurrent connections.

Training is performed with the backpropagation algorithm, which adjusts the network weights by computing the gradient of a loss function with respect to each parameter [102]. This gradient is obtained by applying the chain rule through the network layers, and the parameters are updated using gradient descent. In practice the MLP improves its predictions by iteratively reducing the loss until convergence.

To promote generalisation and avoid overfitting the MLP applies several concepts that are standard in neural network training. These ideas ensure that the model does not become overly sensitive to noise in the training data. The essential components for this thesis are

- controlling the overall capacity of the network through the number of layers and neurons
- the use of non-linear activation functions, which allow the model to learn complex patterns while maintaining numerical stability
- regularisation techniques such as weight penalties or dropout that reduce sensitivity to small fluctuations in the data
- learning rate scheduling and mini batch training which improve convergence behaviour and robustness

In this thesis, the MLP is used as a supervised regression model that maps sea-state properties to operability related response metrics. By combining non-linear transformations with gradient based optimisation the model is able to capture relationships that are difficult to learn with tree based methods alone.

Predictive Capability of Summary Statistics Based Learning Models

This part of the study evaluates how well different model classes predict maxima of operability related responses from sea-state summary inputs only. The analysis serves two goals. First it identifies which target variables are learnable from aggregated environmental parameters without additional temporal context. Second it quantifies how predictive performance changes with the amount of training data.

The comparison across Random Forest and the MLP is informative because each model encodes structure differently. Tree-based models partition the input space in a hierarchical way and are effective when responses align with threshold-like interactions, meaning that small changes in sea-state parameters can trigger sudden shifts between response regimes rather than smooth variations. The MLP learns smooth non-linear transformations through gradient based optimisation and is effective when responses vary smoothly with the inputs. By evaluating these models on the same tasks and under the same data budgets the study isolates which responses are robust to model choice and which responses are sensitive to architectural assumptions or data volume.

The results from this stage guide the next steps in two ways. They indicate which responses are sufficiently determined by summary statistics and which responses likely require information about phase relationships and temporal evolution.

5.2. Time Series-Based Learning

The TSBL approach in this thesis is designed to investigate how the use of full wave and response time traces can contribute additional insight into the dynamic behaviour of the coupled system. By utilising the complete temporal evolution from the simulations, this approach aims to capture aspects of the system's behaviour that cannot be represented through summary statistics alone, including how responses vary overtime.

Conceptually, the input wave elevation is represented as a continuous signal $\mathbf{x}[t]$ over the interval $t \in [0, T]$, where T is the duration of the simulation window. The corresponding vessel response is written as $\mathbf{y}[t]$, which includes both rigid-body motions and mooring line tensions. In practice, these continuous signals are sampled at a fixed interval Δt to produce discrete time series $\mathbf{x}[1{:}T]$ and $\mathbf{y}[1{:}T]$.

The system consists of two vessels, each modelled in six degrees of freedom (surge, sway, heave, roll, pitch, yaw), so the total motion DOFs are $2 \times 6 = 12$.

Additionally, mooring lines connecting the vessels contribute to (in this learning) scalar tension values.

Through data inspection and from SSBL it comes forward that in the case the fender compression load is exceeded, there will always be a different critical parameter that will be exceeded first. For that reason, and to reduce complexity of the suggested models, it is decided to *not* take the fender compression loads into account in the TSBL methods of this thesis.

Thus, for $N_{\text{lines}} = 16$, the total response dimension per timestep is:

$$\mathbf{y}[t] \in \mathbb{R}^{12+16} = \mathbb{R}^{28}.$$

The learning task is to map wave elevation and environmental conditions to a high-dimensional response vector that captures coupled vessel motions and mooring loads across time.

The objective is to learn a mapping:

$$f_\theta : \mathbf{x}[1:T] \longmapsto \hat{\mathbf{y}}[1:T],$$

where f_θ is a parametric model (e.g., a machine learning model) that predicts the full response trajectory given the input sequence. This mapping is influenced by environmental factors that vary across simulations, including wave height, wave period and direction, current speed and direction, and wind speed and direction.These parameters define the sea state and loading conditions, which significantly affect the vessel's dynamic response. The TSBL models in this thesis are trained only on sea states containing a single sea-wave system; swell waves and any combined sea–swell conditions are excluded from the time-series learning. So for the TSBL models the 8317 sea-wave conditions, as discussed in Appendix B, are used.

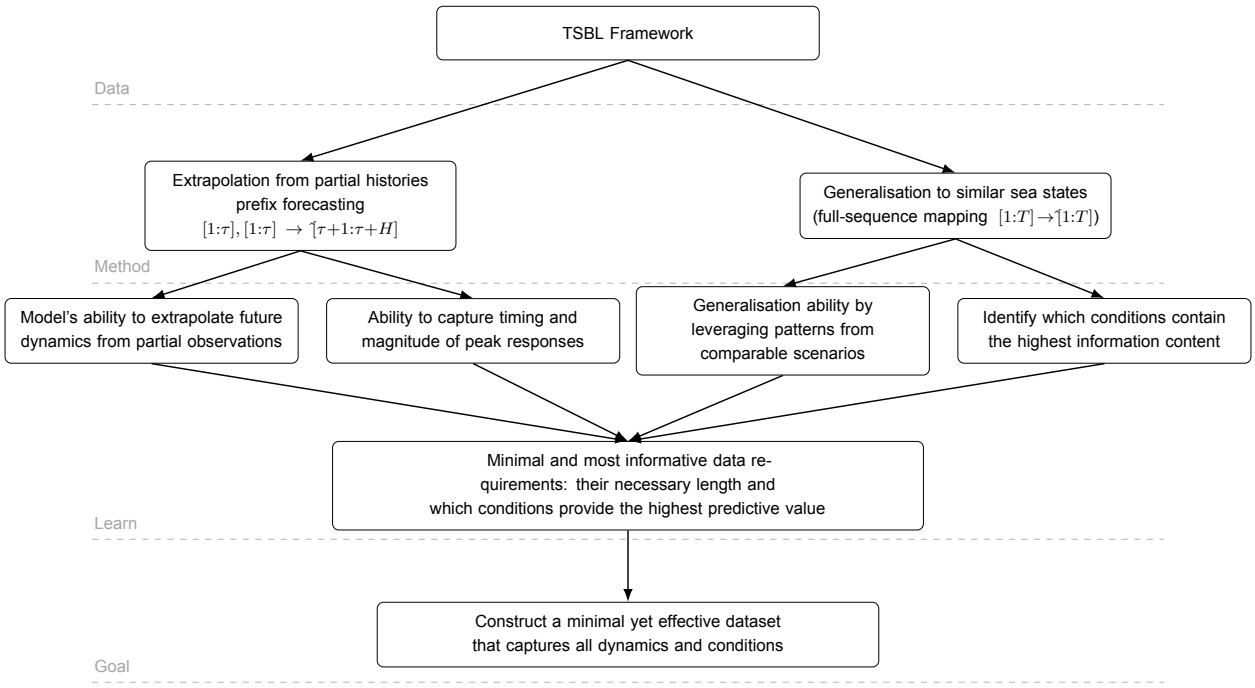


Figure 5.2: High-level overview of TSBL workflow.

For the TSBL part of this thesis, two distinct prediction methods are proposed, as also displayed in Figure 5.2:

1. **Forecasting from an initial segment:** Given an observed prefix $\mathbf{x}[1:\tau]$ and $\mathbf{y}[1:\tau]$ with $\tau < T$, the model predicts the remaining response $\hat{\mathbf{y}}[\tau+1:T]$.

   $$f_\theta : \big(\mathbf{x}[1:\tau], \mathbf{y}[1:\tau]\big) \longmapsto \hat{\mathbf{y}}[\tau+1:T].$$

   This approach evaluates the ability to extrapolate future dynamics from initially observed segment of the wave and motion history and capture timing and magnitude of peak responses.

2. **Generalisation from similar conditions:** Using full sequences from comparable sea states, the model predicts responses for a new condition without access to its outputs:

   $$f_\theta : \mathbf{x}[1:T] \longmapsto \hat{\mathbf{y}}[1:T], \quad \text{with } \mathbf{x} \sim \text{similar to training cases.}$$

   This tests the model's capacity to leverage patterns from related scenarios and identify which conditions contain the highest information content.

Together, these strategies aim to identify minimal yet informative data requirements, evaluate generalisation capability, and determine which conditions provide the highest predictive value. The ultimate goal is to construct a compact dataset that captures essential dynamics and variability for operability analysis.

5.2.1. Response signal decomposition

Vessel motions and mooring loads in waves occur at different time scales because they are driven by distinct physical mechanisms. The total response of the coupled FLNG–LNGC system can therefore be expressed as the superposition of multiple frequency components [97]:

- **Wave-frequency (WF) component:** Responses arising from *first-order wave excitation forces*, which are linear in wave amplitude and directly correlated with the incident wave spectrum. WF components dominate short-term dynamics and govern oscillatory vessel motions and relative motions at the interface.
- **Low-frequency (LF) component:** Responses induced by *second-order wave drift forces*, originating from quadratic wave interactions. These forces act over much longer time scales than WF oscillations and are critical for global vessel offsets, mooring line tensions, offloading operability, and station-keeping performance.

Combining these components into a single, undifferentiated signal introduces challenges for data-driven models:

1. **Different time scales:** LF drift responses evolve over tens of seconds to minutes, whereas WF responses occur at periods of approximately 5–25 seconds. A single model attempting to learn both simultaneously must capture fundamentally different temporal dependencies, which can degrade learning efficiency and generalisation performance.
2. **Distinct physical drivers:** LF behaviour depends strongly on second-order hydrodynamic effects and mooring system restoring characteristics, while WF behaviour is governed by first-order excitation forces, inertia, and hydrodynamic damping. Treating these components separately allows the models to learn more interpretable mappings that align with the underlying physics.

Accordingly, frequency-based filtering is applied to the raw motion responses and to the mooring line tension time series:

$$\mathbf{y}(t) = \mathbf{y}_{\mathrm{LF}}(t) + \mathbf{y}_{\mathrm{WF}}(t),$$

where $\mathbf{y}_{\mathrm{LF}}$ contains the components below a prescribed cut-off frequency and $\mathbf{y}_{\mathrm{WF}}$ contains the remaining higher-frequency content. This decomposition is applied consistently to vessel motions, relative motions, and mooring line tensions.

The approach follows the theoretical separation of first- and second-order wave loads described in the aNySIM Theory Documentation [97] and in [98]. By resolving the response into LF and WF components prior to learning, the resulting models are better able to capture coupled dynamics across scales and to respect the physical structure of offshore multibody behaviour.

As a result, all methodology steps described in this chapter are performed separately for the LF and WF components of the original response signals.

5.2.2. Segment-Based Sequence to Sequence Learning

For the segment-based sequence-to-sequence learning method of this thesis, the following is proposed: a encoder–decoder sequence-to-sequence neural predictor that forecasts multi-degree-of-freedom (DOF) vessel motions over a horizon of H steps using past motion and wave elevation measurements. The model is trained on short, overlapping windows extracted from the clean training region (first 70% of the simulation) and is applied recursively in overlapping chunks to generate long-horizon forecasts for the unseen final 30%.

The suggested model contains the following parts:

- Motion encoder: multi-layer LSTM that encodes past motion and past wave elevation into a memory (hidden and cell states).
- A wave feature extractor: a 1D CNN that takes a chunk of future wave data and turns it into a compact vector (an embedding).
- Motion decoder: LSTM initialised with the encoder memory. It takes the combined input (wave point-elevation data + environment info) repeated across the horizon and outputs the future motion sequence.
- Recursive chunking algorithm: repeatedly applies encoder+decoder on overlapping chunks to predict long horizons while preserving causality.

The encoder–decoder architecture predicts the future motion trajectory by combining two sources of information: (i) the past motion and wave history, compressed into an internal memory by the encoder, and (ii) the future wave characteristics summarised by the CNN embedding. The encoder captures the dynamics of the vessel in the short-to-medium term, while the decoder generates the entire future motion sequence of one horizon of H steps at once, instead of predicting each step one by one.

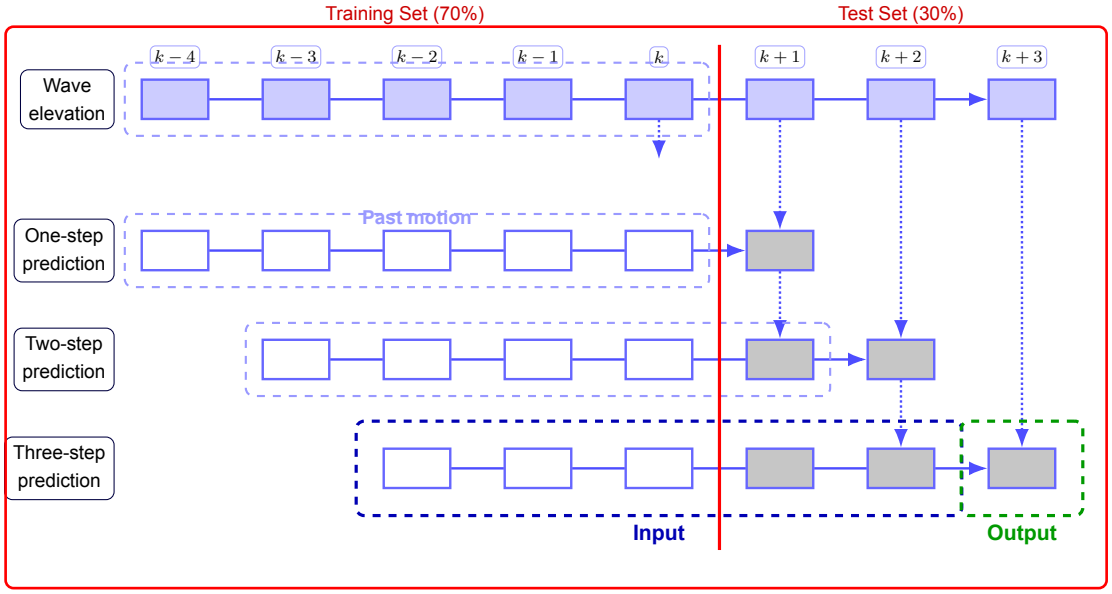


Figure 5.3: Recursive prediction visualization with train/test split and in- and output windows.

The recursive prediction mechanism, illustrated in Figure 5.3, operates by repeatedly feeding overlapping input windows into the encoder–decoder model to extend the prediction horizon beyond the length of any single chunk. During training, the model sees only short sequences from the first 70% of the

data (left side of the figure), learning how past motion, wave features, and environmental conditions jointly determine the next few steps.

At test time (right side), this learned mapping is applied recursively. After generating the first prediction horizon (e.g., $k{+}1$), the output is partially reused as part of the next input window. This produces the overlapping structure shown in the diagram, where predicted values (gray nodes) are injected back into the sequence for the next forward pass. By sliding the window forward and re-applying the same encoder-decoder block, the model constructs a long-horizon forecast while ensuring that each prediction step depends only on information available up to that point

Each new window combines real historical observations with the model's previous predictions, keeping the trajectory grounded in actual dynamics while enabling the forward propagation of its own forecast. This recursive overlapping process allows the system to generate multi-step trajectories such as $(k{+}1,\, k{+}2,\, k{+}3)$ even though training was limited to shorter windows. The stacked one-step, two-step, and three-step prediction paths in the figure illustrate how longer horizons emerge from deeper recursive chains. The subsequent subsections elaborate on the specific architectural design choices of the model.

5.2.2.1. Encoder LSTM

Role and Purpose: The encoder LSTM processes the past window of vessel motions and wave elevation and compresses this information into a compact internal state. This internal state captures the recent motion history of the vessel together with the corresponding wave, providing a concise description of the dynamic condition at the end of the past window.

Inputs and Construction: For each sample b, the encoder receives a past window of length P. At each timestep t, the motion vector $\mathbf{m}_b^{\text{past}}[t]$ and the wave elevation $w_b^{\text{past}}[t]$ are combined into a single input vector:

$$\mathbf{u}_b[t] = \text{concat}\big(\mathbf{m}_b^{\text{past}}[t],\, w_b^{\text{past}}[t]\big) \in \mathbb{R}^{D_m+1}.$$

Here, $\text{concat}(\cdot,\cdot)$ simply means putting two vectors next to each other to form one longer vector. Collecting these inputs over all timesteps forms the tensor

$$\mathbf{U} \in \mathbb{R}^{B\times P\times (D_m+1)},$$

which ensures that motion and wave information are processed jointly.

Operation: At each timestep, the LSTM maintains two internal vectors, namely the hidden state $\mathbf{h}_t$ and the cell state $\mathbf{c}_t$. Together, these vectors store short-term and long-term information extracted from the past window. The encoder processes the input sequence one timestep at a time. At each step, the LSTM updates its hidden and cell states based on the new input and its previous internal state:

$$(\mathbf{h}_t, \mathbf{c}_t) = \text{LSTM}_{\text{enc}}\big(\mathbf{u}_b[t],\, (\mathbf{h}_{t-1}, \mathbf{c}_{t-1})\big), \qquad t = 1, \ldots, P.$$

Through these updates, the LSTM retains information about slow vessel responses and remains sensitive to short-term variations.

Output Representation: After processing all P timesteps, the final LSTM states, $\mathbf{h}_P$,$\mathbf{c}_P$, form a compact representation of the vessel's dynamic condition at the end of the past window. These states initialise the decoder for predicting future motions.

5.2.2.2. Convolutional Wave Compressor

Role and Purpose: A one-dimensional convolutional network compresses the future wave elevation sequence to a fixed-size embedding. The goal is to summarise the dominant temporal structures of the incoming sea state over the prediction horizon without passing the full raw sequence to the decoder.

Inputs and Construction: The CNN receives the future wave segment

$$\mathbf{w}_b^{\text{fut}} = [\, w_b[1], \ldots, w_b[H] \,]^\top \in \mathbb{R}^H.$$

It is processed through a sequence of n_{blk} convolutional blocks. Each block has a convolution with kernel size k_i, followed by a nonlinear activation and a downsampling step with stride s_i. The overall transformation is defined as

$$\phi : \mathbb{R}^H \to \mathbb{R}^{N_w}, \qquad \mathbf{z}_b = \phi(\mathbf{w}_b^{\text{fut}}),$$

where $\mathbf{z}_b$ denotes the resulting wave embedding.

Operation: The wave sequence is processed by the convolutional blocks as

$$\mathbf{u}^{(i)} = \text{Down}_{s_i}\big(\sigma(\text{Conv}_{k_i}(\mathbf{u}^{(i-1)}))\big), \qquad \mathbf{u}^{(0)} = \mathbf{w}_b^{\text{fut}}.$$

The early blocks extract short-range features, while deeper blocks capture longer-range patterns. Downsampling reduces the temporal resolution approximately as

$$H_i \approx \left\lfloor \frac{H_{i-1}}{s_i} \right\rfloor,$$

forcing the network to aggregate information over progressively wider temporal intervals.

Output Representation: After the final block, the feature map is flattened and projected to an N_w-dimensional vector:

$$\mathbf{z}_b = W_\phi \,\text{Flatten}(\mathbf{u}^{(n_{\text{blk}})}) + \mathbf{b}_\phi.$$

The resulting embedding $\mathbf{z}_b$ provides a compact summary of the future wave signal and serves as conditioning information for the decoder.

5.2.2.3. Decoder LSTM

Role and Purpose: The decoder LSTM predicts the vessel motions over the future horizon H. It integrates the encoded past dynamics and the compressed representation of the future wave conditions, ensuring that the predicted motion sequence reflects both the current vessel state and the upcoming wave forcing.

Inputs and Construction: The decoder receives the wave embedding

$$\mathbf{z}_b \in \mathbb{R}^{N_w},$$

which is broadcast across all timesteps in the prediction horizon:

$$\mathbf{v}_b[t] = \mathbf{z}_b, \qquad t = 1, \ldots, H.$$

The decoder is initialised using the encoder output:

$$(\tilde{\mathbf{h}}_0, \tilde{\mathbf{c}}_0) = (\mathbf{H}_b^{\text{enc}}, \mathbf{C}_b^{\text{enc}}).$$

Operation: At each future timestep, the decoder updates its internal memory using the LSTM recurrence:

$$(\tilde{\mathbf{h}}_t, \tilde{\mathbf{c}}_t) = \mathrm{LSTM}_{\mathrm{dec}}\big(\mathbf{v}_b[t],\, (\tilde{\mathbf{h}}_{t-1}, \tilde{\mathbf{c}}_{t-1})\big), \qquad t = 1, \ldots, H.$$

Within a single forward pass, the decoder generates all future motion estimates of horizon H, making it a direct multi-step predictor.

Output Representation: Each hidden state is mapped to a motion prediction using a linear layer:

$$\widehat{\mathbf{m}}_b[t] = W_{\mathrm{out}}\, \tilde{\mathbf{h}}_t + \mathbf{b}_{\mathrm{out}}.$$

Collecting the predictions for all timesteps gives the future motion trajectory

$$\widehat{\mathbf{m}}_b = [\, \widehat{\mathbf{m}}_b[1], \ldots, \widehat{\mathbf{m}}_b[H] \,].$$

This trajectory depends on both the encoded past behaviour of the vessel and the expected future wave conditions.

For long-horizon forecasting, this module is used recursively in combination with the encoder. After producing the first horizon, the predicted motions are concatenated with the next segment of the wave sequence to form a new “past window.” This updated window is passed again through the encoder to obtain a new memory state, which then initialises the decoder for predicting the next horizon. In this way, the model rolls forward chunk by chunk, using its previous predictions as part of the input for the next prediction interval.

Autocorrelation Function

As the aim of this methodology is to identify a minimal yet representative dataset that covers the full range of environmental conditions and resulting critical response parameters, it is important to understand how predictable the response signals are in time. The segment-based learning approach used in this work attempts to predict the final portion of a three-hour simulation using only an initial time window. Whether this is feasible depends partly on the intrinsic memory of the response signal: if the signal shows little correlation with its past, then forecasting its future evolution becomes difficult for any model.

However, unlike classical time-series forecasting methods that rely solely on past values of the response, the method used in this thesis also incorporates information on the incoming wave. This means that the predictions are not only driven by the signal’s internal memory, but also by the learned relationship between wave excitation and the system’s dynamic response. Even so, the intrinsic memory time remains important: a longer memory time implies a stronger dependence between earlier and later parts of the response, which makes the prediction task more feasible, even when wave information is included.

To assess how well the predictions of the proposed model align with the predictability of the underlying response, a signal analysis is performed. If a response decorrelates rapidly with respect to its own past, the remaining portion of the simulation contains little information about the initial segment alone, which can make prediction more difficult unless the response is strongly and directly driven by the incoming wave signal.

The predictability, or memory, of each response signal is quantified using the autocorrelation function (ACF),

$$R_{yy}(\tau) = \frac{1}{N} \sum_{t=1}^{N-\tau} \frac{(y_t - \bar{y})(y_{t+\tau} - \bar{y})}{\sigma_y^2}, \tag{5.2}$$

which measures how strongly the signal correlates with its past values at a time lag τ. According to the Wiener–Khinchin theorem, the rate at which the ACF decays is governed by the bandwidth of the power spectral density: narrow-band responses show slowly decaying ACFs, whereas broadband responses lose memory much more rapidly [101], [103].

For oscillatory signals, such as vessel motions, the effective memory time τ_m is determined using the standard $1/e$ envelope criterion,

$$|R_{yy}(\tau_m)| \approx e^{-1}, \tag{5.3}$$

which avoids ambiguity caused by zero crossings and provides a robust, frequency-independent measure of signal memory [103].

A long memory time indicates that the future part of the response remains strongly related to the initial segment, supporting the feasibility of segment-based learning. Conversely, a short memory time indicates that the response decorrelates quickly from its own past, meaning that the initial segment contains limited information about the remainder of the signal. In such cases, accurate long-horizon prediction becomes challenging unless the response is strongly and directly driven by the incoming wave elevation, which can provide an alternative source of predictive structure. This analysis therefore supports the methodology by clarifying when a reduced training dataset and a segment-based prediction strategy are likely to perform well, and when the underlying dynamics place inherent limits on the predictability of the full three-hour response across the different environmental conditions considered in this study.

5.2.3. Full Time Series-Based Learning

In this thesis, for full time series–based learning, we propose joint prediction of vessel motions (12 DOF combined for two vessels) and mooring–line tensions (16 lines) driven by measured wave elevation time series and environmental conditions. The model learns the full mapping from wave measurements and conditions to motions and tensions using:

- Convolutional residual feature extractor for wave patterns: a 1D convolutional network with residual connections that processes the raw wave elevation signal and extracts multi-scale features, including local steepness and frequency variations.
- Condition encoder for environmental variables: a small MLP that embeds the environmental parameters into a compact vector so that their influence on vessel and mooring responses is incorporated at every timestep.
- Transformer encoder for temporal context and long-range dependencies: an attention-based sequence model that integrates information across the full time series, enabling the network to capture non-local temporal interactions between wave features, environmental conditions, and evolving vessel responses.
- Attention-based decoders for motions and tensions: two decoder modules—one for motions and one for mooring line tensions—that use trainable vectors to attend to the encoded sequence and generate predictions for all timesteps simultaneously.

Together, these components allow the model to represent the time-dependent interactions between waves, vessel motions, and mooring forces across the full simulation window.

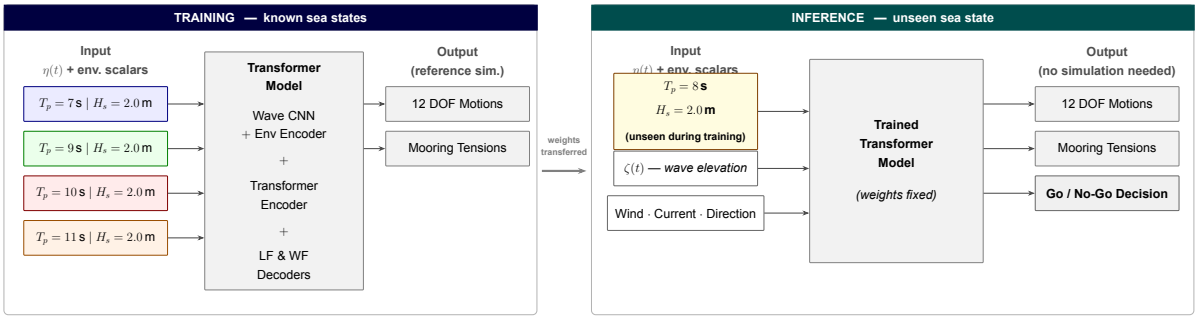


Figure 5.4: High-level overview of the full time series-based learning framework. During training the transformer learns a mapping from wave elevation and environmental conditions to 12 DOF vessel motions and mooring tensions across known sea states ($T_p = 7, 9, 10, 11\,\mathrm{s}$; $H_s = 2.0\,\mathrm{m}$). At inference time only the wave signal and environmental scalars are needed; no reference simulation is required.

In the following sections, the different modules of the suggested model are discussed in further detail.

5.2.3.1. Convolutional Wave Compressor

This part of the model is explained before in Section 5.2.2.2, Convolutional Wave Compressor).

5.2.3.2. Condition Encoding

Role and Purpose. The condition encoder maps the environmental condition vector into a compact representation suitable for the temporal model. Environmental variables such as wave direction, wind speed, and current velocity strongly influence hydrodynamic loads and mooring behaviour. Encoding these conditions ensures that their effects are incorporated consistently across all timesteps.

Inputs and Construction. The environmental condition vector is defined as

$$\mathbf{c} = \left[H_{s,sea},\ T_{p,sea},\ \theta_{sea},\ V_{\mathrm{wind}},\ \theta_{\mathrm{wind}},\ V_{\mathrm{curr}},\ \theta_{\mathrm{curr}}\right]^{\top}.$$

where $H_{s,sea}$ denotes the significant sea wave height, $T_{p,sea}$ the peak sea wave period, θ_{sea} the sea wave direction, U_{wind} the wind velocity, θ_{wind} the wind direction, U_{curr} the current velocity, and θ_{curr} the current direction.

Operation: The condition vector is processed by a small MLP:

$$\mathbf{c}_e = \mathrm{Enc}_{\mathsf{cond}}(\mathbf{c}) \in \mathbb{R}^{N_c}.$$

The MLP learns interactions among the environmental parameters and the outputs and compresses them into an embedding that can be integrated efficiently into the sequential model.

Output Representation: The encoder outputs the environmental condition embedding $\mathbf{c}_e$, which is broadcast across all timesteps and used in the temporal fusion stage.

5.2.3.3. Temporal Fusion and Transformer Encoder

Role and Purpose: The temporal fusion module brings together three sources of information:

- the per-time wave embedding from the wave compressor,
- the global wave embedding,
- the environmental-condition embedding.

These are merged into a single feature sequence and processed by a Transformer encoder, which learns long-range temporal relationships needed to predict vessel motions and mooring-line tensions.

Inputs and Construction: For each timestep n, the model constructs a fused feature vector by concatenating the three inputs:

$$\mathbf{z}_b[n] = \mathrm{concat}\big(\mathbf{w}_b[n],\ g_b,\ \mathbf{c}_{e,b}\big) \in \mathbb{R}^{F}.$$

Here,

- $\mathbf{w}_b[n] \in \mathbb{R}^{N_w}$ is the per-time wave embedding,
- $g_b \in \mathbb{R}^{N_g}$ is the global wave embedding broadcast to all timesteps,
- $\mathbf{c}_{e,b} \in \mathbb{R}^{N_c}$ is the environmental-condition embedding, also broadcast over time.

Thus the fusion dimension is

$$F = N_w + N_g + N_c.$$

Transformer Encoding: The fused sequence $\mathbf{z}_b[n]\}_{n=1}^{T}$ is passed through a Transformer encoder:

$$\mathbf{Z}_b[1:T] = \mathrm{TransformerEnc}\big(\{\mathbf{z}_b[n]\}_{n=1}^{T}\big),$$

where each output vector

$$\mathbf{Z}_b[n] \in \mathbb{R}^{d_{\mathrm{model}}}$$

contains information aggregated from all timesteps via self-attention. The resulting sequence

$$\mathbf{Z}_b = [\mathbf{Z}_b[1], \ldots, \mathbf{Z}_b[T]]$$

acts as the shared memory for both decoder modules.

5.2.3.4. Decoders: Motion and Tension

Role and Purpose. Two decoder modules generate the final outputs:

- the motion decoder, which predicts the 12-DOF vessel motions,
- the tension decoder, which predicts the 16 mooring-line tensions based on the predicted motions and the environmental conditions.

Both decoders use the encoded sequence $\mathbf{Z}_b[1:T]$ as their source of temporal information.

Motion Decoder: Inputs and Construction. The motion decoder introduces one learnable vector for each timestep,

$$E_{\text{motion}} \in \mathbb{R}^{T\times F},$$

which serves as the decoder's input sequence. These vectors allow the decoder to extract motion-related patterns from the encoded features.

Motion Decoder: Operation. The motion decoder processes these timestep vectors together with the encoded sequence:

$$\mathbf{D}_b[1:T] = \text{MotionDecoder}(E_{\text{motion}},\ \text{memory} = \mathbf{Z}_b).$$

Motion Decoder: Output Representation. Each decoded vector is transformed into a 12-component motion prediction, producing the full time series of vessel motions.

Tension Decoder: Inputs and Construction. To predict tensions, the predicted motions are combined with the environmental embedding:

$$\mathbf{mcb}[n] = \text{concat}(\hat{\mathbf{n}}_b[n],\ \mathbf{c}_{e,b}).$$

This sequence forms the input information used by the tension decoder.

Like the motion decoder, the tension decoder introduces one learnable vector per timestep:

$$E_{\text{tension}} \in \mathbb{R}^{T\times F}.$$

Tension Decoder: Operation. The tension decoder processes its timestep vectors together with the combined motion–condition sequence:

$$\mathbf{D}_b^T[1:T] = \text{TensionDecoder}(E_{\text{tension}},\ \text{memory} = \mathbf{mcb}).$$

Tension Decoder: Output Representation: Each decoded vector is transformed into the 16 mooring-line tensions, giving a complete time series of tension predictions for all lines.

Full Time-Series Learning Analysis

The purpose of the full time series based learning is to determine how well the model can reproduce the temporal structure of a full response signal under different combinations of wave height and peak period, and to understand which types of motions require more or less diverse training data. Unlike the segment-based method, the prediction here is not constrained by the intrinsic memory time of the response because the model outputs the entire three-hour time series. Instead, performance depends on how similar the test response behaviour is to the responses represented in the training set and on how well the model learns the relationship between the environmental inputs and the resulting time-series dynamics.

To quantify this similarity, the Pearson correlation coefficient is used [104],

$$\rho = \frac{\sum_t (y_t - \mu_y)(\hat{y}_t - \mu_{\hat{y}})}{\sigma_y \sigma_{\hat{y}}}, \tag{5.4}$$

where y_t is the reference response, $\hat{y}_t$ is the predicted response, and μ and σ denote their respective means and standard deviations. The coefficient ρ measures linear similarity between the two time series by quantifying how their deviations from the mean vary together. The numerator computes the covariance between the signals, while the denominator normalises this covariance by the product of their standard deviations, ensuring that

$$-1 \leq \rho \leq 1.$$

A value close to $+1$ indicates strong positive linear alignment, a value near 0 indicates little or no linear relationship, and a value near -1 indicates strong negative alignment. Because the coefficient is unaffected by amplitude scaling, it is suitable for comparing predicted and reference time series across sea states with different energy levels.
In addition to correlation, simple descriptive statistics are compared between the predicted and reference signals to check amplitude consistency and spread:

- Means ($\mu_{\hat{y}}$ vs. μ_y): assesses overall bias in the predicted response.
- Standard deviations ($\sigma_{\hat{y}}$ vs. σ_y): assesses whether variance (energy) is matched.
- Min/Max envelopes ($\min \hat{y}_t$, $\max \hat{y}_t$ vs. $\min y_t$, $\max y_t$): checks whether the prediction captures the observed amplitude range without systematic clipping or overshoot.

Together with ρ, these statistics provide a compact picture of temporal agreement (phase and shape) and amplitude fidelity (level, spread, and extremes). They also help disentangle cases where two signals are well aligned in time but differ in amplitude, or vice versa.

Here, ρ is used as a data-similarity diagnostic, relating the model's performance on a test case to how similar that test response is to the responses already contained in the training dataset. By analysing ρ alongside the mean, standard deviation, and min/max comparisons across different combinations of, for example, significant sea wave height $H_{s,sea}$, peak sea period $T_{p,sea}$ and sea wave direction, we gain practical insight into what the minimal dataset should contain to ensure good generalisation across the environmental domain.

Overall, the combined use of Pearson correlation and simple statistical comparisons (means, standard deviations, and min/max) provides a clear and practical way to relate training–test similarity to predictive performance, and therefore supports the definition of a minimal set of training simulations for accurate full time-series prediction.

5.3. Evaluation and Methodological Synthesis

All models developed in this thesis follow a unified evaluation protocol to enable a consistent comparison between the SSBL approach and the TSBL approach. Ground truth is provided by MARIN's aNySIM XMF time-domain simulations, treated as authoritative reference data [97]. Each method is trained on a subset of the Generic three-hour simulation dataset and evaluated on an external test set containing unseen combinations of sea states. The methods of evaluation of the predictions are outlined in this subsection.

Quantitative scores: To quantify predictive accuracy, three complementary error measures are used. The coefficient of determination (R^2) [105] is defined as

$$R^2 = 1 - \frac{\sum_{t=1}^{N}(y_t - \hat{y}_t)^2}{\sum_{t=1}^{N}(y_t - \bar{y})^2}, \tag{5.5}$$

and is reported together with the root mean squared error (RMSE) [105],

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{t=1}^{N}(y_t - \hat{y}_t)^2}, \tag{5.6}$$

In Equation 5.5 and Equation 5.6:

- t is the index over samples (e.g. time steps or cases), with $t = 1, \ldots, N$,
- N is the total number of samples in the evaluation set,
- y_t is the ground-truth value of the response at index t,
- $\hat{y}_t$ is the model prediction at index t,
- $\bar{y}$ is the sample mean of the ground-truth values

The R^2 score measures the proportion of variance in the true response that is captured by the model. A value of $R^2 = 1$ indicates perfect agreement between predictions and ground truth; values near zero indicate that the model performs similarly to predicting the sample mean; and negative values indicate performance worse than this trivial baseline. Because it is dimensionless, R^2 provides a convenient measure of how well the model explains the system's variability.

The RMSE provides an absolute measure of prediction error in the physical units of the response. By taking the square root of the average squared deviation, it gives a score that is directly comparable to the scale of the underlying variable. It penalises large errors strongly, making RMSE informative when the model must capture extreme responses or short peaks.

Operability assessment (evaluation): In addition to numerical accuracy, both the SSBL model and the TSBL models are evaluated on their capacity to reproduce operational decisions consistent with the simulation-based reference study. For every test case, predicted responses are subjected to the same operability thresholds used as listed in Table 4.1 [83]. The resulting "operable vs. non-operable" classification is compared to the ground-truth decision derived from aNySIM outputs. This evaluation step verifies whether good regression performance translates into correct go/no-go decisions and highlights typical confusion patterns (false safe vs. false unsafe).

Consistency across methods: The evaluation procedure is identical across the two methodological streams. In the SSBL approach, inputs are restricted to environmental descriptors, and outputs are

maxima/minima of responses relevant for operability. This setup tests how far a reduced representation can support reliable decision-making at low computational cost. In the TSBL approach, full three-hour time traces of waves and responses are used as inputs, and the model predicts complete response traces. Besides the R^2 and RMSE, we also check how well the predictions capture the time-varying behaviour of the aNySIM results, including both low- and wave-frequency components.

Evaluation and contribution to the research question: Taken together, the evaluation procedures in this methodology are structured to investigate how the two modelling approaches address different aspects of the central research question:

How can machine learning-based methods support and accelerate operability assessment for side-by-side offloading?

The SSBL models are used to investigate whether operability screening can be performed directly from environmental descriptors and to assess how far such simplified representations may reduce the computational effort typically required by simulations during planning.
The TSBL models are used to assess whether data driven models can reproduce full response trajectories with enough fidelity for physics aware decision making. This includes capturing the timing of extreme events and the interactions between different degrees of freedom, while also reducing the dependence on large scale time domain simulations.
All models are evaluated on a test set containing unseen sea-state conditions to examine their ability to generalise to new conditions and to identify the conditions under which machine-learning alternatives can provide reliable predictive performance. The operability assessment (evaluation) quantifies the agreement between model-derived go/no-go classifications and those obtained from aNySIM simulations, including the analysis of characteristic confusion patterns (false safe versus false unsafe). Together, these methodological components provide a structured framework for determining how data-driven models disrupt conventional simulation-based operability assessment workflows, while preserving the levels of reliability required in multi-body offshore operations.

6

Execution

This chapter presents the practical implementation of the methods introduced in chapter 5. Whereas the methodology outlines the conceptual framework for both Summary-Statistics-Based Learning (SSBL) and Time-Series-Based Learning (TSBL), the present chapter details the decisions, design iterations, and processing steps that enabled their execution. The aim is not only to document the workflow but also to ensure reproducibility and transparency for future research.

The execution consists of three main components. section 6.1 describes the preprocessing of summary and time-series datasets, including feature engineering and data cleaning procedures. subsection 6.1.2 outlines the development of the SSBL pipeline, which is based on summary environmental and hydrodynamic descriptors. section 6.2 details the TSBL approaches, consisting of (i) motion-signal decomposition, (ii) segment-based sequence-to-sequence learning, and (iii) full time-series learning. These three TSBL variants differ in temporal resolution and data requirements but share several processing steps and optimization procedures. Finally, subsubsection 6.2.4.3 summarises cross-cutting tools such as hyperparameter optimization and validation strategies.

6.1. Summary-Statistics-Based Learning (SSBL)

This section details how the summary-statistics models introduced in chapter 5 were implemented in practice. It outlines the preprocessing of environmental input features, the partitioning strategy for training and evaluation, and the design of the dual-head MLP used to predict the operability-relevant maxima from these inputs. Together, these steps form the complete workflow used to assess the capability and limitations of the SSBL approach.

6.1.1. Data and Preprocessing

The SSBL models use MARIN's simulation outputs, which include: (i) summary statistics (min/max over 3 h per sea state) of vessel motions and interface loads, and (ii) the corresponding environmental descriptors (sea and swell H_s/T_p, wind and current magnitudes and directions).

Directional variables (sea, swell, wind, current directions) are encoded via sine/cosine pairs to preserve circularity and avoid discontinuities at $0/360°$. All other inputs are kept as observed scalars.

Input Features: The input vector comprises the environmental parameters only: sea and swell H_s/T_p, wind speed and direction, and surface current speed and direction; yielding eight scalars. The four directions are decomposed into sine/cosine, giving a 16-dimensional input. This choice keeps the predictors purely exogenous, i.e., no response information is allowed to leak into the inputs.

Output Targets: The model predicts the subset of MARIN outputs that governs operability. Specifically, we use:

- Roll of FLNG and LNGC (min and max over 3 h);
- Relative manifold motions in surge and sway (min and max);
- Maximum line tension across all 16 SBS mooring lines (single extreme per run);
- Maximum fender compression force across all five fenders (single extreme per run).

6.1.2. Training, Validation, and Test Strategy

The dataset is partitioned using a two-stage strategy that combines k-means clustering for representative sampling with stratified splitting for unbiased evaluation.

Stage 1: Representative K-Means Sampling K-means clustering partitions the input feature space into k clusters by iteratively assigning points to the nearest centroid and updating centroids as cluster means [106]. This identifies representative regions of the environmental and operational parameter space.

The optimal cluster count k^* is selected by maximizing the silhouette score:

$$k^* = \arg \max_{k \in [k_{\min}, k_{\max}]} \text{Silhouette}(k) \tag{6.1}$$

The silhouette score measures cluster quality by evaluating both intra-cluster cohesion (how close points are to their assigned cluster) and inter-cluster separation (how distant clusters are from each other). Higher values indicate better-defined clusters. The score is computed on a subsample of 5,000 rows to reduce computational cost.

Once k^* is determined, a representative sample is constructed by proportionally sampling rows from each cluster. If cluster i contains n_i rows, then $\frac{n_i}{N} \times S$ rows are selected from it, where N is the dataset size and S is the target sample size (50% of the full dataset). This ensures the training set captures the same distribution of operating conditions as the full dataset.

Stage 2: Internal Train/Validation/Test Split The representative sample is further split into three disjoint subsets:

- **Training set:** 42.5% of full dataset (85% of representative sample)
- **Validation set:** 7.5% of full dataset (15% of representative sample)
- **Test set:** 50% of full dataset (remaining unsampled data)

The training set optimizes model weights, the validation set enables hyperparameter tuning and early stopping, and the test set provides an unbiased estimate of generalization performance.

6.1.3. Multi-Head Multilayer Perceptron (MLP) Architecture

The architecture is designed as a dual-head neural network, with separate predictive heads for forces (line tensions and fender forces) and motions (vessel orientations and manifold distances). This separation allows the model to specialize in predicting fundamentally different physical phenomena.

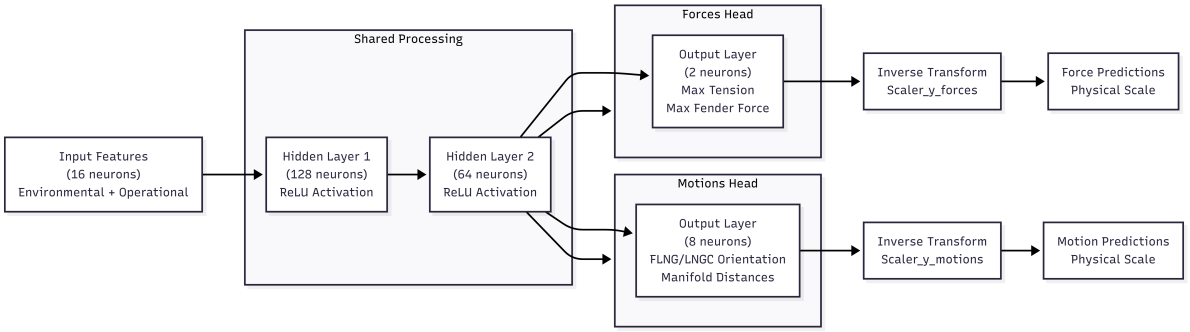


Figure 6.1: Overview MLP Architecture

Network Structure: Each head is a fully connected feedforward MLP with the following specifications:

- **Input Layer:** 16 neurons, corresponding to the 16 engineered input features (including sin/cos directional decompositions)
- **Hidden Layers:** Configuration of [128, 64] neurons per layer
- **Output Layer:** Forces head outputs 2 neurons (max line tension, max fender force); motions head outputs 8 neurons (4 orientation min/max pairs, 4 manifold distance min/max pairs)
- **Activation Function:** ReLU (Rectified Linear Unit) for hidden layers, linear activation for output layer
- **Solver:** Adam optimizer with learning rate $lr = 10^{-3}$ and L2 regularization coefficient $\alpha = 10^{-4}$

Loss Function and Regularization: The model minimizes Mean Squared Error (MSE) on scaled target variables, with early stopping applied to prevent overfitting. Specifically:

$$\mathcal{L} = \frac{1}{N}\sum_{i=1}^{N}\|\hat{\mathbf{y}}_i - \mathbf{y}_i\|_2^2 + \alpha\|\mathbf{W}\|_2^2 \tag{6.2}$$

where $\mathbf{W}$ represents the network weights and $\alpha = 10^{-4}$ is the L2 penalty coefficient.

6.2. Time Series-Based Learning (TSBL)

This section describes how time-series methods are used to predict vessel motions and mooring loads. First, responses are split into low-frequency (LF) and wave-frequency (WF) components with Butterworth filters so each model focuses on the behaviour it is designed to capture. We then present two routes: a segment-based sequence-to-sequence approach for short-horizon forecasts from overlapping windows, and a full time series approach that processes the entire record with shared inputs and dedicated LF and WF branches.

6.2.1. Data preprocessing

The hydrodynamic simulations of MARIN use initialization periods before the vessel reaches dynamic equilibrium with environmental forcing. Althoug every seastate is simulated for 3.5 hours, only the last 3 hours are used for operablity assessment [83].

To ensure that only realistic, steady-state operational behaviour is retained for model training, the preprocessing pipeline systematically removes initialization and boundary transients:

Transient trimming: To ensure that the learning models operate on physically meaningful steady-state behaviour, the initial and final transient regions of each simulation are removed before further processing. The first 1800 timesteps are discarded, corresponding to 30 minutes at the 1 Hz sampling rate, in order to eliminate the start-up phase in which the hydrodynamic response has not yet reached a stationary regime. Similarly, the last 30 timesteps (30 s) are excluded to avoid boundary effects that arise when the environmental forcing is terminated at the end of the simulation. These final seconds often exhibit numerical artefacts or unphysical decay behaviour unrelated to true system dynamics. After removing both regions, the resulting usable time series consists of 10 772 timesteps.

Feature Selection and Engineering

As discussed in Appendix B, the original generic dataset contains cases with sea waves, swell waves, and combined sea–swell conditions. Because the point wave elevation record represents a single free-surface elevation time series, the combined sea–swell cases cannot be decomposed into two separate wave signals. To maintain physical consistency and to keep the computational effort within reasonable limits, only the sea-wave cases are used for the TSBL modelling. After filtering the dataset accordingly, a total of 8317 distinct sea states remain.

The TSBL models operate on all seven environmental (Table 4.2 & Table 4.3) parameters together with the point wave elevation signal and the local water depth h. Including h enables the trained models to expand to new geographical locations with different bathymetric conditions.

The environmental input vector used in this work is therefore defined as:

$$\{\, H_{\mathrm{s,sea}},\ T_{\mathrm{p,sea}},\ \theta_{\mathrm{sea}},\ V_{\mathrm{wind}},\ \theta_{\mathrm{wind}},\ V_{\mathrm{cur}},\ \theta_{\mathrm{cur}},\ h \,\}.$$

Directional environmental parameters are encoded directly using their angular values. For each directional variable θ_{env}, two features are created:

$$\theta_{\mathrm{env}} \ \longrightarrow\ \big(\cos\theta_{\mathrm{env}},\ \sin\theta_{\mathrm{env}}\big).$$

This provides a smooth, continuous representation of all wave, swell, wind, and current directions and avoids discontinuities at the $0^\circ \leftrightarrow 360^\circ$ boundary.

Output Target Definition

The TSBL models predict a total of 28 outputs:

- 12 vessel motions (6-DOF for the FLNG and 6-DOF for the LNGC),
- 16 mooring line tensions (Lines 1–16).

Table 6.1: Model output channels: vessel motions (FLNG and LNGC) and mooring line tensions.

Output Group	Components	Unit
FLNG motions (6 DOF)	Surge, Sway, Heave, Roll, Pitch, Yaw	m / rad
LNGC motions (6 DOF)	Surge, Sway, Heave, Roll, Pitch, Yaw	m / rad
Mooring line tensions	Lines 1–16	kN

Excluded outputs: Fender compression forces are available in the simulation data but are not included as learning targets. In all MARIN simulations, whenever a fender exceeded its allowable limit, a different operability criterion exceeded its limit earlier. Including them would therefore introduce unnecessary model complexity without improving predictive performance.

Relative Motions are not predicted directly because they are deterministic transformations of absolute motions and are computed post-hoc when needed.

6.2.2. Butterworth Signal Decomposition

This preprocessing step separates each motion and tension time series into LF and WF components using Butterworth Infinite Impulse Response (IIR) filters. The decomposition enables the LF and WF learning models to be trained on frequency bands that match their physical behaviour: slowly varying, static responses for LF, and phase-dependent oscillatory responses for WF. A cutoff frequency of $0.03\,\mathrm{Hz}$ is chosen so that long-period swell motions would fall within the WF band rather than the LF band. Although swell waves are not included in the TSBL models developed in this thesis, this cutoff remains suitable for future extensions of the method where swell–induced oscillations may need to be represented within the WF band.

The environmental conditions corresponding to the frequency spectra in Figure 6.2 are listed in Table 6.2. Figure 6.2 shows frequency spectra of representative FLNG surge and heave motions with the 0.03 Hz cutoff used for LF-WF separation, which reflect the expected distinction between wave-frequency and low-frequency responses discussed in section 4.7: a dominant WF peak associated with first-order excitation in heave, and a much smaller LF band for surge caused by slow-drift second-order forcing.

Table 6.2: Environmental conditions corresponding to the frequency spectra in Figure 6.2.

Wind vel.	Wind dir.	$H_{s,\mathrm{sea}}$	$T_{\mathrm{p,sea}}$	Sea dir.	Curr. vel.	Curr. dir.
1 m/s	90°	2 m	8 s	180°	0.12 m/s	90°

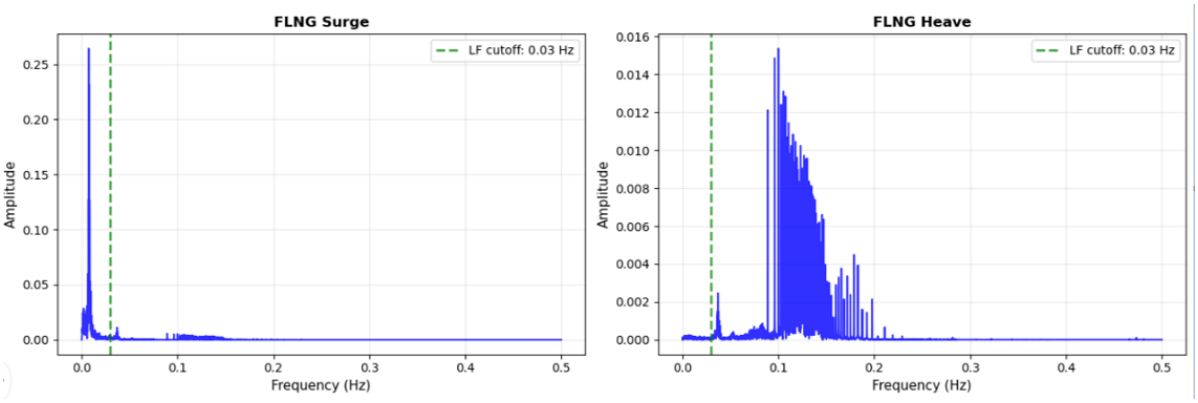


Figure 6.2: Frequency spectra of FLNG surge (left) and heave (right) with the 0.03 Hz low-frequency cutoff used for LF-WF decomposition.

A Butterworth filter is used because it provides a maximally flat passband [107] without ripples or selective amplification. Forward–backward filtering is applied to avoid phase distortion so that LF and WF components remain time aligned.

The decomposition consists of the following steps:

1. Butterworth filtering: A low-pass Butterworth filter extracts the LF component. The WF component is obtained as the residual:
$$x_{\mathrm{WF}}(t) = x(t) - x_{\mathrm{LF}}(t).$$
2. Forward–backward filtering: Filtering is performed to remove phase shift and preserve timing.
3. Reconstruction check: For each signal, LF + WF is compared to the original time series and the RMS error is computed to verify numerical consistency.

Figure 6.4 illustrates the LF and WF components of an FLNG surge signal and demonstrates that their sum reconstructs the original record.

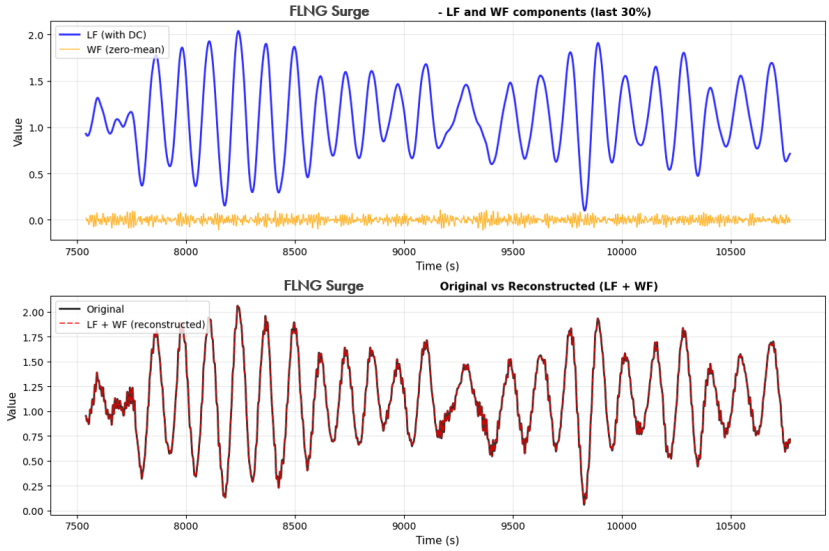


Figure 6.3: Example LF–WF decomposition for FLNG surge. Top: LF component (with mean) and zero-mean WF component. Bottom: original signal and reconstructed LF + WF.

6.2.3. Segment-Based Sequence-to-Sequence Learning

As discussed in chapter 5, this work employs a segment-based sequence-to-sequence architecture to forecast multi-degree-of-freedom (DOF) vessel motion over a finite prediction horizon by combining past motion and wave information with a compact representation of upcoming wave forcing.

In the segment-based sequence-to-sequence model, we predicted only the vessel motions (12 DOFs across both vessels). Tensions were intentionally excluded at this stage, since reliable tension modelling requires accurate motion predictions first.

All experiments use 3-hour simulations with total duration $T_{\text{total}} = 10{,}772\,\text{s}$. A chronological split is applied so that the first 70% is used for training ($T_{\text{train}} = 7{,}540\,\text{s}$) and the final 30% is reserved for testing ($T_{\text{test}} = 3{,}232\,\text{s}$), preventing leakage between regimes.

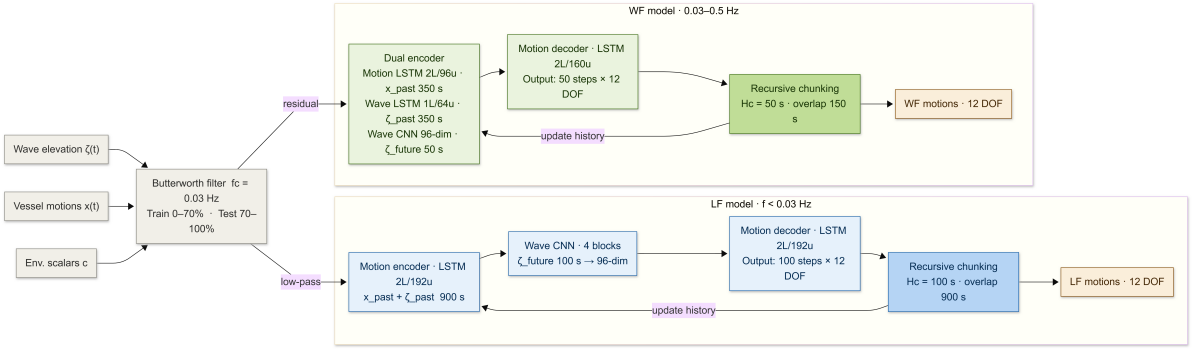


Figure 6.4: Overview of the segment-based sequence-to-sequence architecture showing the WF model (upper branch, $H_c = 50$ s, overlap 150 s) and LF model (lower branch, $H_c = 100$ steps, overlap 900 s), both using dual LSTM encoders with recursive chunking.

Both LF and WF models share the same sequence-to-sequence template but differ in three aspects: (i) feature engineering, (ii) encoder topology, and (iii) segment dimensions (memory window, prediction horizon, recursive rollout).

6.2.3.1. Low-Frequency Model

Low-frequency responses in the range 0–$0.03\,\text{Hz}$ evolve slowly and are dominated by equilibrium among environmental loads and mooring restoring forces. The model therefore uses a long memory and is less sensitive to high-frequency content.

Data representation:

$$\begin{aligned} \text{Past motion (900 s)}: \ & \mathbf{x}_{\text{past}} \in \mathbb{R}^{900\times 12} \quad \text{(two vessels, 6 DOF each)}, \\ \text{Past wave elevation (900 s)}: \ & \mathbf{w}_{\text{past}} \in \mathbb{R}^{900\times 1}, \\ \text{Future wave elevation (100 s)}: \ & \mathbf{w}_{\text{future}} \in \mathbb{R}^{100\times 1}. \end{aligned}$$

The LF model was assigned a longer input window, because low-frequency responses originate from second-order drift forces and long-period coupled modes whose evolution spans several hundreds of seconds Appendix B. Capturing these dynamics requires a temporal context larger than that needed for wave-frequency prediction, making the extended 900 s sequence essential for representing the slow build-up and decay of LF motion components.

What the layers do:

- *LSTM encoder (2 layers, 192 units each).* Running memory that reads the $900\,\mathrm{s}$ history and remembers slow trends. Two stacked layers allow the second layer to work on a cleaner summary produced by the first, improving long-range recall.
- *Wave CNN (4 conv blocks, then a 96-number summary).* Small temporal filters and pooling compress the $100\,\mathrm{s}$ future elevation into a short vector that describes the overall shape and strength of the incoming wave group rather than every sample.
- *Decoder LSTM (2 layers, 192 units each) + linear head.* Starting from the encoder's memory and given the wave summary at each step, the decoder writes out the next 100 predicted time steps. The final linear layer simply maps internal numbers to the 12 motion outputs.

Recursive chunking: For the $3{,}232\,\mathrm{s}$ test sequence, predictions are generated in $H_c = 100\,\mathrm{s}$ chunks. After each chunk, the motion history is updated to keep the most recent $900\,\mathrm{s}$ (observed + predicted).

Critically, it uses no validation set. Instead:

- Training region: 0–70% of time series (windows of 1000 steps with 90% past / 10% future)
- Test region: 70–100% (entirely unseen; predicted recursively in 100-step chunks)
- stopping: Based on training loss plateau (no information leakage from test set)
- Recursive prediction: Chunks are overlapped (900 steps) to maintain temporal continuity and prevent drift in long-horizon prediction

This strategy ensures a clean temporal split: the model has never seen any data from the test region during training, making performance assessment on 70–100% truly representative of generalization to unseen future conditions.

Training Configuration and Loss: The LF predictor employs an encoder-decoder architecture trained with recursive prediction to handle long-horizon forecasting without data leakage.

Table 6.3: Low-frequency model training hyperparameters (encoder–decoder with recursive rollout).

Setting	Value / Rationale
Epochs	300 (sufficient for encoder–decoder convergence)
Batch size	32 (stable batch statistics for the available dataset)
Learning rate	10^{-3} (standard Adam starting point)
Weight decay	10^{-4} (light L2 regularisation)
Early stopping	Patience $= 50$ epochs on training-loss plateau
LR schedule	ReduceLROnPlateau, factor 0.5, patience $= 5$ epochs
Gradient clipping	Max-norm $= 1.0$ (prevents exploding gradients)

Training Configuration: The low-frequency model is trained using a recursive prediction approach. This means that during testing the model repeatedly uses its own previous predictions as input. Because of this, the entire test set must remain completely unseen during training. As a result, the optimisation process and early stopping monitor the training loss only, rather than a separate validation loss.

The model is trained with the Adam optimiser and a learning rate of 10^{-3}, which allows the network to make fast progress at the start of training. A small weight decay term (10^{-4}) is included to reduce overfitting by discouraging the model from relying too strongly on any single set of weights. If the training loss stops improving for five epochs, the learning rate is reduced by half to help the model settle into a better minimum. If there is still no improvement after a longer period (50 epochs), training is stopped to avoid wasting computation.

Finally, gradient clipping (using a max-norm of 1.0) is applied to prevent occasional large parameter updates. This is especially important for recursive prediction, where unstable updates can quickly accumulate and cause the forecast to drift or explode over long horizons.

Optimizer and scheduling: The LF model is trained using the Adam optimizer with adaptive learning-rate scheduling. Since recursive prediction requires the entire test set to remain completely unseen, the optimisation process monitors the training loss directly. A `ReduceLROnPlateau` scheduler decreases the learning rate by a factor of 0.5 whenever the training loss fails to improve for five consecutive epochs. This allows the optimiser to take progressively smaller and more precise steps as convergence is approached, ensuring stable refinement of the solution trajectory.

Loss Function: The model is trained using a simple Mean Squared Error (MSE) loss. For each batch, the loss measures how far the predictions deviate from the true motion values. It is defined as

$$\mathcal{L} = \frac{1}{BTD} \sum_{i=1}^{BTD} (\hat{y}_i - y_i)^2 , \tag{6.3}$$

where

- B is the batch size (number of samples in one batch),
- T is the number of predicted time steps per sample,
- D is the number of output channels (the 12 motion DOFs),
- y_i is the true value at position i in the flattened batch–time–channel tensor,
- $\hat{y}_i$ is the corresponding model prediction.

Using the flattened index i keeps the notation compact while still averaging over all samples, time steps, and output channels.

6.2.3.2. Wave-Frequency Model

WF responses in the range 0.03–$0.5\,\mathrm{Hz}$ oscillate at the incident wave period and have no meaningful static equilibrium. Hydrodynamic excitation scales with wave acceleration, motivating explicit physics-informed features in the WF model.

Data representation and features: Each wave sample is represented as

$$\boldsymbol{\zeta}[t] = \left[\zeta(t),\, \dot{\zeta}(t),\, \ddot{\zeta}(t)\right]^\top \in \mathbb{R}^3.$$

where elevation, velocity, and acceleration provide the relevant kinematic information.

The model inputs are:

$$\begin{aligned}\text{Past motion (350 s)} &: \mathbf{x}_{\text{past}} \in \mathbb{R}^{350\times 12},\\ \text{Past wave features (350 s)} &: \mathbf{\Phi}_{\text{past}} \in \mathbb{R}^{350\times 3},\\ \text{Future wave features (50 s)} &: \mathbf{\Phi}_{\text{future}} \in \mathbb{R}^{50\times 3}.\end{aligned}$$

These sequence lengths match the characteristic time scales of WF dynamics: a 350 s history captures wave groups and immediate vessel response, while a 50 s future segment provides the short prediction horizon.

Model architecture: The WF model uses a multi-input encoder–decoder structure:

- *Motion encoder (LSTM, 2 layers, 96 units).* Learns the vessel's natural WF behaviour (inertia, damping, restoring).
- *Wave encoder (LSTM, 1 layer, 64 units).* Extracts the structure and phase of the incoming wave signal.
- *Wave CNN (same blocks as LF, 3 channels).* Processes elevation, velocity, and acceleration through temporal convolutions and downsampling to obtain a compact embedding of the future 50 s wave segment.
- *Decoder (LSTM, 2 layers, 160 units) + linear head.* Combines the motion encoder state, wave encoder state, and CNN wave summary to generate the next 50 WF motion steps across all 12 DOFs.

This dual-encoder setup mirrors the physical decomposition of "system + forcing".

Recursive chunking: During testing, the model predicts in segments of

$$H_c = 50\ \text{s},$$

with $150\ \text{s}$ overlap. This improves stability for oscillatory signals. After each chunk, the motion history is updated with both observed and predicted values.

Table 6.4: WF hyperparameters.

Setting	Value
Epochs	400
Batch size	32
Learning rate	1.5×10^{-3}
Weight decay	10^{-4}
Warmup	10 epochs (no LR changes)
Early stopping	Patience $= 50$ epochs
LR schedule	ReduceLROnPlateau (factor 0.5, patience $= 12$)
Gradient clipping	Max-norm $= 1.0$

Training configuration. The WF predictor employs the multi-input LSTM encoder-decoder architecture described above. Training prioritises amplitude and phase matching over smoothness, reflecting

the physics of direct wave-driven oscillations. All optimisation monitors the training loss only, because recursive prediction requires the full test region to remain unseen.

Loss function:

$$\mathcal{L} = 0.60\,\mathcal{L}_{\mathrm{MSE}} + 0.25\,\mathcal{L}_{\mathrm{vel}} + 0.15\,\mathcal{L}_{\mathrm{accel}}. \tag{6.4}$$

- MSE loss (60%): captures the main amplitude and phase of the WF signal.
- Velocity loss (25%): constrains the first derivative for realistic phase behaviour.
- Acceleration loss (15%): improves the frequency content of the prediction.

Rotational DOFs are weighted twice the translational ones. Early stopping monitors training loss only, since recursive prediction requires unseen test data.

Design rationale: The WF model uses a higher learning rate (1.5×10^{-3} vs. 10^{-3} for LF) because WF response scales almost directly with wave forcing, allowing faster convergence. A 10-epoch warmup stabilises early updates under high-frequency noise.

6.2.4. Full Time Series Based Learning

As discussed in chapter 5, the full time series-based learning of the model learns the full mapping from wave measurements and conditions to motions and tensions full time series signal. This section specifies the combined LF–WF prediction system that operates on the full $10{,}772\,\mathrm{s}$ time series. Both branches share input processing and then diverge into LF-specific and WF-specific components. Below the full workflow of the LF and WF model combined is displayed in Figure 6.5 .

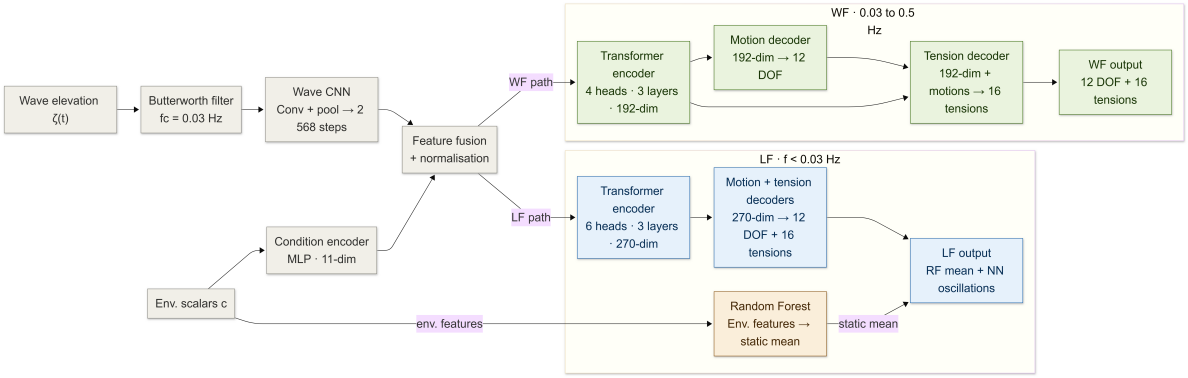


Figure 6.5: Overview of the full time series learning architecture.

In this subsection the models are displayed as how they are structured in the seperate analysis in chapter 7. For chapter 8, the

All inputs are normalized via channel-wise z-score standardization using statistics computed exclusively from training data preventing leakage.

Wave CNN feature extraction:

- *Convolution layer.* Slides a short window along the time axis to detect patterns such as wave crests and groups.
- *Pooling layer.* Downsamples by keeping the largest response in each small region, which reduces noise and shortens the sequence.
- *Projection (1×1).* A simple reshaping step that compresses many channels into a compact vector per time step.

Environmental encoding and fusion: Environmental conditions are encoded in 11 dimensions (8 scalars + 3 directional pairs). Wave CNN embeddings and environment vectors are fused. Normalization uses a StandardScaler fitted on the training split only.

Transformer Encoder:

- *Embedding dimension.* The number of values carried per time step (e.g., 270 for LF). Larger values hold more detail but are heavier to train.
- *Attention heads.* Independent “views” that compare one time step to all others to find what matters. More heads allow the model to focus on different patterns (e.g., long waves vs short bursts) at the same time.
- *Feed-forward layer (4× rule).* A small two-layer network applied at each time step to mix and refine information; its inner size is set to four times the embedding by design.

- *Dropout and layer norm.* Standard stabilisation steps: dropout reduces overfitting; layer normalisation keeps scales consistent.

6.2.4.1. Low-Frequency Model

To ensure physical interpretability, the LF model predicts its output as the sum of:

$$\hat{\mathbf{y}}_{\mathrm{LF}}(t) = \underbrace{\mathbf{y}_{\mathrm{RF}}}_{\text{static mean}} + \underbrace{\mathbf{y}_{\mathrm{NN}}(t)}_{\text{zero-mean oscillations}} .$$

The mean component is provided by a Random Forest trained solely on environmental features, producing a stable equilibrium offset representative of low-frequency drift. On top of this baseline, the neural decoders reconstruct the oscillatory LF variations captured by the Transformer. This division of responsibilities mirrors the physical composition of LF responses: a dominant drift term and a slow-modulation component.

The LF branch is designed to represent the slowly varying drift-driven behaviour of the system. After the shared encoder, two specialised decoder modules operate on the fused representation. These modules are distinct in purpose but complementary in reconstructing the LF part of the response.

Four additional statistical features are derived as extra input:

$$\hat{H}_s = 4\sqrt{\langle \zeta^2 \rangle_{\text{window}}} \quad \text{(significant wave height proxy)} \tag{6.5}$$

$$F_{\text{2nd}} = \hat{H}_s^2 \quad \text{(second-order forcing)} \tag{6.6}$$

$$F_w = V_{wind}, \quad F_c = V_{current} \quad \text{(wind and current magnitudes)} \tag{6.7}$$

—

Motion decoder:

- *Structure.* Deeper recurrent stack with residual links for long-horizon stability.
- *Continuity.* Velocity-based penalty encourages smooth LF evolution.
- *Peak timing.* Acceleration correlation reduces drift in slow-drift peak locations.

Tension decoder:

- *Structure.* Shallow network suited to slowly varying LF load behaviour.
- *Safety bias.* Stronger penalty on underprediction to avoid false-safe outcomes.
- *Smoothness.* Peak-correlation term suppresses unphysical sharp force variations.

LF loss formulation: The LF loss combines three contributions:

$$\mathcal{L}_{\mathrm{LF}} = \mathcal{L}_{\mathrm{MSE}} + 0.1\,\mathcal{L}_{\mathrm{vel}} + \mathcal{L}_{\mathrm{tens}}.$$

The weighted MSE applies higher weights to rotational motions; the velocity term enforces smoothness; and the tension loss contains the strong asymmetric penalty discouraging underprediction. Together, these terms guide the LF branch toward physically consistent slow-drift behaviour.

SS

6.2.4.2. Wave-Frequency Model

The WF branch reconstructs the short-period, wave-induced component of the vessel response. It operates on the same fused wave–environment representation as the LF model, but focuses on capturing rapid oscillations, phase alignment, and sharp motion features characteristic of WF dynamics.

Motion decoder:

- *Structure.* UNet-style decoder integrating Transformer embeddings with skip connections from the wave CNN.
- *Wave detail retention.* Skip connections ensure that local wave shape, slopes, and crests influence WF motion prediction.
- *Activation.* All convolutional blocks use batch normalisation and ReLU to maintain stable high-frequency reconstruction.

Motion-conditioned tension decoder:

- *Design.* Receives both Transformer features and the predicted WF motions.
- *Physical consistency.* Tensions are conditioned directly on motions, reflecting the fact that WF loads arise from instantaneous kinematics.
- *Smooth reconstruction.* The decoder structure suppresses artefacts while preserving sharp peaks associated with wave impacts.

WF loss formulation: The WF loss is defined as:

$$\mathcal{L}_{\mathrm{WF}} = \mathcal{L}_{\mathrm{MSE}} + 0.4\,\mathcal{L}_{\mathrm{tens}},$$

with rotational DOFs weighted twice the translational DOFs. The tension loss applies a sixfold penalty to underprediction to reduce the risk of false-safe outcomes. No velocity-based term is included, as WF dynamics naturally contain rapid oscillations that should not be overly smoothed.

6.2.4.3. Final Model Development

To configure the final LF and WF models for application to the complete dataset, a dedicated hyperparameter optimization stage was carried out using Optuna's Tree-structured Parzen Estimator (TPE) [108]. For each trial, the selected training data were split into an 85% training portion and a 15% validation portion, ensuring a consistent and unbiased evaluation of candidate configurations. Early stopping monitored the validation loss to prevent overfitting and to identify the most stable parameter sets.

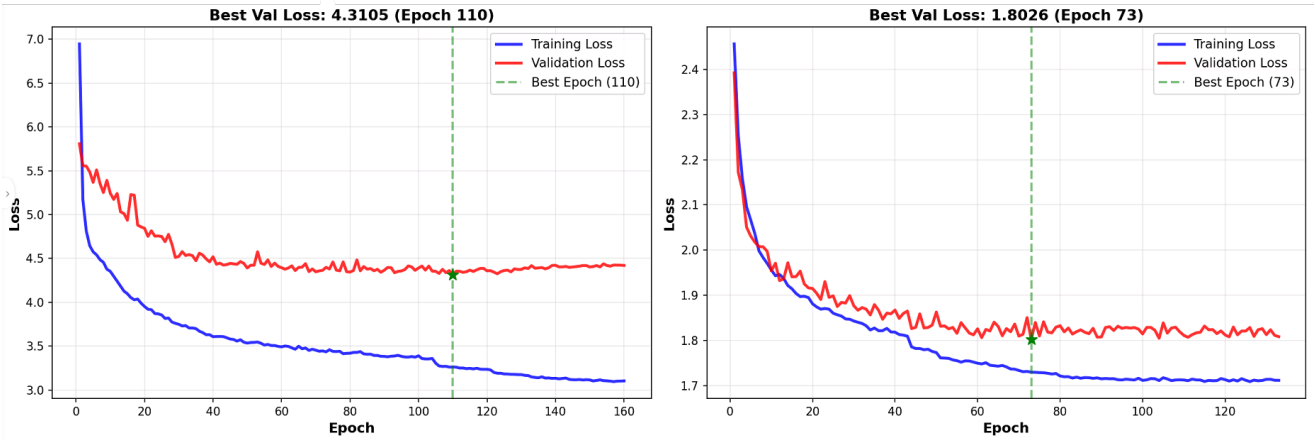


Figure 6.6: Training and validation loss curves for the LF (left) and WF (right) models. The dashed line indicates the epoch with the lowest validation loss for each model. Both curves illustrate the convergence behaviour of the models under the selected optimization configuration.

The loss curves in Figure 6.6 summarise the behaviour of the LF and WF models throughout training. The LF model converges more slowly and exhibits a notable gap between training and validation loss, reflecting the difficulty of learning slow-drift behaviour from limited data. In contrast, the WF model shows smoother convergence, tighter alignment between training and validation curves, and an earlier stabilisation point, indicating a more stable optimisation landscape. These models were deployed in chapter 8.

7

Preliminary Results Analysis

This chapter presents the results of the machine-learning models developed in this thesis for operability assessment of side-by-side offloading. Both summary-statistics-based learning (SSBL) and time-series-based learning (TSBL) approaches are evaluated.

7.1. SSBL Models

For the SSBL models, two distinct machine learning techniques are implemented: Random Forest (RF) and a Multi-Layer Perceptron (MLP). These approaches represent different model architectures, each providing complementary insights into how machine learning captures operability-relevant relationships from reduced summary simulation outputs.

Each model is trained using increasing fractions of the full dataset: 1%, 5%, 10%, 20%, 50%, and 80% of the total generic dataset including sea waves and swell waves, as presented in chapter 4. This setup allows a direct evaluation of data efficiency, showing how predictive skill evolves as more training data becomes available. Across all methods, performance improves with data quantity, but even small subsets provide surprisingly strong predictive accuracy. Below, the RF and MLP results are presented in detail.

7.1.1. Random Forest Results

The RF results in Fig. 7.1 show a clear and consistent improvement as the training fraction increases. On the scaled RMSE plot, errors decrease rapidly between the 1% and 10% splits, after which the curves flatten, indicating diminishing returns once a moderate amount of data is available. Variables such as line tension and maximum surge exhibit the highest absolute RMSE values, reflecting their larger dynamic range, but their relative decrease with additional data mirrors that of the other parameters.

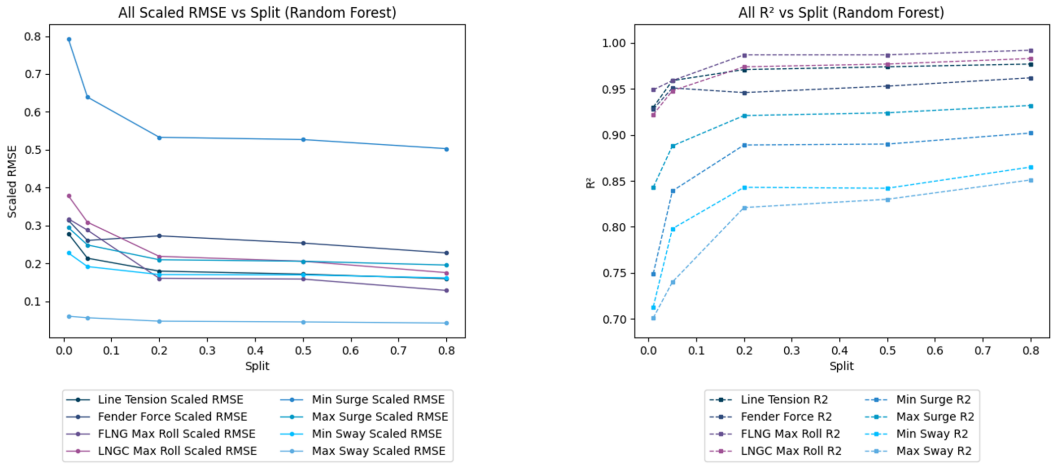


Figure 7.1: Random Forest performance: scaled RMSE (left) and R^2 (right) across training-data splits.

The R^2 curves reveal complementary behaviour. Even with only 5–10% of the dataset, most parameters already reach $R^2 > 0.9$. Roll- and sway-related metrics achieve values above 0.95 early on and remain stable as more data is added. At 80% training data, nearly all variables reach R^2 values between 0.95 and 0.99. This shows that Random Forest is highly data-efficient and captures nonlinear relationships in the summary statistics with limited training data.

7.1.2. MLP Results

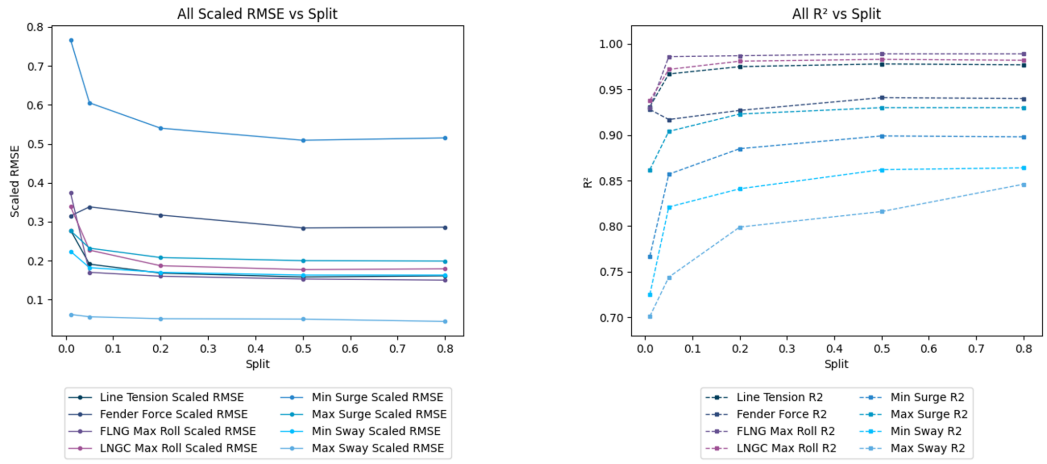


Figure 7.2: MLP performance: scaled RMSE (left) and R^2 (right) across training-data splits.

The MLP model shows similar learning behaviour to the Random Forest, but with slightly lower stability at very small training fractions. As seen in Fig. 7.2, scaled RMSE decreases sharply between the 1% and 10% splits, while the R^2 score improves from values around 0.75–0.85 at 1% to above 0.9 once 10% of the dataset is used. Beyond the 20% split, the MLP converges to performance levels comparable to RF, with most parameters achieving R^2 values above 0.95.

Across all training splits, the spread between the best and worst performing variables remains small, indicating that the MLP captures cross-variable relationships consistently. Once sufficient data is available (20% or more), the MLP matches the Random Forest in both error reduction and correlation

strength.

7.1.3. SSBL Operability Assessment

To evaluate whether the summary-based MLP model can support SSBL operability assessment, the key operability-relevant parameters are examined: roll (orientation limits), relative surge and sway (manifold clearance), and line tension and fender loads (mooring and contact safety). The following figures show the

performance for each class of parameters, together with MARIN's operability limits.

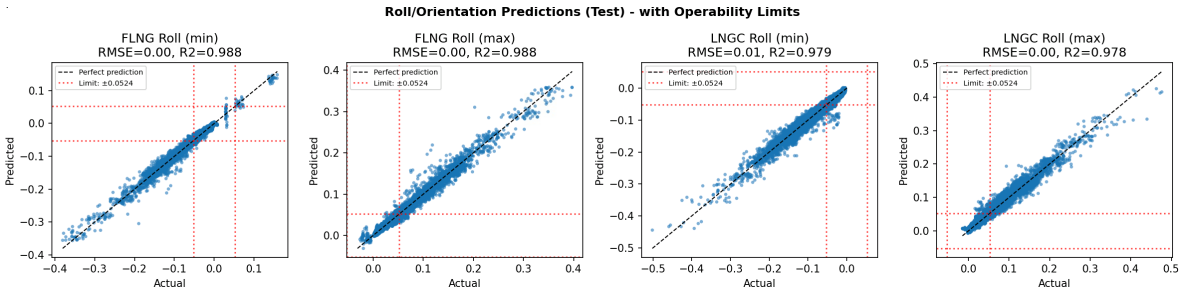


Figure 7.3: MLP predictions for minimum and maximum roll for both vessels, including operability limits.

Roll Predictions and Orientation Limits: Figure F.1 shows the predicted minimum and maximum roll angles for the FLNG and LNGC. The predictions align closely with the 1:1 diagonal and achieve very high correlation ($R^2 = 0.979$–0.989).

Because roll motion is a primary operability-limiting parameter, the strong agreement observed for both minimum and maximum roll responses indicates that the summary statistics-based model can reliably support orientation-limit assessment. Even near the operability threshold, the model does not show systematic underprediction of roll amplitudes, ensuring conservative and safe classification behaviour.

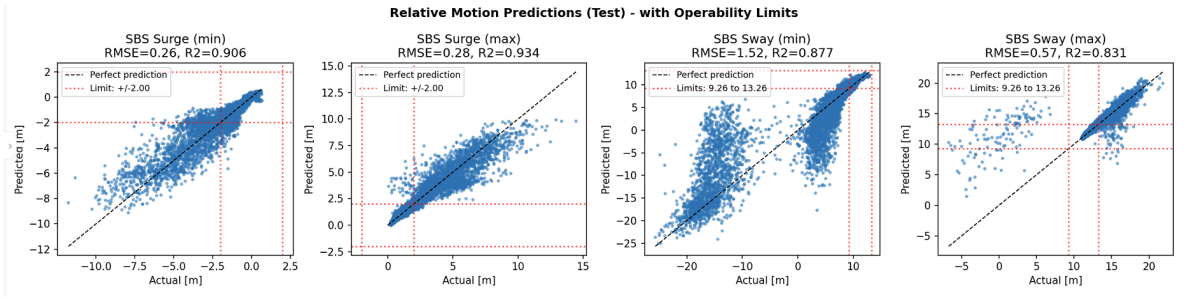


Figure 7.4: MLP predictions for minimum and maximum relative surge and sway, with operability limits indicated.

Relative Surge and Sway Predictions. Figure F.2 presents the relative motion predictions used to assess manifold clearance requirements. Relative surge predictions exhibit strong performance ($R^2 = 0.90$–0.93), with limited, unbiased scatter near the $\pm 2\,\mathrm{m}$ operability boundary. Both low-motion and high-motion regimes are captured accurately, indicating that the model can reliably identify conditions under which longitudinal clearance limits are exceeded.

The minimum relative sway response shows the largest scatter among the predicted quantities. This metric represents the instantaneous signed difference between the LNGC and FLNG manifold posi-

tions in sway, and its minimum value over the simulation is driven by the combined low-frequency (LF) and wave-frequency (WF) response of the coupled system. Inward relative motion activates several strongly non-linear mechanisms. Fender engagement introduces a highly non-linear force–compression relation with additional frictional effects [75]. At the same time, the SBS mooring system becomes increasingly active in out-of-phase horizontal modes, which exhibit short natural periods and low damping [77]. Small variations in LF excitation or gap-flow-induced hydrodynamics can therefore result in disproportionately large differences in instantaneous inward sway.

In addition, hydrodynamic interactions within the narrow gap modify both first- and second-order forces. Wave shielding, enhanced sway drift, and suppression of free-surface motion by the lid influence the timing and amplitude of the LF response envelope [75]. The presence of these interacting mechanisms leads to a non-unique mapping between environmental inputs and minimum relative sway, which explains the observed mid-range scatter.

In contrast, maximum relative sway corresponds to outward motion away from the fenders. In this direction, the response is governed primarily by low-frequency drift loading and mooring restoring forces, without the activation of fender contact or strong gap-induced effects. Although the mooring system remains non-linear, the outward response follows a consistent dependence on environmental parameters such as significant wave height, peak period, and wave direction. This leads to a clearer and more repeatable relationship between inputs and response, which explains the tighter alignment observed for the maximum sway predictions. Similarly, minimum surge does not involve fender engagement or pronounced gap-flow interaction, and hydrodynamic coupling effects are weaker than in sway [75]. As a result, the minimum surge response exhibits a more direct dependence on the environmental inputs, yielding improved prediction consistency compared to minimum sway.

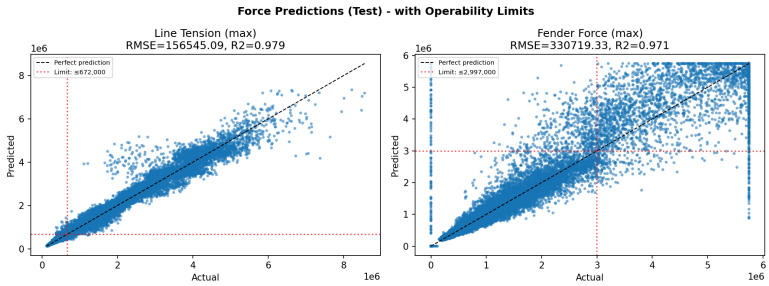


Figure 7.5: MLP predictions for maximum line tension and maximum fender force with operability limits included.

Line Tension and Fender-Force Predictions: Figure F.3 shows the predictions for the maximum line tension and the maximum fender force, which are two critical parameters for SSBL operability assessment. The MLP achieves excellent correlation, with R^2 values between 0.971 and 0.977, and reproduces the full force range with high fidelity. Importantly, the predicted values remain conservative near the MARIN-imposed limits, as the model does not systematically underpredict loads close to the threshold.

The fender-force predictions also show strong agreement, despite the increased scatter observed at high load levels. This behaviour reflects the nonlinear nature of contact mechanics. Overall, these results indicate that the summary-based model captures the dominant relationship between vessel motions and load responses sufficiently well to support operability decision-making.

Operability Classification

The overall operability classification performance of the summary-based MLP model is shown in Fig. 7.52.

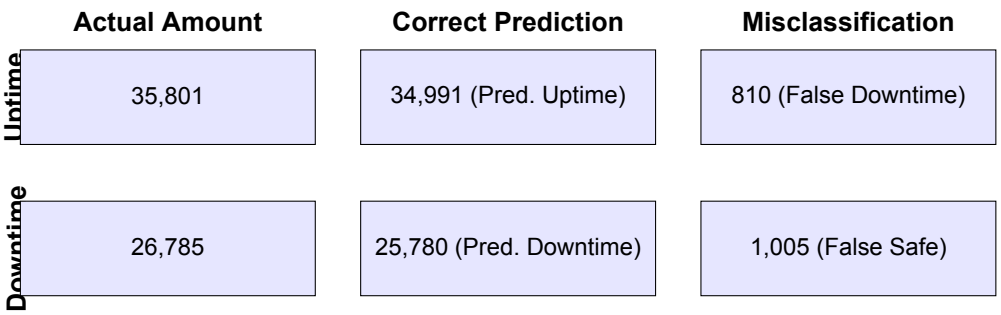


Figure 7.6: Confusion-matrix summary of uptime/downtime classifications for the test set (62,586 simulations).

Out of the 62,586 test simulations the model predicts in line wiht the assumend MARIN baseline with an overall accuracy of 97.1%. The confusion-style matrix shows the distribution of errors: 810 False Unsafe predictions, where the model conservatively labels an operable case as inoperable, and 1,005 False Safe predictions, where a violation is missed and a non-operable case is classified as operable. Because False Safe predictions represent the safety-critical failure mode, their rate remains low relative to the total number of downtime cases. The balance between predicted uptime and downtime closely matches the true distribution, indicating that the model does not show a systematic bias toward over- or under-prediction of operability. Taken together with the high accuracy across all parameter-level predictions, these results demonstrate that the summary-based MLP provides a reliable basis for first-pass SSBL operability screening.

SSBL: Key False Safe Patterns and Interpretation

Table 7.1: Representative environmental regimes producing False Safe (FS) cases across roll and manifold responses.

Phenomenon / Outputs	**Param 1**	**Param 2**	**Param 3**
Beam–sea LF amplification (relative surge, LNGC roll, FLNG roll)	heading $= 90°/270°$	$H_s^{sea} \approx 2.5\,\mathrm{m}$	$T_p^{sea} = 6$–$8\,\mathrm{s}$
Crossed sea–swell LF–WF coupling (relative surge/sway, roll contributions)	$T_p^{sea} = 6$–$10\,\mathrm{s}$	$T_p^{swell} = 12$–$20\,\mathrm{s}$	beam/quartering
Transverse beam sensitivity (relative sway, roll alignment)	heading $= 90°/270°$	$H_s^{sea} = 1.0$–$2.0\,\mathrm{m}$	$T_p^{sea} = 8$–$10\,\mathrm{s}$

Beam/near-beam amplification (roll and surge): False safe clustering in roll and relative sway occurs predominantly at beam headings with $H_s^{sea} \approx 2.5\,\mathrm{m}$ and $T_p^{sea} = 6$–$8\,\mathrm{s}$. In this regime, SBS multi-body interaction strengthens: gap-flow velocities increase, radiation damping is altered, and out-of-phase horizontal modes couple into roll and relative surge. Such beam–sea amplification and its sensitivity to proximity and radiation coupling are well established in multi-body hydrodynamics [75, 76]. This explains why LNGC roll peaks at $90°$ and FLNG roll peaks at $270°$—they are the mirrored active-beam sides of the SBS configuration.

Crossed sea–swell interaction: Conditions with $T_p^{sea} = 6$–10 s combined with swell periods $T_p^{swell} = 12$–20 s produce overlapping LF drift and WF roll/pitch responses. The swell lies near the natural roll period (15–16 s), enhancing WF roll, while the sea governs LF surge variability. Studies consistently report that these crossed conditions generate the strongest probability of downtime, line load exceedance and relative motion violations [51, 77]. The same mechanism explains simultaneous FS occurrences across Manifold X, Manifold Y, and both roll signals.

Sensitivity in beam and near-beam seas: Manifold Y FS events cluster at $90°/270°$ with moderate $H_s^{sea} = 1$–2 m and $T_p^{sea} = 8$–10 s. These headings activate out-of-phase sway/yaw modes with limited viscous damping and strong gap-flow forcing. Single-body models cannot replicate the resulting sway-drift peaks; explicit multi-body diffraction and damping treatment are essential [75, 51]. Roll FS counts align with this same directional sensitivity, reinforcing the shared physical origin.

These results show that the MLP summary-based model can provide accurate predictions of operability-critical metrics and forms a robust foundation for SSBL operability screening, especially when computational efficiency is prioritised. The absence of systematic underprediction near operability thresholds ensures conservativeness and supports safe decision-making.

7.1.4. Interpertation of Results

Both the Random Forest and MLP models show strong capability in predicting operability-relevant summary parameters from environmental inputs. The learning curves indicate that only 5–10% of the dataset is needed for high accuracy, with diminishing gains beyond 20–50%. RF is slightly more robust at very small data fractions, while the MLP reaches comparable performance once sufficient data is available.

For operability assessment, the summary-based MLP performs well: it predicts the key minimum and maximum responses required for safety checks and does not systematically underestimate values near operability thresholds.

However, summary-based models also have clear limitations. Because they operate on aggregated quantities—they cannot reconstruct how the maximum responses arise. As a result, they cannot capture the underlying time-domain evolution, intermediate peaks, wave-group effects or interactions that lead up to the reported maximum motions, loads, or clearances. These aspects require full time-series learning and remain outside the scope of summary-based learning.

In short, the summary models support accurate and conservative operability assessment, but do not provide insight into the dynamic processes that produce the limiting responses.

7.2. Time Series-Based Models

In the following the different result for the Time Series based models will be discussed. The predicted time series for every discussed result are for simulation following Table 7.2.

Table 7.2: Environmental condition set including significant wave height.

Wind vel.	Wind dir.	$H_{s,\mathrm{sea}}$	$T_{\mathrm{p,sea}}$	Sea dir.	Curr. vel.	Curr. dir.
1 m/s	90°	2 m	8 s	180°	0.12 m/s	90°

7.2.1. Segment-Based Sequence to Sequence Learning

As discussed in chapter 5, the segment-based sequence-to-sequence models are trained on the first 70% of each 3-hour simulation, where both the incoming wave and the resulting motions are provided. In the remaining 30%, only the incoming wave is given and the model must predict the vessel motions without further access to the true response history.

Low-frequency and wave-frequency components of the responses are separated following the procedure described in subsection 6.2.2. The following sections present and analyze the prediction results for both components. The analysis uses the autocorrelation function to evaluate how well the model captures the temporal structure of the motions and to assess the suitability of this approach for SBS offloading operability prediction.

First, the low-frequency predictions are examined, followed by the wave-frequency predictions. A single sea state is used for illustrating the method, as the fundamental behaviour of the forecasting approach should not depend on the specific environmental condition: the model predicts vessel motions based on past motion history and the incoming wave, and its essential capabilities remain the same across sea states.

7.2.1.1. Low-Frequency Results

As low-frequency surge and sway responses are particularly relevant for operability in SBS offloading, these are discussed. The following Figure 7.7 and Figure 7.12 present the recursive predictions for each vessel, allowing a direct evaluation of how accurately the model reproduces the dominant low-frequency behaviour over the full forecast horizon.

As the modelled sea state corresponds to head seas, the sway response is much smaller than the surge response. In head-sea conditions, surge is directly excited by the incident waves, whereas sway is only generated indirectly through diffraction, radiation effects, and mechanical coupling between the vessels.

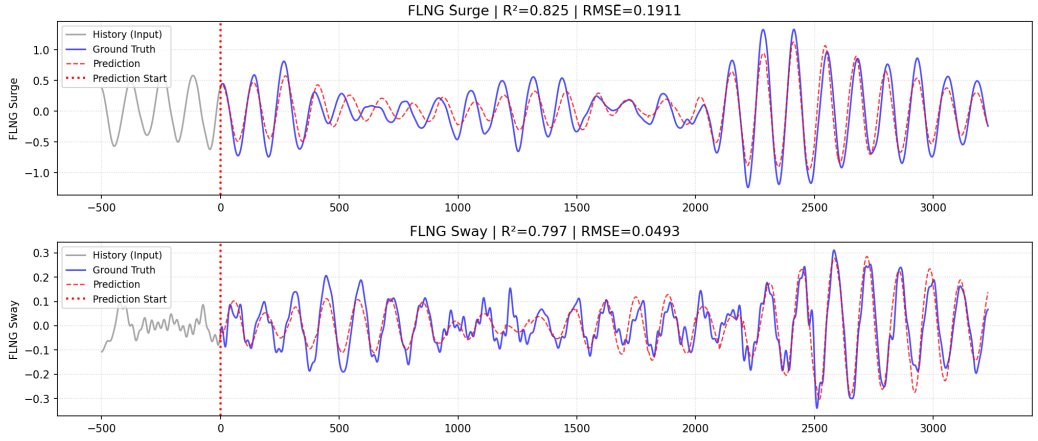


Figure 7.7: Recursive low-frequency predictions for the FLNG.

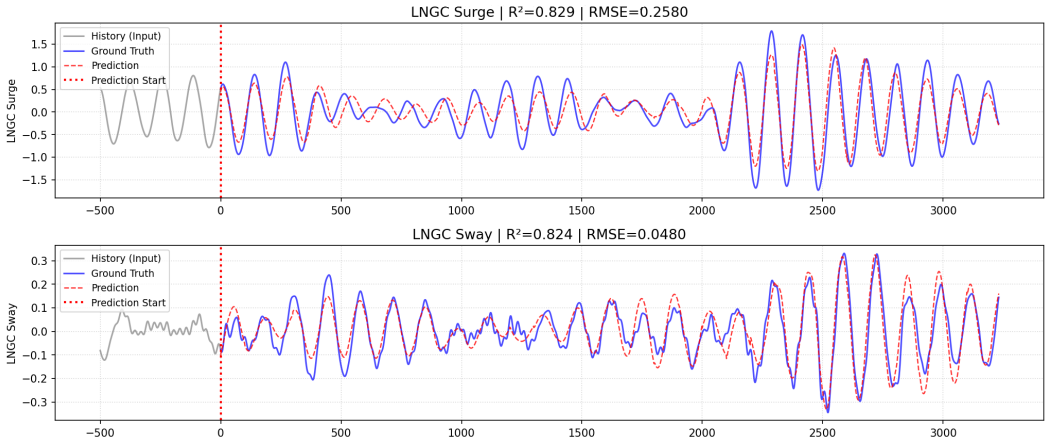
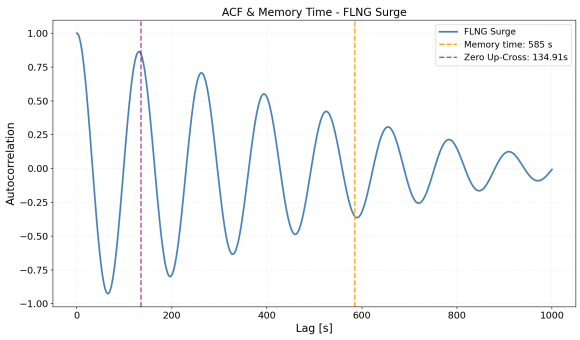


Figure 7.8: Recursive low-frequency surge and sway predictions for the LNGC.

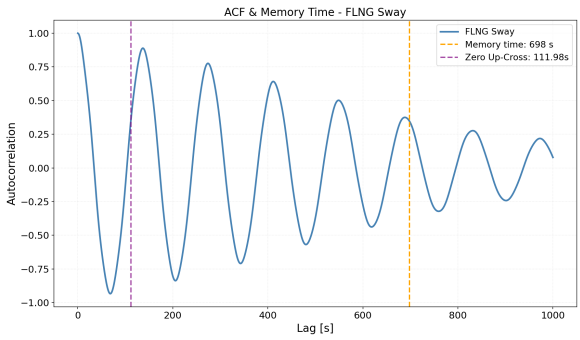
The recursive rollout plots for the low-frequency components reveal strong periodic behavior in surge and sway motions, which the model captures effectively over the full prediction horizon. Surge, sway exhibit clear oscillatory patterns which follow the true phase and amplitude. This indicates that the dominant low-frequency dynamics are highly predictable based on past motions and maintain coherence over time.

Low Frequency Autocorrelation Function Analysis: To examine the influence of low-frequency dynamics on motion predictability, two representative ACFs are shown for the surge and sway degrees of freedom. Both signals are dominated by low-frequency drift forces. The memory time τ_m is determined using the $1/e$ envelope criterion, which identifies the lag at which the ACF envelope decreases to approximately 0.37 of its initial value Equation 5.3.

Surge The surge ACF shows a slowly decaying low-frequency oscillation with a memory time of roughly 585 s. The autocorrelation remains coherent over approximately four oscillation cycles, indicating that the surge signal retains long-range temporal dependence. The oscillation period remains stable across the full lag range, confirming that the underlying dynamics are consistent. The dominant low-frequency period obtained from the zero-crossing method is approximately 135 s, which corresponds closely to the natural surge period and fits well within the observed memory horizon.



(a) FLNG Surge: $\tau_m \approx 585$ s



(b) LNGC Sway: $\tau_m \approx 698$ s

Figure 7.9: Low-frequency ACFs and memory times computed using the $1/e$ envelope method.

Surge motion is primarily driven by second-order wave drift forces acting along the vessel's longitudinal axis. Because hydrostatic restoring in this DOF is very weak, the response is governed mainly by the mooring system, resulting in slow variations and gradual energy dissipation. This leads to an extended memory and a predictable low-frequency signal, where future motion is strongly influenced by several preceding oscillatory cycles.

Sway The sway ACF exhibits a similar low-frequency character, but with an even longer memory time of about 698 s. As with surge, the ACF maintains coherent oscillations over six cycles. The oscillation period is stable across the entire lag window, reflecting a consistent lateral low-frequency response. The dominant period from the zero-crossing method is approximately 112 s, which aligns with the expected natural sway period and fits proportionally within the longer memory time.

Sway motion is likewise dominated by second-order wave drift forces and experiences very low hydrostatic restoring. The restoring contribution therefore comes mainly from the mooring system, leading to slow, long-period lateral oscillations with weak damping. This correlation structure makes the sway signal predictable and suitable for long-horizon forecasting.

Taken together, the ACF results for surge and sway show long memory: both motions retain consistent low-frequency structure over approximately six full oscillation cycles. This extended memory window means that future states remain strongly correlated with past states, which benefits long□horizon forecasting.

7.2.1.2. Wave-Frequency Results

In this section, the wave-frequency responses for heave and roll are examined. These two motions are are dominated by WF behaviour, with roll being particularly important because it governs operability limits in SBS offloading. The following results show how well the model captures these fast, wave-driven oscillations for both the FLNG and the LNGC.

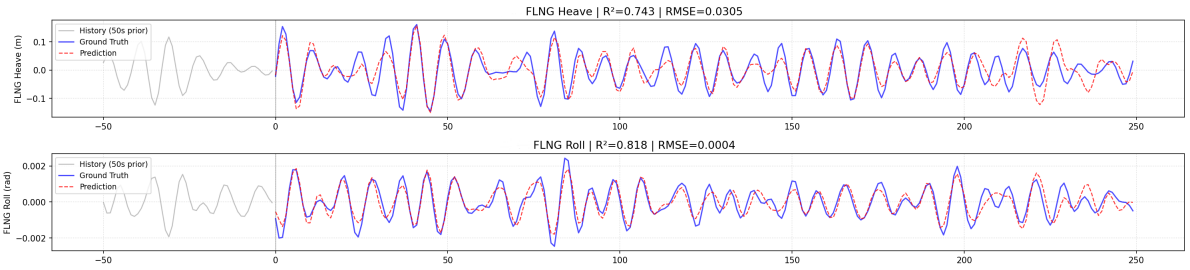


Figure 7.10: Recursive wave-frequency predictions for the FLNG (early window, first 250 s).

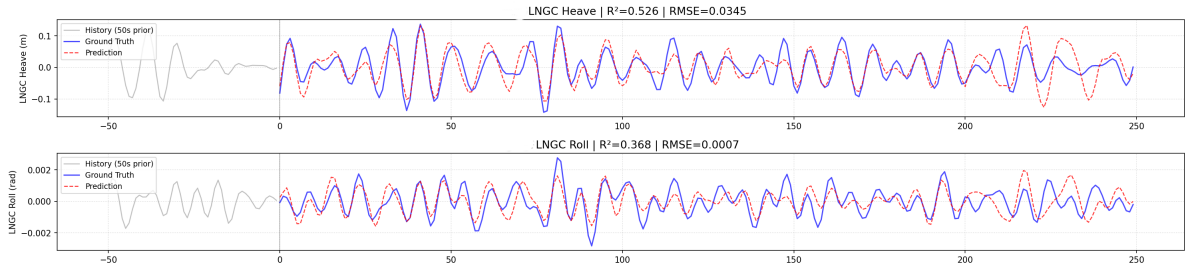


Figure 7.11: Recursive wave-frequency predictions for the LNGC (early window, first 250 s).

The recursive rollout plots for the wave-frequency components show that the model captures the oscillatory behaviour well at the start, but accuracy can decreases as the prediction horizon grows (**??**, **??**). In the early window (first 250 s), both heave and roll align well with the true signal, with heave maintaining stable. For the LNGC, roll begins to accumulate errors earlier and deviations become noticeable.

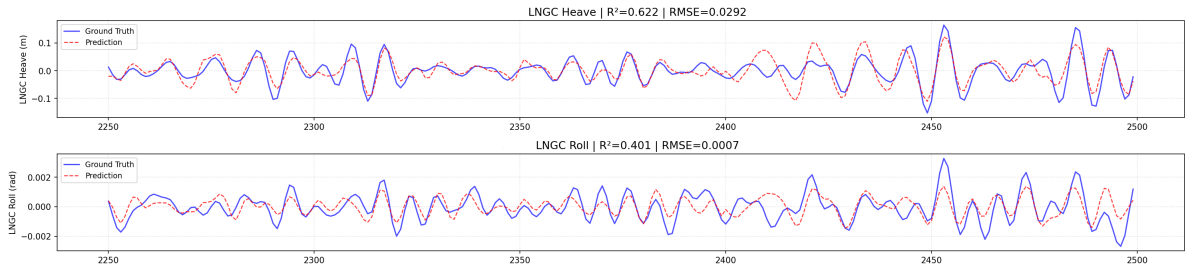


Figure 7.12: Recursive wave-frequency heave and roll predictions for the LNGC (late window, 2250–2500 s).

Because the recursive method feeds each predicted value back into the model, small inaccuracies at every step compound over time, eventually produce noticeable shifts in both timing and amplitude. However, this drift develops slowly. The wave frequency components respond almost instantaneously to the incoming waves, which the model continues to "see" through the environmental inputs. As a result, the predicted signal can remain largely in phase with the true motion even multiple cycles ahead.

For the FLNG, the wave-frequency predictions stay relatively stable throughout the sequence, with only occasional dips in correlation where accumulated error becomes visible. Overall, the strong direct link between incident waves and wave-frequency vessel response helps the model maintain coherence further into the recursive rollout than would be expected from a purely autoregressive forecast.

As shown in Figure 7.14 and Figure 7.15, both vessels show intermittent drops in R^2 for the wave-frequency roll, heave, and pitch predictions. These drops occur simultaneously across both vessels,

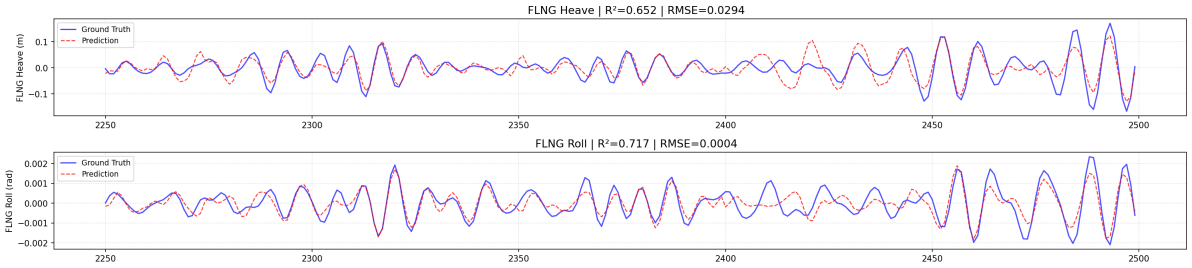


Figure 7.13: Recursive wave-frequency heave and roll predictions for the FLNG (late window, 2250–2500 s).

which is consistent with a shared cause in the incoming wave signal. The waves are generated using a JONSWAP spectrum [109], which naturally produces wave groups with local variations in amplitude and phase. It is hypothesised that during transitions between wave groups, the instantaneous wave excitation changes in a way that the model has not sufficiently learned to track, causing a temporary loss of phase alignment with the reference signal for both vessels at the same time. Such phase errors are particularly impactful for the R^2 metric, which is sensitive to timing mismatches even when amplitude is well reproduced. The LNGC roll (green) shows more severe and prolonged drops than the FLNG, which may reflect its higher sensitivity to gap-flow and multibody interaction effects that are not fully captured from the single-point wave elevation input alone.

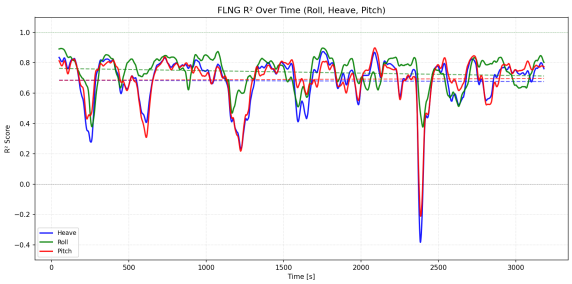


Figure 7.14: R^2 over time for recursive wave-frequency FLNG predictions.

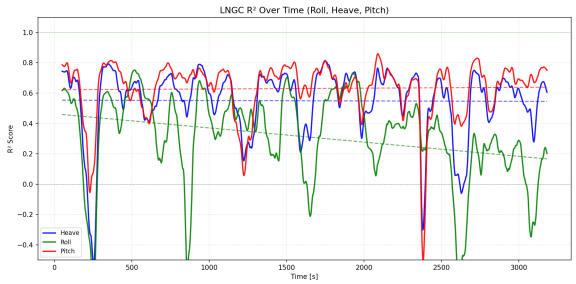


Figure 7.15: R^2 over time for recursive wave-frequency LNGC predictions.

In short, wave-frequency motions are predicted well for short periods, but over longer horizons the recursive approach leads to increasing errors. Motions with more regular wave patterns, such as heave, remain accurate for longer, while roll motion prediction drift sooner. This shows that predicting wave-frequency dynamics over long periods is difficult especially if you want to significantly reduce the simulation time.

Wave Frequency Auto Correlation Analysis: Below, wave-frequency ACFs of FLNG and LNGC roll are shown to evaluate their influence on prediction capability. At wave frequencies, the model can start to drift over long prediction horizons because it is very sensitive to timing and phase. Motions with cleaner and more regular auto correlation oscillations stay accurate for longer, while less regular ones lose accuracy sooner.

FLNG roll: Figure 7.16 shows a clear wave-frequency oscillation with a dominant period of ≈ 8 s. The correlation decays rapidly to the $1/e$ level after ≈ 9 s, so the wave-frequency memory contains only $\sim$1 cycle.

This short-memory is consistent with roll being not directly being excited by the waves in head seas: first-order head waves act mainly in heave and pitch, and any roll content arises through coupling (e.g.

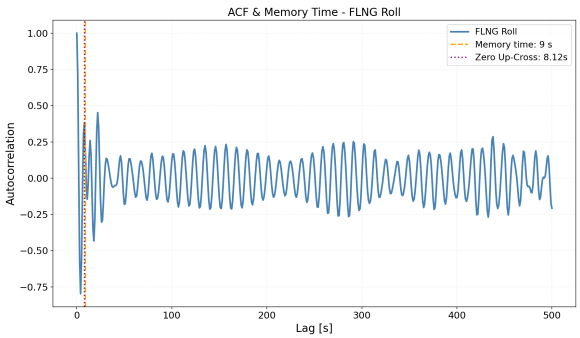


Figure 7.16: ACF & memory time for FLNG roll (wave-frequency).

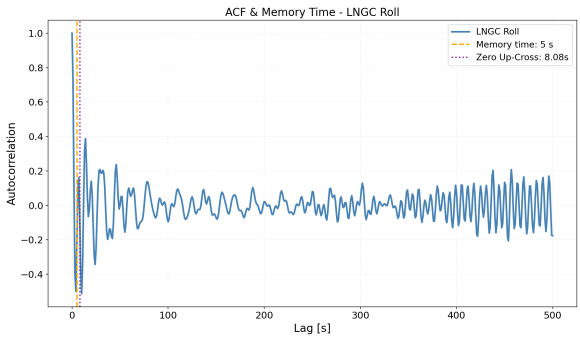


Figure 7.17: ACF & memory time for LNGC roll (wave-frequency).

cross-coupled moments and mechanical coupling), while strong hydrostatic restoring quickly damps correlation in time. As can be seen in section C.1 the memory time of heave and pitch memory time already extends to $\sim$2 cycles.

LNGC roll: In Fig. 7.17, the LNGC exhibits a similar dominant period ($\approx$ 8 s) but an even shorter memory time (≈ 5 s). Oscillations are less regular due to SBS effects: fenders, mooring-line geometry and gap-flow interactions. As with the FLNG, roll is governed by in-direct first-order wave excitation and strong hydrostatic restoring, but the SBS coupling accelerates loss of correlation, further limiting the useful memory time.

Unlike the LNGC, the FLNG acts as the heavier, moored reference body, which gives it a more regular and stable roll ACF. The LNGC roll ACF, by contrast, shows faster decay and more irregular oscillations, reflecting its greater sensitivity to multibody interaction forces in the SBS configuration.

The ACF behaviour helps explain the differences seen in the prediction performance of the roll signal for the two vessels. Because the FLNG roll ACF is more regular and decays more slowly, its wave-frequency roll response remains coherent over a longer time, which aligns with the more stable recursive predictions observed in the rollout. In contrast, the LNGC roll ACF shows faster decay and more irregular oscillations, indicating shorter memory and less phase consistency. This reduced temporal coherence is reflected in the prediction plots, where LNGC roll begins to accumulate timing and amplitude errors earlier than the FLNG.

7.2.1.3. Training Data Requirements for Wave- and Low-Frequency Models

A learning-curve analysis was performed to assess how much of the available simulation dataset is required to achieve stable and accurate predictions for the wave-frequency and low-frequency response components. The two models exhibit distinct data requirements, reflecting the different physical mechanisms governing wave-frequency and low-frequency vessel motions.

For the wave-frequency model, the FLNG roll motion is taken as the reference parameter for evaluating prediction quality. For each training fraction, the model is trained for a fixed number of 100 epochs to provide a consistent and efficient comparison of performance. As shown in Figure 7.18, prediction accuracy improves rapidly with increasing training set size and stabilises at approximately 40% of the available simulations. Beyond this point, reductions in RMSE and improvements in R^2 become marginal, indicating diminishing returns. This behaviour is consistent with the short intrinsic memory of wave-frequency motions, which are dominated by first-order wave excitation and are primarily governed by instantaneous phase alignment with the incoming wave signal.

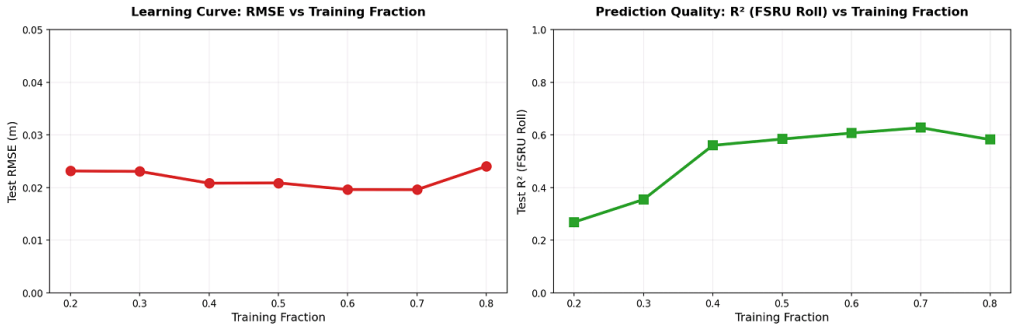


Figure 7.18: Learning-curve analysis for the wave-frequency time-series model using FLNG roll as reference parameter. All models are trained for 100 epochs per training fraction.

In contrast, the low-frequency segment-based LSTM model requires a larger fraction of the dataset to reach stable performance. In this case, FLNG surge motion is used as the reference parameter, as it is representative of second-order drift-dominated behaviour. The model is again trained for 100 epochs for each training fraction. As shown in Figure 7.19, a strong reduction in RMSE is observed at relatively small training fractions, while the R^2 values continue to increase up to approximately 60% of the dataset. This reflects the longer intrinsic memory and system-level nature of low-frequency motions, which are governed by second-order drift forces and multibody coupling effects. Adequate representation of these response patterns therefore requires broader coverage of the environmental parameter space.

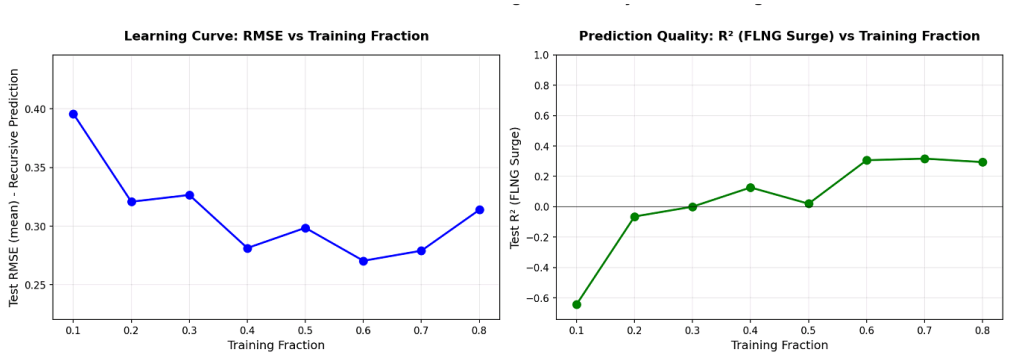


Figure 7.19: Learning-curve analysis for the low-frequency segment-based LSTM model using FLNG surge as reference parameter. Models are trained for 100 epochs at each training fraction.

Overall, the analysis shows that approximately 40% of the available simulations is sufficient for accurate wave-frequency prediction, whereas reliable low-frequency prediction requires closer to 60% of the dataset. These findings support the use of reduced training sets for wave-frequency modelling, while highlighting the higher data demands associated with learning low-frequency, drift-dominated dynamics.

The rapid convergence of the wave-frequency model can be explained by the amount of learnable structure present within a single time series. Figure 7.20 shows a representative wave-elevation record together with its envelope and detected wave groups over a three-hour simulation. The wave-group envelope is obtained from the analytic signal using the Hilbert transform, which provides a physically meaningful description of amplitude modulation in narrow-band random wave processes [**Falnes2020**]. Hundreds of wave groups and local extrema are present within a single simulation, providing many

repeated excitation patterns from which the wave-frequency model can learn.

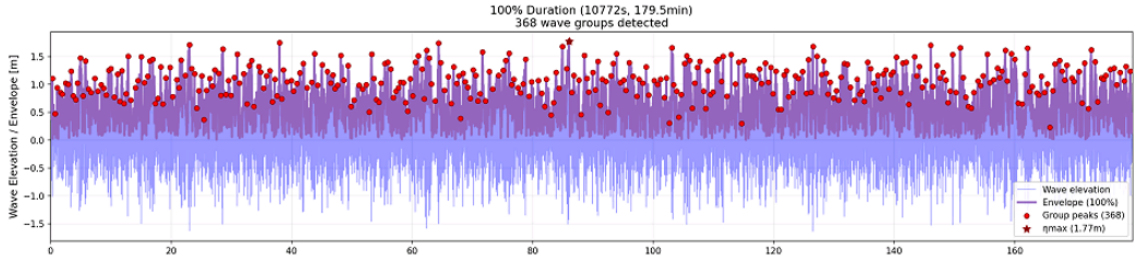


Figure 7.20: Wave-elevation time series with envelope and detected wave groups over a three-hour simulation. .

For irregular seas, the significant wave height is defined as

$$H_s = 4\sigma_\zeta,$$

where σ_ζ is the standard deviation of the wave elevation. The expected maximum crest height over a finite record can be approximated following Longuet–Higgins [110] as

$$\zeta_{\max} \approx \sigma_\zeta \sqrt{2 \ln N},$$

where N is the number of statistically independent wave cycles. For a three-hour simulation with wave periods of 8 s, yielding an expected maximum crest of approximately $1.8\ \mathrm{m}$.

In the simulations considered here, the largest wave group is already observed within the first 50% of the time series, and within the first 25% the maximum wave elevation reaches approximately $\zeta \approx 1.75\ \mathrm{m}$, close to the expected three-hour maximum. Because all simulations use the same random wave seed, this behaviour is consistent across cases for a given $T_{\mathrm{p,sea}}$. Consequently, the most energetic wave-frequency excitation events are encountered early and repeatedly, supporting why the wave-frequency model converges rapidly and requires only a limited fraction of the full dataset.

7.2.1.4. Interpretation of the results

The analysis highlights clear differences between the predictability of low-frequency (LF) and wave-frequency (WF) motions. WF components exhibit short memory, retaining only about 1–2 oscillation cycles before the autocorrelation decays. This limited temporal dependence makes long-horizon prediction more challenging and leads to faster growth of timing and amplitude errors. Nevertheless, because WF motions follow the incident wave almost instantaneously, the predictions remain reasonable over short horizons.

In contrast, LF motions evolve slowly and show a much longer memory of approximately 4 cycles. This extended coherence allows the model to maintain stable multi-step predictions over longer periods. For reference, a 3000 s horizon contains roughly 30 LF cycles but around 300 WF cycles. The LF predictor is therefore required to propagate the signal over far fewer cycles, and because the underlying LF dynamics change gradually, long-term forecasting remains comparatively robust. If the LF prediction were propagated over hundreds of cycles, similar to the WF case, the model would also begin to diverge due to compounding errors, especially since LF motions are not directly tied to the instantaneous incoming wave. Moreover, the WF components have many more cycles available within the first 70% training window, giving the model substantially more examples to learn the relationship between wave input and motion response. In contrast, the LF component contains far fewer cycles in the same training interval, making it more dependent on the intrinsic memory of the physical signal itself rather than on large amounts of training data.

Component	Period	Memory (cycles)	Training cycles (70%)	Test cycles (30%)
Wave-frequency (WF)	$\approx$ 10 s	$\approx$ 1–2	$\approx$ 756	$\approx$ 324
Low-frequency (LF)	$\approx$ 400 s	$\approx$ 4	$\approx$ 19	$\approx$ 8
Implications for prediction: WF: short memory $\Rightarrow$ faster error growth over long horizons; still reasonable at short horizons due to direct wave forcing. LF: longer memory & slower evolution $\Rightarrow$ more stable multi-step forecasts; fewer cycles to propagate over long horizons.				

Table 7.3: Summary of dominant period, memory, and cycle counts (for a 3 h record: 70% training = 7 560 s; 30% test = 3 240 s).

Overall, the segment-based forecasting concept proves accurate for predicting new motion segments from past motion and the incoming wave. However, in the context of operability assessment across a wide range of environmental conditions, this approach offers limited practical benefit on its own. Each sea state still requires a partly numerical simulation to generate the motion history used for the forecast, meaning that the total computational effort is not substantially reduced. For this reason, while the method performs well on a given case, it does not yet provide the efficiency gains needed for broad operability analyses.

7.2.2. Full Time Series-Based Learning

In the full time-series-based learning approach, an unseen sea state is predicted based on training data from sea states with similar conditions. The model is trained on the full 3 hour simulation time series of the input wave elevation time series and the output time series of vessel responses. in the subsequent subsections, the same time series is predicted with training on different variations of training sets. This way it is possible to test reliably, whether the model can generalize to new H_s, T_p.

This allows us to evaluate for which parameters these models can generalize (predict when unseen) and for which parameters a learnable relationship exists for predicting unseen 3 hour simulation responses. Again, a separate low frequency and wave frequency model are applied to account for the first order wave force and second order wave drift force separately

In this section, only the motion signals that are directly relevant for operability assessment are shown. Line-tension signals are predicted as well, but are intentionally not displayed here. In both the numerical MARIN model and the time series models used in this thesis, the mooring tensions are computed as a direct consequence of the vessel motions. It is therefore essential to first establish that the motion predictions are accurate, before evaluating the corresponding line-tension behaviour.

Generalizing to new $H_{\mathrm{s,sea}}$

To assess whether the model can generalize to new values of $H_{\mathrm{s,sea}}$, the transformer is trained on full 3-hour simulations under identical environmental conditions, with only the significant wave height varied between cases. Wind and current are purposely chosen with minimal magnitude and aligned to avoid interaction with the incoming waves, ensuring that the response is dominated solely by head-wave excitation rather than additional environmental effects. The conditions for the training and test data are listed in Table 7.9.

	Train						**Test**
$H_{\mathrm{s,sea}}$ (m)	0.75	1.00	1.50	2.50	3.00	3.50	**2.00**

Wind vel.	Wind dir.	$T_{\mathrm{p,sea}}$	Sea dir.	Curr. vel.	Curr. dir.
1 m/s	90°	8 s	180°	0.12 m/s	90°

Table 7.4: Significant wave heights and constant environmental parameters across training and predicted cases.

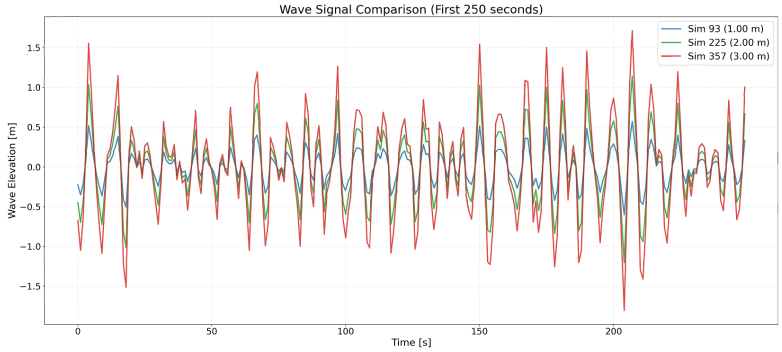


Figure 7.21: Comparions of wave signal for different $H_{\mathrm{s,sea}}$.

Low-Frequency Predictions

It can be seen in Figure 7.22 that surge and sway—the important low-frequency motions—are predicted well for the FLNG: the model follows the overall evolution for both DOFs and the maximum responses in the time window are captured by the predicted signals.

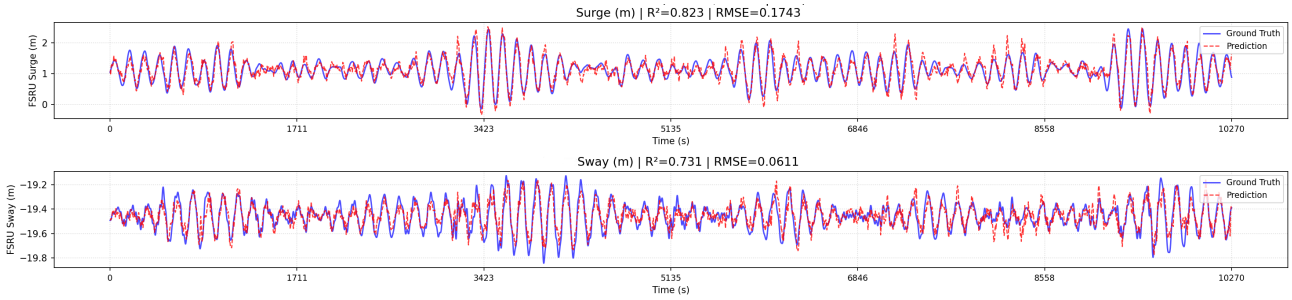


Figure 7.22: FLNG surge and sway predictions vs. reference for unseen $H_{\mathrm{s,sea}}$.

As can be seen for both the FLNG and LNGC, surge is predicted better with $R^2 \approx 0.8$, while sway also shows strong performance with $R^2 \approx 0.7$. No material difference is observed between FLNG and LNGC predictive skill for these low-frequency DOFs.

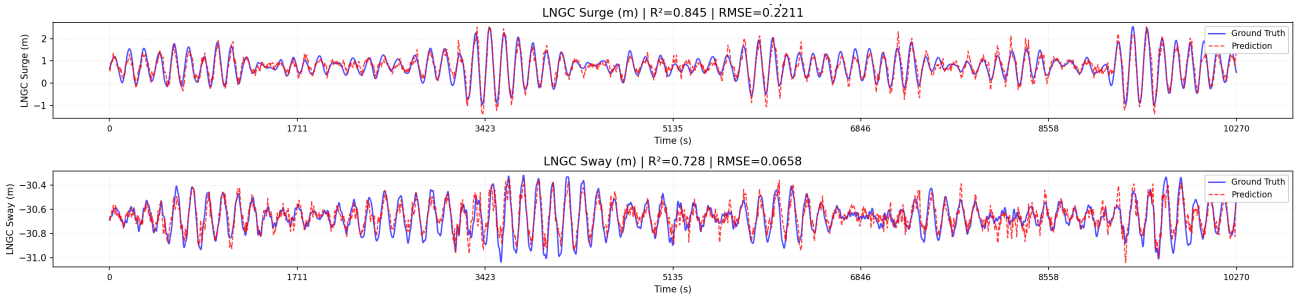


Figure 7.23: LNGC surge and sway predictions vs. reference for unseen $H_{\mathrm{s,sea}}$.

For wave signals with different $H_{\mathrm{s,sea}}$, the incoming wave simply scales in amplitude, as shown in Figure 7.21. This implies that, in pure head-sea conditions and without side-by-side interaction effects, motions that are directly forced by the incoming waves—surge, heave, and pitch—would also scale proportionally. In Figure 7.25, the surge response indeed follows this trend, but small deviations appear.These differences arise from hydromechanical and mechanical coupling effects in the SBS configuration, meaning that an accurate model must capture these nonlinear interactions rather than relying solely on scaling of the response it learned in the exact same seastate, but wiht a different significant wave height.

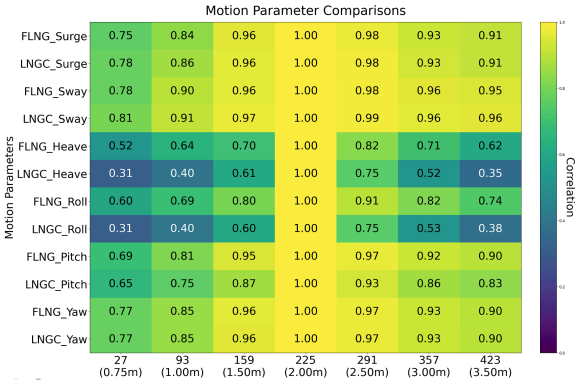


Figure 7.24: LF motion cross-correlation for varying $H_{\mathrm{s,sea}}$.

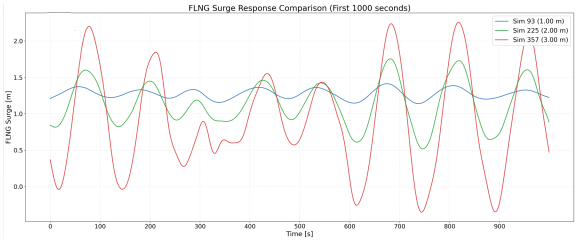


Figure 7.25: Surge motion (LF) comparison across $H_{\mathrm{s,sea}}$ (FLNG).

The low-frequency part responds well, to new wave height $H_{\mathrm{s,sea}}$ which whem looking at Figure 7.53, Figure 7.24 & Figure 7.25 is to be expected as the all the waves are modelled under the same seed and the model basically just has the scale the low frequency response seen during training.

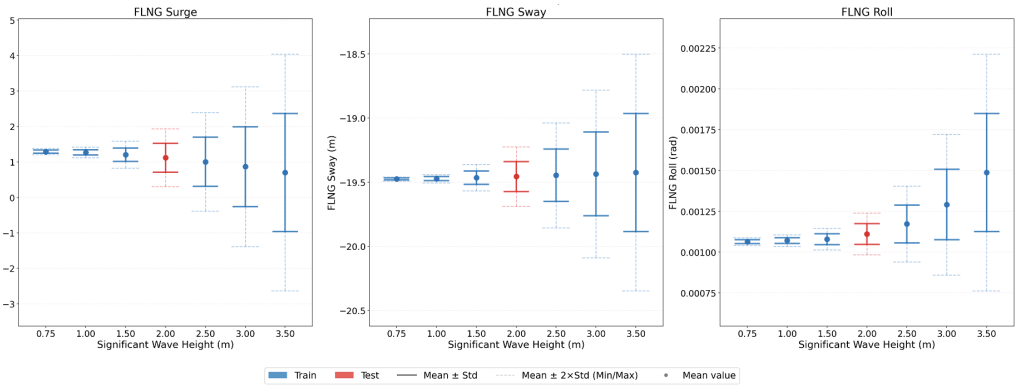


Figure 7.26: Comparison of wave signal for different $H_{\mathrm{s,sea}}$, mean $\pm$ std bars (with min–max ranges).

Figure 7.26 shows how the low-frequency motion statistics vary with significant wave height $H_{\mathrm{s,sea}}$. Because all wave signals were generated with the same random seed, their temporal shape is identical and only the amplitude increases with $H_{\mathrm{s,sea}}$. This is reflected in the motion statistics: surge, sway, and roll all exhibit larger spread (standard deviation and min/max range) at higher $H_{\mathrm{s,sea}}$, while the mean values remain nearly unchanged. The increasing error bars therefore represent a pure amplitude scaling of the underlying wave forcing, consistent with the strong cross-correlation observed in Figure 7.24. This explains why the model generalises well to new $H_{\mathrm{s,sea}}$: the LF response shape does not change, only its magnitude.

Wave-Frequency Predictions

For the heave motion, the model predicts the wave-frequency cycles well with $R \approx 0.9$ and low RMSE ≈ 0.02. In roll, agreement remains reasonable with $R \approx 0.6$; small phase shifts and some peak-height differences are visible, which is expected because roll under head seas is not directly forced and arises from a combination of radiated-wave coupling and mechanical interaction in the system. Overall, the model captures shape, timing, and amplitude adequately, indicating that wave-frequency predictions generalize to new $H_{\mathrm{s,sea}}$.

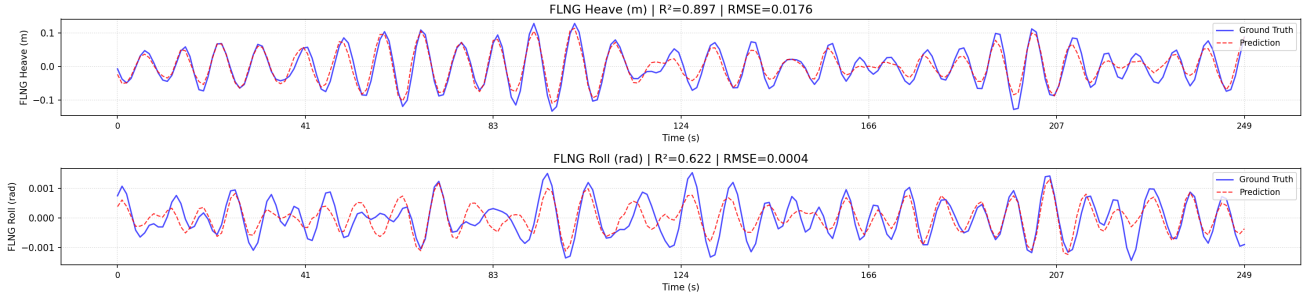


Figure 7.27: FLNG first 250 s: heave and roll predictions vs. reference (unseen $H_{\mathrm{s,sea}}$).

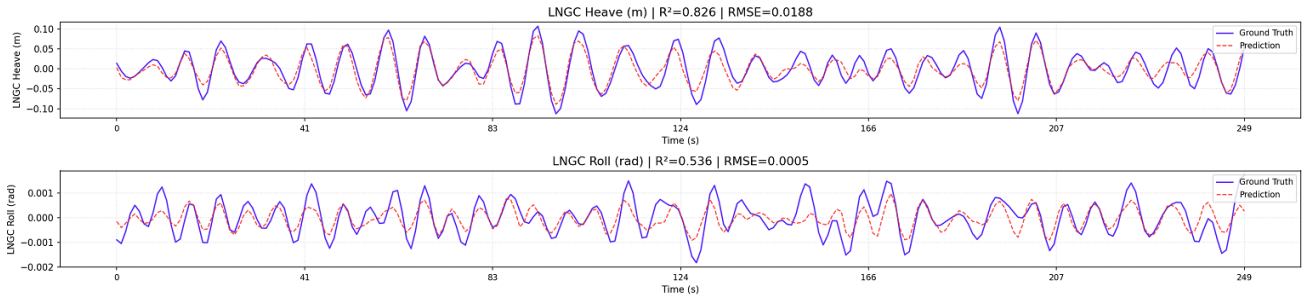


Figure 7.28: LNGC first 250 s: heave and roll predictions vs. reference (unseen $H_{\mathrm{s,sea}}$).

To relate the training motion responses to the test-case responses across LF/WF dynamics, the Pierson correlation is used (see chapter 5). Figures 7.24 - 7.29 show that for the LF response, surge and sway correlations remain high, which explains why these DOFs are predicted well, when predicting for a new significant wave height $H_{\mathrm{s,sea}}$.

At wave frequency, heave and pitch correlations are high, consistent with head seas where these motions are directly forced by the waves; changes in $H_{\mathrm{s,sea}}$ primarily scale the response amplitude rather than its timing. Especially because across these wave input time series the seed is not changed, so a simulation with a bigger wave height jsust scales the incoming wave. In contrast, roll, especially for the LNGC, exhibits a faster drop in correlation because roll in head seas is not directly forced and is more influenced by radiated-wave coupling and mechanical interaction within the SBS system. This behaviour is reflected in the predictions: heave is accurately predicted, while roll shows more frequent phase and peak-height inconsistencies.

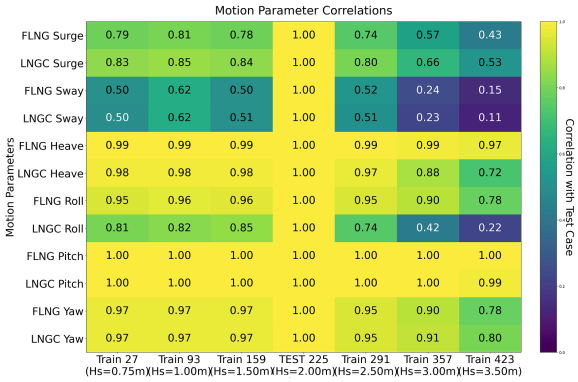


Figure 7.29: WF motion cross-correlation for varying $H_{\mathrm{s,sea}}$.

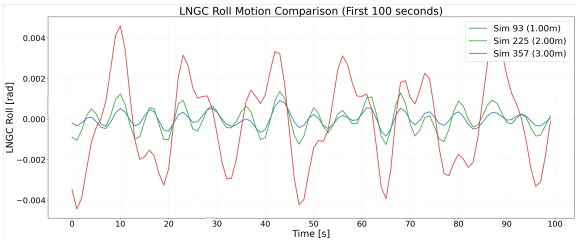


Figure 7.30: Wave-frequency (WF) LNGC roll motion across $H_{\mathrm{s,sea}}$ (FLNG).

To give some reference on the correlation matrices in Figure 7.24 and Figure 7.30, are the incoming

wave for three different $H_{\mathrm{s,sea}}$ and the surge and roll response to see how they differ with varying the incoming waveheight.

Interpertation of Results

In summary, the model generalizes well to unseen values of $H_{\mathrm{s,sea}}$ for both the LF and WFcomponents. For LF surge and sway, the response structure remains almost identical across wave heights because all simulations use the same wave seed; changing $H_{\mathrm{s,sea}}$ mainly scales the incoming wave. This is reflected in the high LF correlations and in the similar surge envelopes across $H_{\mathrm{s,sea}}$, which explains why the model can accurately reproduce these motions.

For the WF response, the model also performs well. Heave remains highly correlated across $H_{\mathrm{s,sea}}$ and is predicted with high accuracy, while roll shows moderate deviations, consistent with its indirect excitation mechanism in head seas. Because $H_{\mathrm{s,sea}}$ does not alter the spectral content of the wave but only its magnitude, the WF relationship between the incoming wave and the resulting motion remains largely unchanged. This allows the model to transfer its learned behaviour effectively.

Overall, $H_{\mathrm{s,sea}}$ is a parameter for which a clear and learnable input–output relationship exists: both LF and WF dynamics preserve their temporal structure when the wave height changes. As a result, the model can reliably predict the vessel motions for new $H_{\mathrm{s,sea}}$ values under otherwise identical conditions.

Generalizing to new $T_{\mathrm{p,sea}}$

To assess whether the model can generalize to new values of $T_{\mathrm{p,sea}}$, the same architecture is trained on full 3-hour simulations in which only the typical period is varied while all other environmental conditions are kept identical. Wind and current conditions are chosen to match those used in subsection 7.2.2. The specific training and test cases are listed in Table 7.9.

	Train				**Test**
$T_{\mathrm{p,sea}}$ **(s)**	7	9	10	11	**8**

Wind vel.	Wind dir.	$H_{\mathrm{s,sea}}$	Sea dir.	Curr. vel.	Curr. dir.
1 m/s	90°	1.5	180°	0.12 m/s	90°

Table 7.5: Typical periods and constant environmental parameters across training and predicted cases.

To better understand the characteristics of the incoming waves, the corresponding dominant period ranges for each $T_{\mathrm{p,sea}}$ are provided in Table 7.6, and the associated wave signals are shown in Figure 7.32. These ranges highlight how varying $T_{\mathrm{p,sea}}$ reshapes the wave spectrum and shifts the distribution of energy across frequencies.

$T_{\mathrm{p,sea}}$ [s]	6.0	7.0	8.0	9.0	10.0	11.0
Period range [s]	2.33–8.42	2.18–9.83	2.11–11.23	2.38–12.64	2.64–14.04	2.90–15.45

Table 7.6: Dominant wave–period ranges for each peak period $T_{\mathrm{p,sea}}$.

this means that, for generalization to new $T_{\mathrm{p,sea}}$ values, the model cannot rely on simple amplitude scaling of the incoming wave. Instead, it must learn how differences in wave frequency content translate

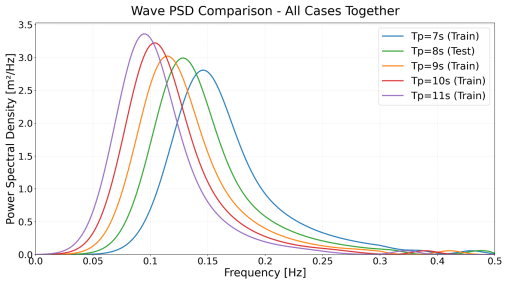


Figure 7.31: Comparison of PSD of different $T_{\mathrm{p,sea}}$

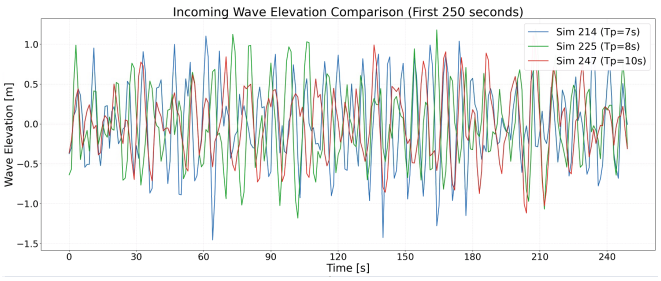


Figure 7.32: Comparison of wave signals for different $T_{\mathrm{p,sea}}$.

into differences in the resulting motion response. In contrast to the $H_{\mathrm{s,sea}}$ case—where waves generated with the same seed mainly differ in amplitude—the typical period fundamentally changes the spectral structure of the input (Figure 7.31), requiring the model to capture a more complex frequency-dependent relationship between the incoming wave and the vessel motions. This aligns with Chapter 3, where second-order drift forces are described as being rooted in quadratic interactions between spectral components; thus, changing $T_{\mathrm{p,sea}}$ restructures the drift-force spectrum rather than simply attenuating or amplifying it.

Low-Frequency Predictions

As can be seen in Figure 7.33, training on different $T_{\mathrm{p,sea}}$ values does not enable the model to produce accurate low-frequency FLNG surge and sway predictions. The model cannot reproduce either the phase or the amplitude evolution of the LF response.

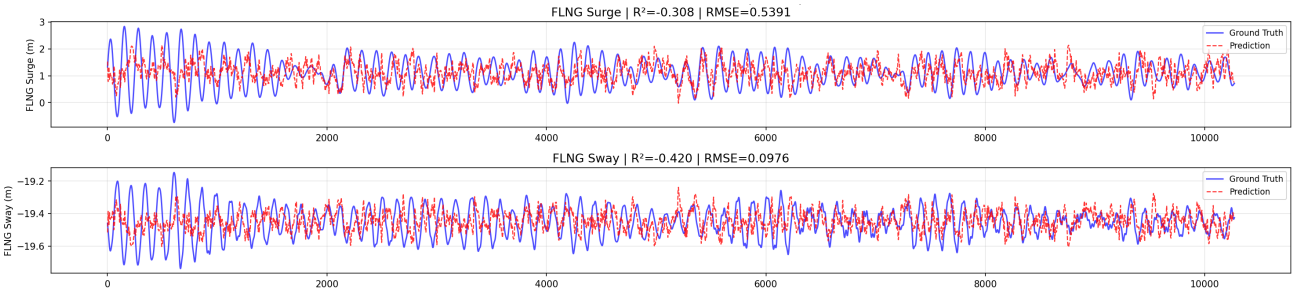


Figure 7.33: FLNG surge and sway predictions for unseen peak period $T_{\mathrm{p,sea}}$.

The LF motion correlation matrix in Figure 7.35 confirms this: the LF components of the training simulations share almost no linear relationship with the test case. This is expected, because each $T_{\mathrm{p,sea}}$ produces a distinct drift-force envelope driven by different spectral interactions. The comparison in Figure 7.36 further shows that LF surge responses for $T_{\mathrm{p,sea}} = 7, 8,$ and 10 differ not only in magnitude but also in envelope shape and drift behaviour. This reflects the theoretical framework in Chapter 3: LF drift forces depend on difference-frequency interactions, and when $T_{\mathrm{p,sea}}$ changes, the underlying frequency pairs $(f_i - f_j)$ that generate drift change as well [75, 67]. The model therefore has no stable cross-$T_{\mathrm{p,sea}}$ structure to learn from.

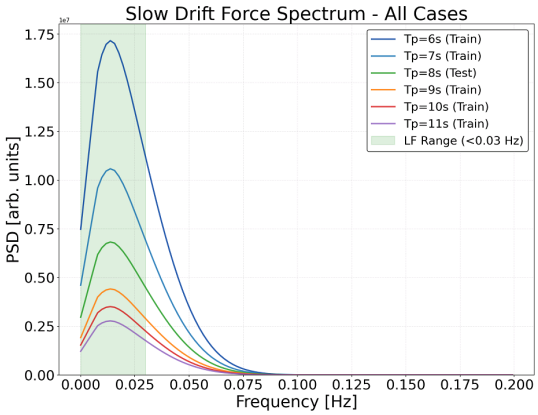


Figure 7.34: Low-frequency drift-force spectra for all training and test cases. The shaded region indicates the low-frequency range used for LF response modelling. Changes in $T_{\mathrm{p,sea}}$ reshape the drift-force spectrum rather than producing a simple amplitude scaling.

This behaviour is directly reflected in the low-frequency drift-force spectra. Figure 7.34 shows the slow-drift force spectra for all training and test cases. Although the low-frequency band is defined consistently, variations in $T_{\mathrm{p,sea}}$ lead to clear changes in both the magnitude and shape of the spectral energy within this band. The dominant low-frequency peaks shift and vary in strength across peak periods, indicating that the underlying drift-force excitation is fundamentally altered rather than uniformly scaled.

As a result, low-frequency vessel responses exhibit different envelope shapes and phase behaviour for each $T_{\mathrm{p,sea}}$, leaving no stable cross-condition structure for the model to learn. This explains the inability of the model to generalize LF motions to unseen peak periods.

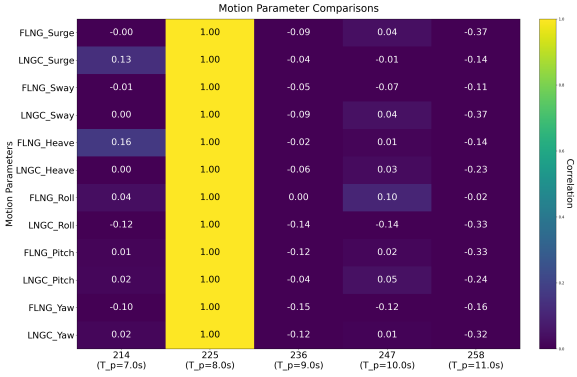


Figure 7.35: LF motion cross-correlation matrix for varying $T_{\mathrm{p,sea}}$.

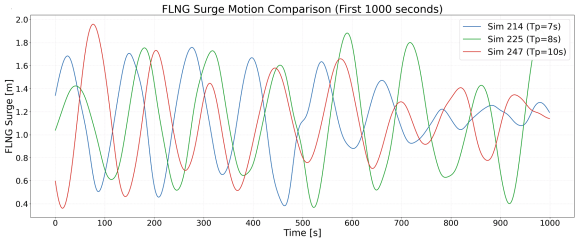


Figure 7.36: Wave-frequency motion correlation matrix for varying $T_{\mathrm{p,sea}}$ (FLNG).

Taken together, these results show why the model is unable to generalize low-frequency motions to unseen $T_{\mathrm{p,sea}}$, with limited training data. The prediction in Figure 7.33 already indicates that the model cannot reproduce the phase or amplitude evolution of the LF response. The correlation matrix in Figure 7.35 confirms this: the LF components of the training simulations share almost no linear relationship with the test case, meaning the model has no consistent cross-condition structure to learn from. This lack of similarity is also visible in Figure 7.36, where the LF surge responses for $T_{\mathrm{p,sea}} = 7, 8,$ and 10 differ not only in magnitude but also in envelope shape and drift behaviour.

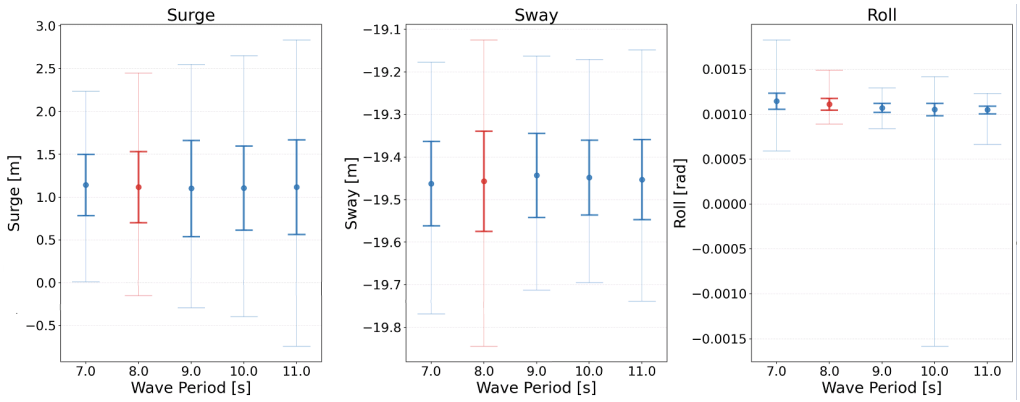


Figure 7.37: Comparison of LF responses for different $T_{p,sea}$, mean $\pm$ std bars (with min–max ranges).

Finally, Figure 7.26 highlights that the LF statistics vary noticeably with $T_{\mathrm{p,sea}}$, further demonstrating that the LF response is not governed by a simple scaling law. LF responses are governed by drift forces and multibody interaction effects that depend on the wave spectrum, so they change fundamentally with $T_{\mathrm{p,sea}}$ [76, 77]. Because these physical mechanisms vary from one $T_{\mathrm{p,sea}}$ to another, there is no consistent pattern for the model to learn, which explains why LF generalization across $T_{\mathrm{p,sea}}$ fails.

Wave-Frequency Predictions

For the wave frequency responses in FLNG Heave and Roll, the model is capable of predicting the signals well. For the heave motion, the model predicts the wave-frequency cycles with R ≈ 0.9 and RMSE ≈ 0.02. In roll, agreement is R ≈ 0.6; clear differeces in peak-height are visible. As discussed before, in head waves, the roll predictions are expected to perform less, as they are goverened by hydromechanical and mechanical coupling, rather than being directly forced

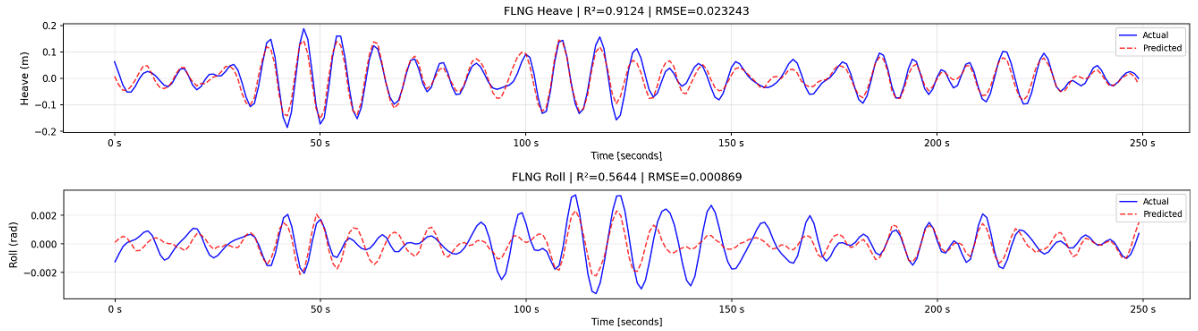


Figure 7.38: FLNG heave and roll predictions for unseen peak period $T_{\mathrm{p,sea}}$.

As can be seen, the prediction quality for the LNGC does not differ significantly from that of the FLNG. The model generally predicts the WF responses well when trained on different $T_{\mathrm{p,sea}}$ values (see Figure 7.39). This is notable, especially when contrasted with the poor LF generalization, because for WF motions the model must learn how differences in wave-frequency content translate into differences in the resulting motion response.

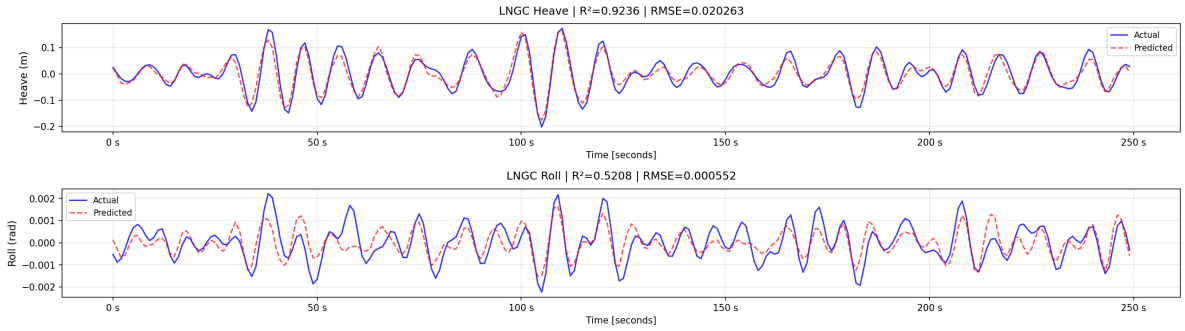


Figure 7.39: LNG heave and roll predictions for unseen peak period $T_{\mathrm{p,sea}}$.

Two factors explain this behaviour. First, the WF response is a direct and near-instantaneous reaction to the incoming wave. In a 3-hour simulation, the sea state contains a wide range of instantaneous periods around the nominal $T_{\mathrm{p,sea}}$ value. For example, a simulation with $T_{\mathrm{p,sea}} = 9$ s contains local periods ranging from 2.38 to 12.64 s. As shown in Table 7.6, all $T_{\mathrm{p,sea}}$ cases in the training set share a large portion of this local-period range. Consequently, the test case with $T_{\mathrm{p,sea}} = 8$ s overlaps substantially with the training cases ($T_{\mathrm{p,sea}} = 7, 9, 10, 11$ s), meaning that much of the incoming-wave content—and its corresponding vessel response—has effectively been seen before during training.

Second, WF motions provide far more training cycles in a 3-hour time series than LF motions. As summarised in Table 7.3, a WF signal contains roughly forty times as many cycles as an LF signal over the same duration. This gives the model many more supervised examples of wave–response relationships to learn from, which helps it to maintain phase and amplitude accuracy over multiple WF oscillations.

Together, these effects explain why WF predictions generalize well across different $T_{\mathrm{p,sea}}$ values, while LF predictions do not: WF behaviour has strong cross-condition similarity and abundant training examples, whereas LF behaviour varies substantially with $T_{\mathrm{p,sea}}$ and offers far fewer cycles from which the model can learn. This is consistent with the hydrodynamic theory in section 3.6: first-order (wave-frequency) forcing is governed by the instantaneous water-particle kinematics at the vessel and depends mainly on the local wave elevation, making it far less sensitive to changes in the spectral energy distribution than the second-order drift forces [72]

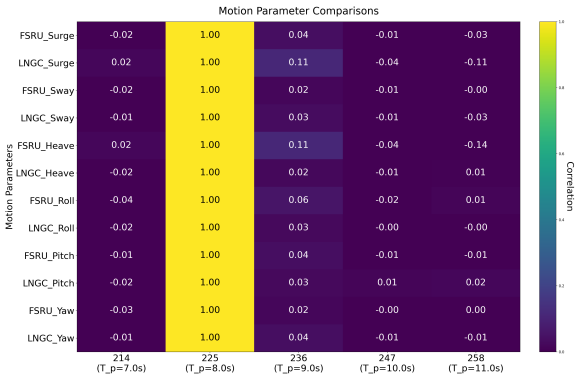


Figure 7.40: WF motion cross-correlation matrix for varying $T_{\mathrm{p,sea}}$.

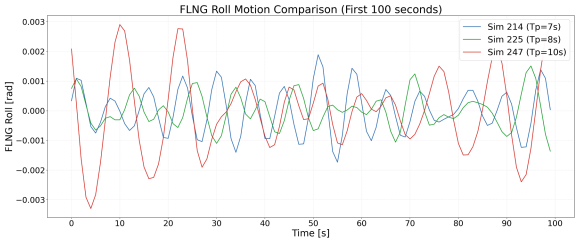


Figure 7.41: Wave-frequency roll motion comparison (FLNG).

Interpertation of Results

The results show a clear distinction between LF and WF generalisation to unseen $T_{p,\mathrm{sea}}$ values. LF responses cannot be reproduced for new peak periods: each $T_{p,\mathrm{sea}}$ produces a different drift-force envelope, and with only a few LF cycles per simulation the model has insufficient data to learn these variations. WF responses generalise well because different $T_{p,\mathrm{sea}}$ cases share substantial local period content, and the large number of WF cycles per simulation provides enough training examples to maintain phase and amplitude consistency. In short, WF dynamics transfer across peak periods while LF behaviour remains specific to the spectral regime represented in training.

Operability Assesment for Head Wave Conditions

The next step moves from component- or parameter-wise testing to operability assessment. Building on the previous findings—strong generalization across $H_{\mathrm{s,sea}}$, limited LF transfer across $T_{\mathrm{p,sea}}$ without broader spectral coverage—we now combine the LF and WF predictors to generate full motion time series and test them against operability criteria (e.g., limits on relative motions, line loads, and fender loads) used in the MARIN simulations. The objective is to evaluate, under head-wave conditions, how well the end-to-end ML pipeline reproduces the operability-relevant outcomes of the MARIN numerical model, rather than only matching individual DOF signals.

Accordingly, we assess *all* head-sea cases in the generic dataset (Table 4.2), totaling **756 simulations**. For each case, the LF and WF models produce the motion and tension time series that are then combined and propagated through the same post-processing steps used by MARIN (relative motions, line loads via the mechanical model) so that the predicted operability can be directly compared against the MARIN-derived operability for the identical condition. This provides a like-for-like benchmark at the level of criteria exceedance, not just signal fit.

7.2.2.1. Dataset Selection

For the operability assessment in head waves, the dataset is constructed using an evenly distributed selection of environmental parameters. Although operability analyses often focus on more severe sea states, the purpose of this initial assessment is to evaluate how a uniformly sampled dataset performs across the full range of conditions. Simulations (cases) close to operability limits are identified and marked with a near limit flag. During the selection of the 300 training cases, the process is designed to ensure a balanced representation of both near-limit and full-range cases in the training set, but not to over-represent near-limit cases.

The resulting parameter coverage is shown in Figures 7.42 and 7.43, which display the distributions of significant wave height, peak period, and wind velocity. The sampling strategy is generated using the k-means clustering method described in the Execution chapter.

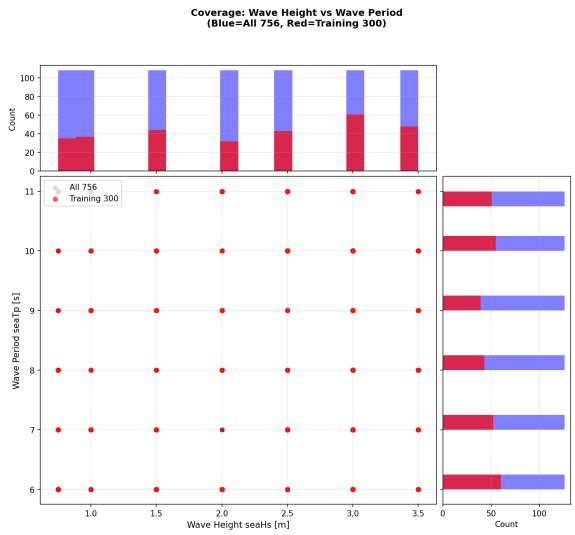


Figure 7.42: Coverage of significant wave height versus peak period ($H_{\mathrm{s,sea}}$–$T_{\mathrm{p,sea}}$).

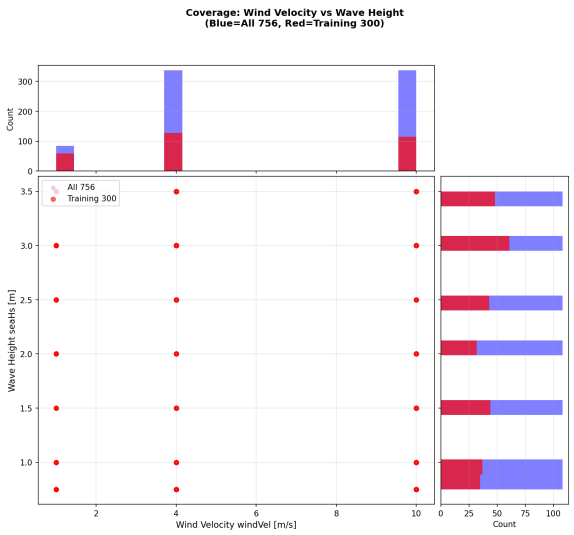


Figure 7.43: Coverage of wind velocity versus significant wave height (U–$H_{\mathrm{s,sea}}$).

7.2.2.2. Calm-Condition Excitations

Figures 7.44–7.47 show the model predictions under low-excitation (head-sea) conditions. These cases feature small response amplitudes in surge, sway, roll, and line tension, making prediction challenging. Because the signal-to-noise ratio is low and the vessel motions contain only weak excitation energy, the model is more sensitive to phase errors and noise in the input sequence. In such conditions, even small deviations in phase alignment can dominate the predicted response.

Table 7.7 summarises the environmental conditions for the calm sea state.

Table 7.7: Calm-condition environmental parameters for test `36490_01XMF_02_000_038_01`.

Wind vel.	Wind dir.	$H_{\mathrm{s,sea}}$	$T_{\mathrm{p,sea}}$	Wave dir.	Curr. vel.	Curr. dir.
1 m/s	90°	0.75 m	9 s	180°	0.12 m/s	90°

These calm conditions correspond to very low excitation energy and therefore produce small-amplitude responses in all degrees of freedom. The following plots (Figs. 7.44–7.47) illustrate the model predictions under this low-energy regime.

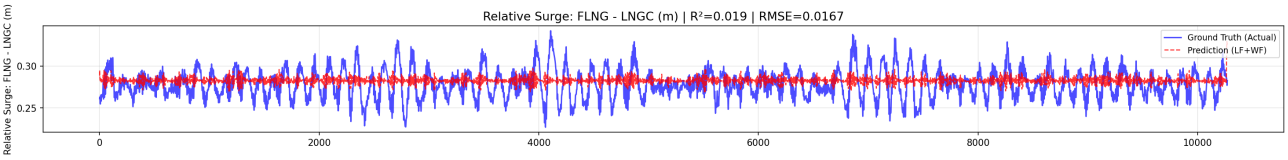


Figure 7.44: Calm-condition relative surge response: FLNG–LNGC prediction vs. reference ($R^2 = 0.039$, RMSE = 0.0167 m). The model captures the mean offset but not the detailed LF structure under low-excitation conditions.

The calm-condition surge response in Fig. 7.44 remains far below the operability limits, with a maximum of only 0.35 m compared to the ± 2 m manifold window. Because surge is dominated by low-frequency (LF) drift motion, very small LF excitation causes the wave-frequency (WF) component to become relatively more visible. In this regime the WF oscillations appear as noise relative to the small LF envelope, which explains the increased apparent scatter in the signals. It can also be observed that, under such low-excitation LF conditions, the LF model tends to predict only the mean surge offset rather than the detailed structure.

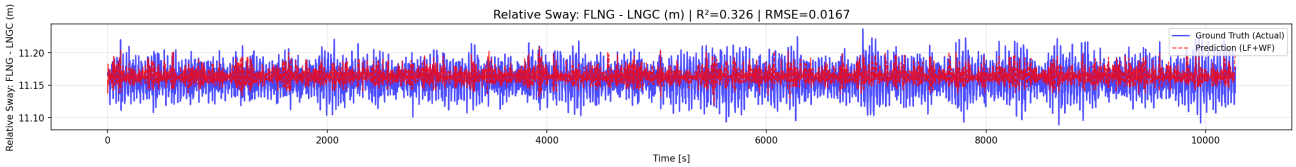


Figure 7.45: Calm-condition relative sway response: FLNG–LNGC prediction vs. reference ($R^2 = 0.326$, RMSE = 0.0167 m). The model reproduces the mean sway offset reasonably well under low-excitation conditions.

The relative sway response in Fig. 7.45 shows an even smaller amplitude of about 0.10 m, consistent with head seas providing minimal sway excitation. With little temporal structure to learn from, minor phase offsets between prediction and reference become visible, although the motions remain well within safe bounds.

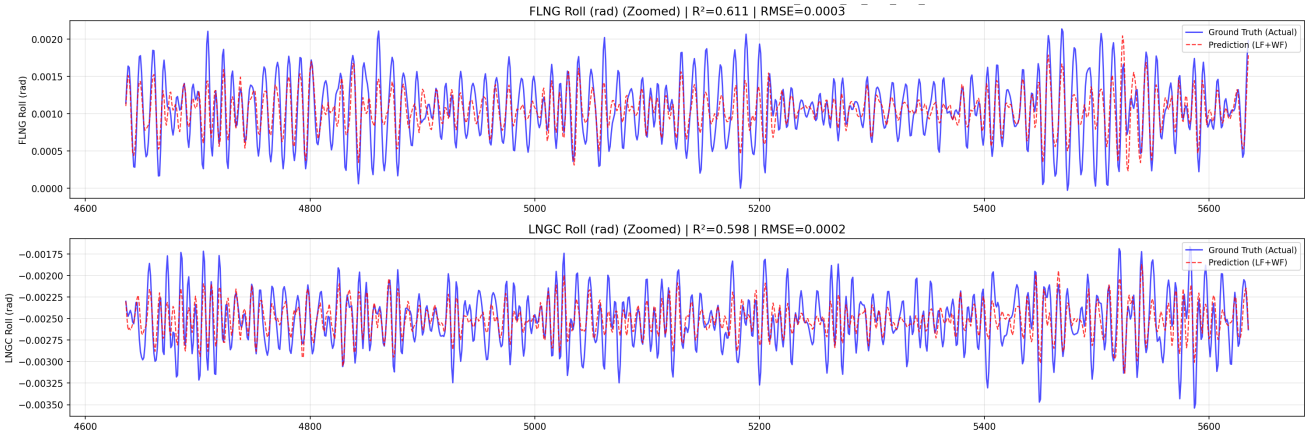


Figure 7.46: Calm-condition roll response: FLNG (top, $R^2 = 0.611$, RMSE = 0.0003 rad) and LNGC (bottom, $R^2 = 0.598$, RMSE = 0.0002 rad) prediction vs. reference.

The roll motion in Fig. 7.46 is dominated by small wave-frequency oscillations, with the reference peaking far below the ±0.0524 rad criterion. Under such weak excitation, phase differences between prediction and reference become more noticeable. The model slightly underpredicts the roll, but the response remains well inside operability limits.

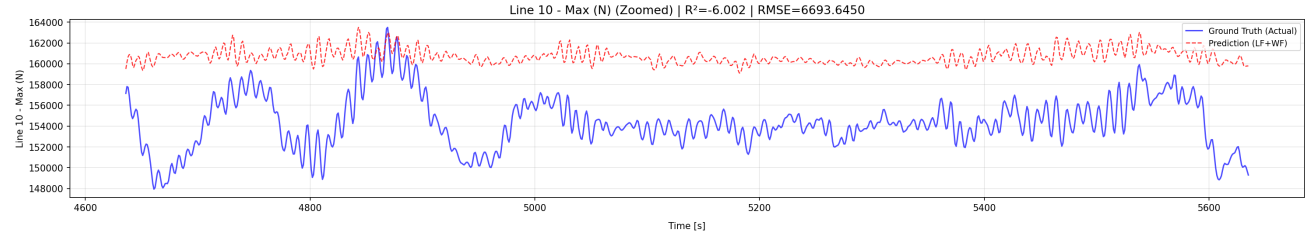


Figure 7.47: Calm-condition mooring line tension response for Line 10: prediction vs. reference ($R^2 = -6.002$, RMSE = 6693.6 N). Despite the low R^2, the model correctly captures the maximum tension level.

Figure 7.47 shows how the model performs the max tension signals. The model only capture the max tensions and barely any oscillations.

In calm seas, the responses are dominated by low excitation energy, small wave-frequency motions, weak drift forces, and limited dynamical variation. Because the transformer relies on temporal patterns, these subtle excitations provide very little structure to learn from. As a result, prediction scatter becomes more visible, even though the underlying trend of the motion is still captured correctly.

7.2.2.3. Critical-Condition Excitations

Figures 7.48–7.51 show predictions under high-excitation head-sea conditions, close to the operability limits. In contrast to the calm cases, the stronger LF and WF components provide clearer temporal structure. The transformer benefits from these stronger physical signals, resulting in more stable and accurate predictions.

Table 7.8 summarises the environmental conditions for the calm sea state.

Table 7.8: Critical-condition environmental parameters for test.

Wind vel.	Wind dir.	$H_{\mathrm{s,sea}}$	$T_{\mathrm{p,sea}}$	Wave dir.	Curr. vel.	Curr. dir.
10 m/s	315°	3.5 m	11 s	180°	0.24 m/s	315°

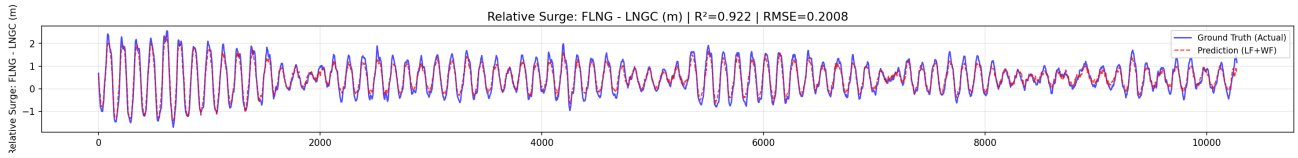


Figure 7.48: Critical-condition relative surge response: FLNG–LNGC prediction vs. reference ($R^2 = 0.922$, RMSE = 0.208 m).

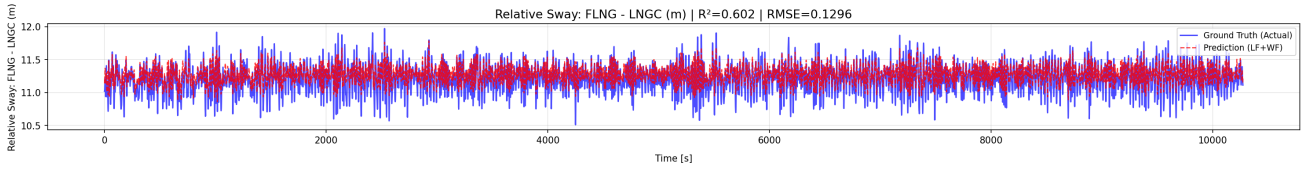


Figure 7.49: Critical-condition relative sway response: FLNG–LNGC prediction vs. reference ($R^2 = 0.602$, RMSE = 0.1296 m).

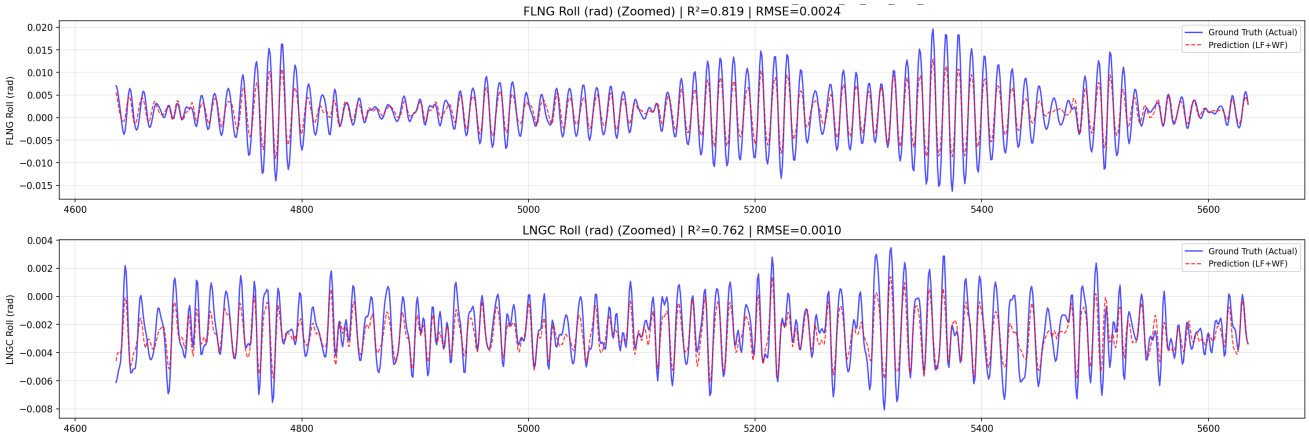


Figure 7.50: Critical-condition roll response: FLNG (top, $R^2 = 0.819$, RMSE = 0.0024 rad) and LNGC (bottom, $R^2 = 0.762$, RMSE = 0.0010 rad) prediction vs. reference.

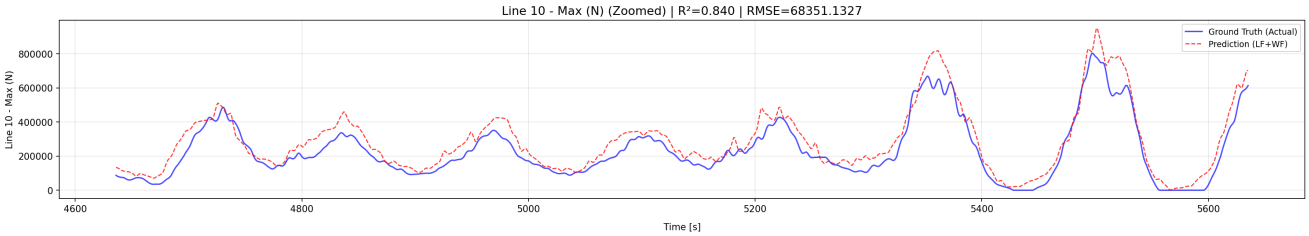


Figure 7.51: Critical-condition mooring line tension response for Line 10: prediction vs. reference ($R^2 = 0.840$, RMSE = 68,351 N).

Why critical conditions are easier to predict: When wave excitation is stronger: - the signal dominates noise, - LF drift motion becomes more pronounced, - WF components are more coherent, - and

the model receives clearer temporal patterns.

Thus, high-excitation cases exhibit improved predictive accuracy.

7.2.2.4. Overall Operability Classification

The transformer-based model performs excellently across all 456 head-sea states, achieving a total accuracy of **99.6%**. Only two misclassifications occur: - one *False Safe* (unsafe state predicted as safe), - and one *False Unsafe* (safe state predicted as unsafe).

This is illustrated in the compact 2×3 confusion matrix in Fig. 7.52. The near-perfect performance highlights that the model captures the dominant features governing operability limits.

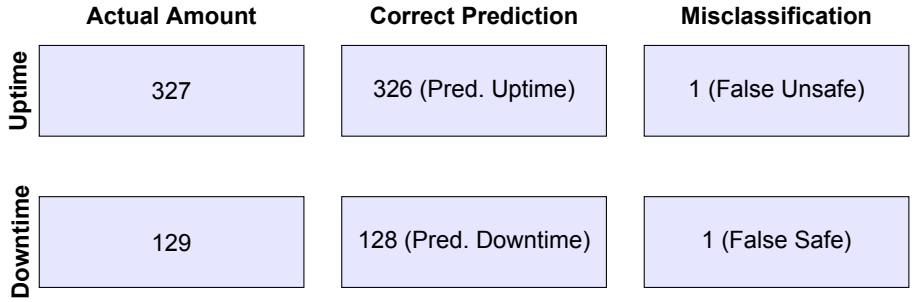


Figure 7.52: Confusion-matrix summary of uptime/downtime classifications for all 456 head-sea states.

These results show strong agreement with MARIN reference values. SBS X is the only criterion showing a modest underprediction trend, but overall physical consistency is preserved.

False Safe Rate Concentrates in High Sea States. Beyond the overall 99.6% accuracy, a critical pattern emerges when misclassifications are examined per operability criterion.

Peak-period effects show a similar trend: when **a response limit is exceeded**, short-period waves ($T_{\mathrm{p,sea}} = 6$–7 s) produce a 35.0% False Safe rate, compared to only 5.3% at $T_{\mathrm{p,sea}} = 11\,\mathrm{s}$. Comprehensive overview tables of difficult to predict environmental conditions are given in Appendix E.

This pattern corresponds closely to findings from hydrodynamic studies on side-by-side moored vessels. Huijsmans, Pinkster and de Wilde (2001) report pronounced low-frequency drift-force peaks around $T_{\mathrm{p,sea}} \approx 6.3$–7.4 s, which produce large transverse drift forces and gap-driven amplification effects [75].

Similarly, Buchner et al. (2001) show for a similar system that in the high-sea condition $H_{\mathrm{s,sea}} = 3\,\mathrm{m}$, $T_{\mathrm{p,sea}} = 7\,\mathrm{s}$, both relative sway and yaw responses increase dramatically even in head waves, driven by wave reflection from the shuttle bow and strong second-order drift forcing [76]. These amplified motions result in elevated tension responses and nonlinear fender–mooring coupling.

The model's 99.6% overall accuracy therefore masks a concentration of risk in the extremes. In practical operations, the model is *safest* in calm conditions (where no violations occur anyway) but requires careful monitoring in high-energy, short-period regimes where operational limits are tightest and physics is most nonlinear.

7.2.2.5. Conditional Accuracy: Severity-Dependent Performance

Performance improves markedly in calm and moderate seas ($H_{\mathrm{s,sea}} < 2.5$ m, $T_{\mathrm{p,sea}} > 9$ s), where False Safe rates remain below 10%. However, in the challenging high-energy regime ($H_{\mathrm{s,sea}} > 3.0$ m, $T_{\mathrm{p,sea}} < 7$ s)— precisely where operability restrictions are most stringent—the model underpredicts violations in nearly 60% of individual cases. This systematic underprediction in difficult conditions, while not violating overall safety (only 1 False Safe in the full test set), suggests that **augmenting training data in severe conditions** would improve robustness, particularly through physics-informed constraints on nonlinear mooring dynamics and gap-resonance effects.

In the operability-assessment regime (456 head-sea cases), this asymmetry manifests as:

- High WF accuracy ($R^2 \sim 0.8$–0.9 for heave/pitch) covering all conditions
- Variable LF accuracy concentrated in high-$H_{\mathrm{s,sea}}$, short-$T_{\mathrm{p,sea}}$ cases (where false-safe errors occur).

For a model trained only on head seas, the operability prediction remains safe and accurate. But extending to other sea directions or spectral regimes would require substantial additional $T_{\mathrm{p,sea}}$ and directional training.

Data Efficiency vs. Physics Fidelity Trade-off. The transformer-based pipeline demonstrates that 300 out of 756 simulations are sufficient for operability classification in head-sea conditions. This efficiency arises because:

- WF components are highly predictable with sparse training due to high cycle-count and cross-condition overlap
- Head-sea geometry isolates drift-force complexity to surge; sway and yaw remain weak

However, near-perfect accuracy masks parameter-level weakness. The parameter specific false-safe rate rises to 60% in high-$H_{\mathrm{s,sea}}$, short-$T_{\mathrm{p,sea}}$ conditions precisely because full-spectrum drift forces are underpredicted.

7.2.2.6. Summary

The model demonstrates strong operability prediction capability across the full dataset of 456 head-sea states. It aligns well with MARIN results, handles both calm and critical conditions reliably, and exhibits very low safety-related misclassification rates. Although further improvement is possible in predicting SBS X and selected mooring-line tensions, the transformer architecture proves highly suitable for data-driven operability assessment.

Generalizing to new θ_{sea}

Now that the model has shown excellent operability assessment performance when the training data for head seas is expanded, the next step is to evaluate its ability to generalize to different incoming wave directions. A change in θ_{sea} fundamentally alters the direction and distribution of the excitation forces acting on the vessels, in a way that is entirely different from how varying $T_{\text{p,sea}}$ modifies the spectral content of the waves.

	Train										**Test**
$\boldsymbol{\theta_{\text{sea}}}$ (deg)	90	105	120	140	160	200	220	240	255	270	**180**

Wind vel.	Wind dir.	$H_{\text{s,sea}}$	$T_{\text{p,sea}}$	Curr. vel.	Curr. dir.
1 m/s	90°	1.5	8 s	0.12 m/s	90°

Table 7.9: Typical periods and constant environmental parameters across training and predicted cases.

The training and test cases are listed in Table 7.9. Only the propagation direction of the incoming waves is changed; all spectral parameters remain identical. As illustrated in Figure 7.53, the wave elevation signal is exactly the same for all considered headings because only the direction vector changes. This ensures that differences in the resulting motions are driven purely by directional effects and not by variations in wave energy.

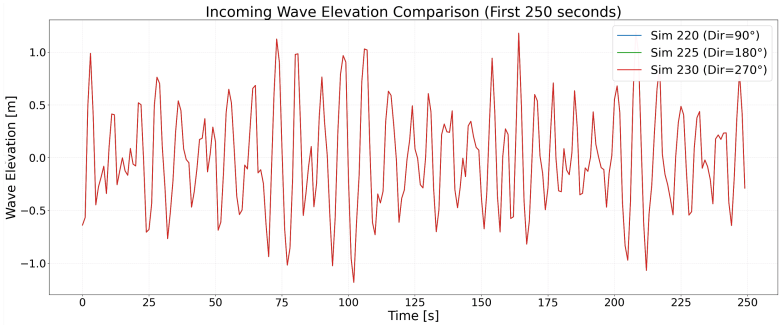


Figure 7.53: Comparions of wave signal for different $H_{\text{s,sea}}$.

7.2.2.7. Low-Frequency Predictions

The low-frequency predictions for the FLNG are shown in Figure 7.54, while the LNGC predictions are displayed in Figure 7.55. For both vessels, the surge response is predicted reasonably well, with $R^2 \approx 0.8$. This behaviour is expected and reflects the directional coherence of longitudinal drift forcing. In head seas, surge is directly excited by second-order wave drift forces aligned with the incident wave direction [75]. Because drift forces maintain a nearly collinear relationship with the vessel's longitudinal axis across all approach headings, the low-frequency envelope structure and phase behaviour show consistency across beam-sea and quartering training cases. Consequently, the model successfully transfers its learned surge dynamics to the unseen head-sea test condition.

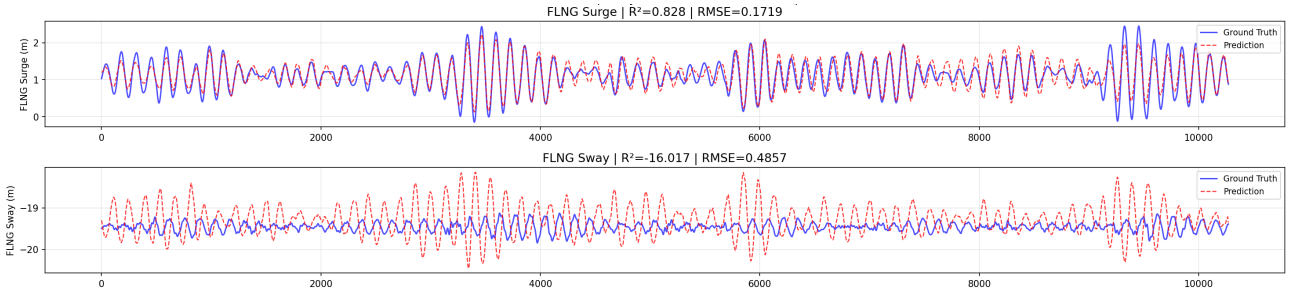


Figure 7.54: FLNG surge and sway predictions vs. reference for unseen θ_{sea}.

In contrast, LF sway is predicted extremely poorly ($R^2 \approx -16$). This reflects the strong heading dependence of low-frequency drift forces. As shown in Figure 7.56, LF sway responses differ substantially between directions: in beam seas ($\theta_{\text{sea}} = 90°, 270°$) strong second-order drift forces act directly in the transverse direction, while in head seas ($180°$) the lateral excitation is minimal [67]. This explains why a model trained on different incoming wave directions and not head seas fails to reproduce the weak sway observed in head seas.

Chapter 3 describes how the narrow gap between the hulls can amplify lateral motion through free-surface oscillations and momentum exchange [73, 74]. These gap-driven interaction effects are strongest when waves approach from the side, effectively acting as a transverse resonator in beam and quartering seas. In head seas, this mechanism becomes inactive, and sway arises only from weaker indirect coupling, such as diffraction and radiation effects from heave and pitch.

This lack of cross-direction similarity is also evident in Figure 7.57, where each heading exhibits a distinct LF sway envelope. The statistical behaviour in Figure 7.58 further highlights a sharp transition around the head-sea direction: surge varies smoothly with heading, whereas sway shows a marked drop in amplitude. This confirms that LF sway in head seas represents a different excitation regime, and therefore cannot be reliably inferred from beam-sea training data.

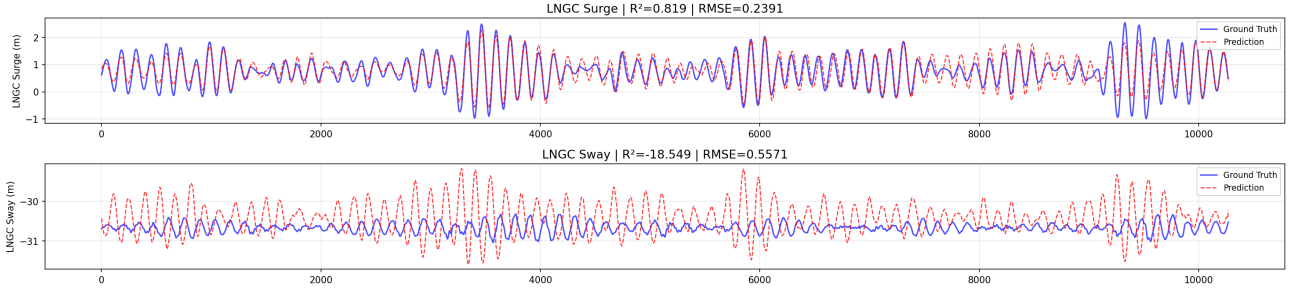


Figure 7.55: LNGC surge and sway predictions vs. reference for unseen θ_{sea}.

Because the training set contains several beam-sea cases with strong sway energy, the model learns that many conditions produce significant lateral drift. When tested on the head-sea case, where the true sway excitation is nearly zero, the model therefore heavily overpredicts sway. The absence of cross-direction structure in the LF band prevents meaningful generalization.

The correlation matrix (Figure 7.56) shows this clearly: the sway response in beam seas has almost no similarity to the sway response in head seas, so the model has no examples that represent the low-excitation regime.

This behaviour agrees with the multibody hydrodynamics discussed in chapter 5. In a side-by-side

configuration, sway and yaw motions are strongly affected by vessel proximity and by second-order drift forces that depend on wave direction [77]. These drift forces and interaction effects change significantly between beam and head seas [75, 76]. Because the physics itself is heading-dependent, the LF sway response in head seas is not just a smaller version of the beam-sea response-it is governed by different mechanisms. The model's difficulty therefore reflects this physical change, not just a statistical limitation.

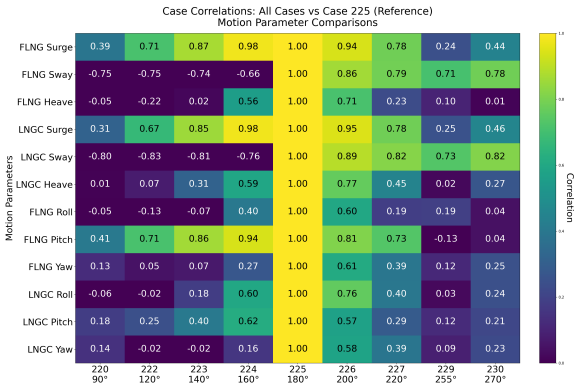


Figure 7.56: LF motion cross-correlation matrix for varying θ_{sea}.

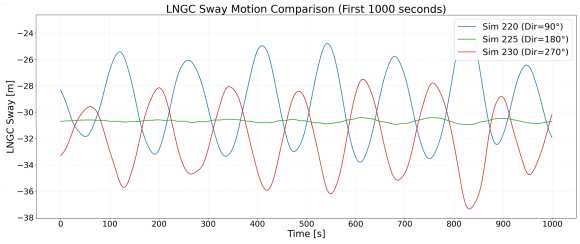


Figure 7.57: LF sway motion comparison across new wave-direction cases (LNGC).

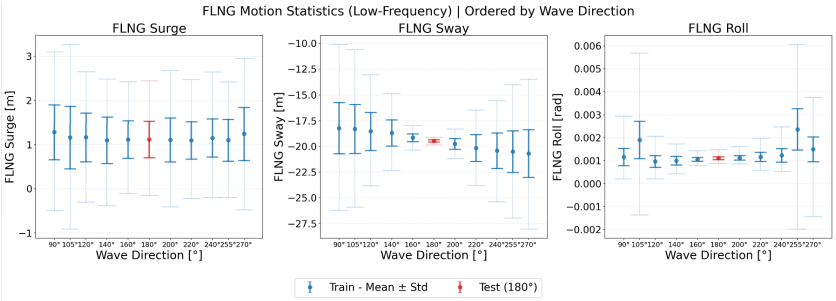


Figure 7.58: LNGC surge and sway predictions vs. reference for unseen θ_{sea}.

7.2.2.8. Wave-Frequency Predictions

From the previous generalization tests on varying $H_{\mathrm{s,sea}}$ and $T_{\mathrm{p,sea}}$, it was learned that the WF model requires substantially less data to generalize well. First-order wave effects provide many more oscillation cycles and exhibit stronger cross-condition similarity, so the amount of training data used for the wave-direction WF model was intentionally reduced. The selected training and test directions are:

	Train					**Test**
$\boldsymbol{\theta_{\mathrm{sea}}}$ (deg)	90	140	200	220	270	**180**

The FLNG and LNGC predictions in Figure 8.1 show heave reproduced with high accuracy and roll with moderate but consistent agreement. This follows from the nature of first-order forcing: WF excitation depends on the instantaneous incoming wave elevation and local water-particle motion rather than the global propagation direction [72]. Heave should remain largely direction-independent, enabling a robust mapping that transfers to unseen headings. Roll behaves as an intermediate case: amplitudes vary

with direction, yet the oscillatory structure remains coherent because the same radiation and diffraction mechanisms govern roll across headings [67].

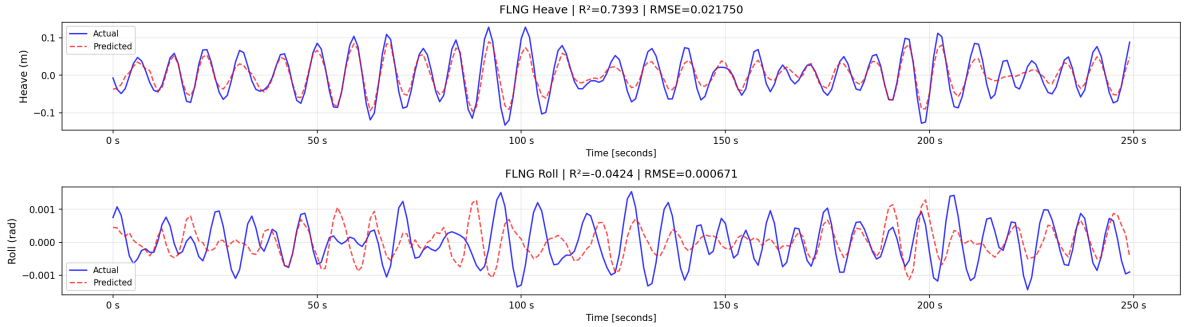


Figure 7.59: FLNG heave and roll WF predictions vs. reference for unseen wave direction θ_{sea}.

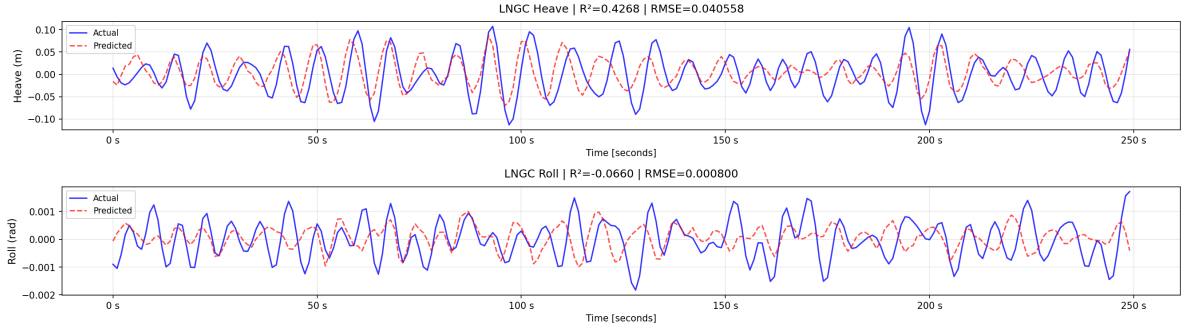


Figure 7.60: LNGC heave and roll WF predictions vs. reference for unseen wave direction θ_{sea}.

The WF cross-correlation matrix in Figure 7.61 shows why roll has more difficulty to deliver steady predictions, as there is limited correlation with the training signals.

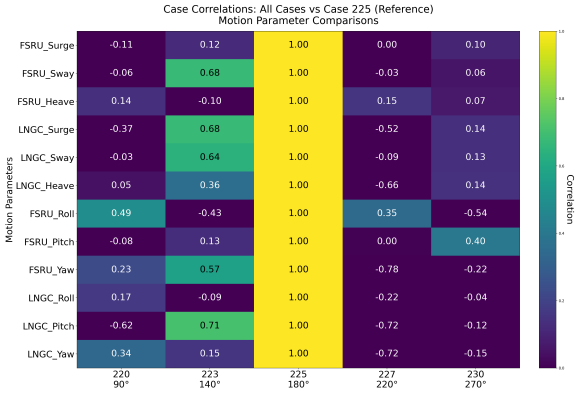


Figure 7.61: WF motion cross-correlation matrix for varying θ_{sea}.

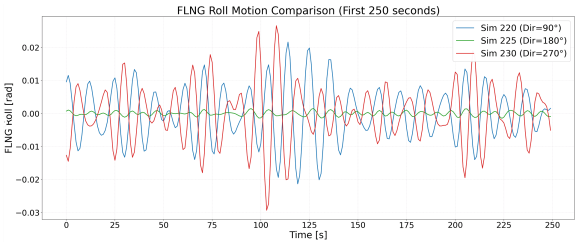


Figure 7.62: Wave-frequency roll motion comparison across new wave-direction cases (FLNG).

The results highlight a clear distinction between low-frequency and wave-frequency generalization to unseen wave directions:

- **Low-frequency motions do not generalize across directions:** LF drift forces depend strongly on θ_{sea} and change magnitude, sign, and envelope shape between beam, quartering, and head

seas. The correlation structure in Figure 7.56 and the motion differences in Figure 7.57 show that LF responses across headings share almost no similarity. Consequently, the model overpredicts sway in head seas and cannot learn a consistent cross-direction LF representation from the limited dataset.

- **Wave-frequency motions generalise across directions with limitations:** WF heave and roll are directly forced by the instantaneous wave elevation, which provides strong cross-condition similarity and abundant training cycles. The WF correlation patterns in Figure 7.61 and the roll comparison in Figure 7.62 confirm that the WF band retains overlap to transfer to unseen headings. However, changes in wave direction modify the vessel response amplitude operators by altering the geometric relationship between the incident wave field and the vessel orientation. These direction-dependent RAO changes, propagated through multibody interaction and radiation-diffraction coupling, limit roll transferability in particular. Heave remains largely direction-independent, while roll amplitude accuracy reduces for headings further from the training set.

In summary, the model can reliably generalize wave-frequency responses to new θ_{sea} values, but low-frequency drift motions do not generalize across headings. LF generalization would require either a broader directional training set or physics-informed modelling of drift forces. WF dynamics, on the other hand, are transferable under relatively sparse directional sampling due to their strong waveform similarity and high cycle count within a 3-hour record.

7.2.3. Interpertation of Results

The results presented in this chapter show a clear distinction between the generalizability of low-frequency (LF) and wave-frequency (WF) vessel motions across environmental conditions. This behaviour aligns with the hydrodynamic mechanisms discussed earlier and with established findings on multibody drift forces and SBS interactions [72, 75, 67, 76].

WF generalization is robust WF motions (heave, roll, pitch) generalize well across unseen $H_{\mathrm{s,sea}}$, $T_{\mathrm{p,sea}}$, and wave directions. This is because:

- 3 h simulations provide hundreds of WF cycles, giving the model abundant training examples;
- different $T_{\mathrm{p,sea}}$ or directional cases still share substantial overlap in local wave periods, reflecting the broad JONSWAP spectral support [109].

As a result, WF predictions remain accurate even for conditions outside the training set.

LF generalization is limited and condition-dependent LF surge, sway, and yaw generalize only in the restricted case of varying $H_{\mathrm{s,sea}}$ with fixed $T_{\mathrm{p,sea}}$, direction, and wave seed. In all other scenarios, LF prediction fails due to:

- the strong dependence of drift forces on the wave spectrum, which changes fundamentally with $T_{\mathrm{p,sea}}$ [75, 67];
- the directional sensitivity of multibody drift forces and gap-flow effects in SBS configurations [76, 77];
- the limited number of LF cycles (10–20 per 3 h), which restricts learnable temporal structure.

Consequently, LF dynamics do not transfer across $T_{\mathrm{p,sea}}$ or θ_{sea} without substantially broader training coverage.

Difficult conditions reveal where errors concentrate Although overall operability classification in head seas is very high (99.6%), the analysis shows that misclassifications are not uniformly distributed. False Safe errors cluster in high-energy conditions: large $H_{\mathrm{s,sea}}$ and short $T_{\mathrm{p,sea}}$, where drift forces and nonlinear SBS interaction effects are strongest [75, 76]. These regimes require targeted data selection for training.

Overall insight WF dynamics are transferable and require relatively little data, while LF dynamics are spectrum- and direction-dependent and require broad environmental coverage to generalize. In head seas, where LF behaviour is consistent, the full operability pipeline performs reliably. Beyond head seas, accurate operability prediction will require expanding the training space, particularly in $T_{\mathrm{p,sea}}$ and directional variability, or incorporating physics-based priors on drift forces and multibody interactions.

8

Final Results

The previous chapters established how machine-learning models behave when trained on summary statistics (SSBL) or time series (TSBL) and how these two paradigms differ in their ability to generalise across environmental conditions. The time series generalisation studies on $H_{s,\text{sea}}$, $T_{p,\text{sea}}$, and θ_{sea} revealed a clear asymmetry: while wave-frequency responses generalised reliably across conditions, low-frequency drift motions only generalised for simulation with new significant wave heights $H_{s,\text{sea}}$. Despite these limitations, the operability assessment in head seas demonstrated that near-perfect agreement with the MARIN numerical model is achievable even when several operability-limiting parameters are predicted with only moderate accuracy. This finding motivates a deeper investigation into the minimal amount of data required to achieve reliable operability prediction across diverse environmental conditions.

The present chapter therefore addresses the following sub-question:

- How can we construct a minimal dataset which captures all conditions and system dynamics?

To answer this question, this chapter applies the lessons learned throughout the earlier analysis. In particular, it uses the observed dependencies of LF and WF generalisation to design a minimal training dataset that is representative across the key environmental axes, yet significantly smaller than the dataset used in the head-sea operability assessment. By intentionally reducing the amount of training data, the models are subjected to a more challenging scenario, allowing the comparison between SSBL and TSBL to be evaluated under conditions where data availability is limited.

This final chapter thus aims connects the theoretical, methodological and practical aspects of the research into a cohesive assessment of data-efficient operability prediction.

8.1. Dataset Selection

The objective is to construct a minimal training dataset that can predict operability across diverse sea states using the smallest number of representative time series (or summary statistics) that still capture the full range of vessel–environment interactions.

The following insights obtained throughout this work guide the selection:

- **Training fraction:** The SSBL study (section 7.1) showed that using approximately 20% of the total generic dataset yields most of the achievable performance, with only marginal improvement from additional training data.
- **Generalisation gaps (TSBL):** The TSBL results (section 7.2) showed that the LF model exhibits reduced generalisation when extrapolating to unseen peak periods $T_{p,\mathrm{sea}}$ and wave directions θ_{sea}. These dimensions therefore require deliberate coverage in the training set.
- **Head-sea operability:** The head-wave analysis (Figure 7.2.2) demonstrated that a subset of approximately 300 out of 756 sea states was sufficient to predict operability in head seas, provided the subset spanned the full response range.
- **Response-range coverage:** Accurate operability prediction requires that the training set spans the full observed range of outputs, including motions, relative motions, and line tensions, not only the input parameter space. This includes coverage of the variation in mean surge and sway positions, as accurate predicted offsets strongly influence the resulting relative motions.

Hydrodynamic context: The operability assessment in head waves revealed that short peak periods ($T_{p,\mathrm{sea}} \approx 6$–$7\,\mathrm{s}$) combined with high sea states ($H_{s,\mathrm{sea}} \geq 3\,\mathrm{m}$) are particularly challenging, with pronounced nonlinear responses and elevated mooring tensions. This aligns with hydrodynamic findings for side-by-side offloading: Huijsmans, Pinkster, and de Wilde (2001) report pronounced low-frequency drift-force peaks at $T_p \approx 6.3$–$7.4\,\mathrm{s}$, producing large transverse drift and gap-driven amplification [75]; Buchner et al. (2001) show that for $H_s = 3\,\mathrm{m}$, $T_p = 7\,\mathrm{s}$, relative sway and yaw increase strongly even in head seas due to wave reflection and second-order drift forcing [76]. These amplified motions lead to higher tensions and nonlinear fender–mooring coupling.

8.1.1. Selection Methodology

Starting from the full generic set of 8317 simulations, we used the insights described above together with an analysis of how well different subsets cover both the input space and the response range. Based on this, we selected a compact training subset of $N_{\mathrm{train}} = 1500$ simulations. The selection follows a two-stage procedure to ensure (i) deliberate coverage of challenging regimes in $T_{p,\mathrm{sea}}$ and θ_{sea}, and (ii) explicit coverage of the full output response range (motions, relative motions, and tensions).

1. **Stratified sampling over sea-state axes.** We partition the data into bins in the $(T_{p,\mathrm{sea}}, \theta_{\mathrm{sea}})$ plane. Each bin receives an explicit quota so short periods and all headings are represented. Within bins, weighted random sampling prioritises difficult regimes using the scheme in Table 8.1. Calm conditions ($H_{s,\mathrm{sea}} < 1.0\,\mathrm{m}$) are excluded; higher sea states and shorter periods receive larger weights.
2. **Response-space refinement (K-means).** From the stratified pool, we cluster simulations in a response space built from manifold relative motions and key motions (e.g., FLNG/LNGC surge and sway). For each cluster, we keep the simulation closest to the centroid. This step ensures broad coverage of extremes and regions close to operability limits (large relative motions, high tensions, FLNG and LNGC roll).

Parameter weighting scheme. The weighting emphasises short periods and high sea states (Table 8.1). Calm sea states are excluded; progressively higher weights are assigned to increasing $H_{s,\mathrm{sea}}$ and decreasing $T_{p,\mathrm{sea}}$.

Table 8.1: Parameter weighting scheme for dataset selection.

Significant Wave Height ($H_{s,\mathrm{sea}}$)

Range	Weight
< 1.0 m	0.0 (excluded)
1.0–1.5 m	0.7
1.5–2.0 m	1.0
2.0–2.4 m	1.5
2.4–2.6 m	4.0
2.6–3.0 m	4.5
≥ 3.0 m	5.5

Peak Period ($T_{p,\mathrm{sea}}$)

Range	Weight
6–7 s	20.0
7–8 s	10.0
8–9 s	4.0
9–10 s	2.0
10–11 s	1.0
> 11 s	0.5

Directional rationale: Although beam and quartering seas often induce the strongest excitation, the present spread-moored configuration is operated predominantly in (near) head seas. To avoid directional bias and promote generalisation, θ_{sea} is given uniform weight so all sectors are sampled evenly. If the project context implied predominantly beam/quartering seas, a turret-moored concept (weathervaning) would be more appropriate, and the directional weights would be adapted accordingly.

Full operability emphasis: We apply an additional multiplier to simulations close to operability boundaries: extreme manifold positions, large roll angles ($> 2.5°$), and high line tensions (approaching the 50% MBL limit). This ensures the model sees sufficient edge-case behaviour, where prediction accuracy most affects safety decisions.

To illustrate how well the selected $N_{\mathrm{train}} = 1500$ simulations represent the variability in the full dataset, several key inputs and responses are displayed below. These plots show how the training subset distributes across the main environmental and response dimensions relevant for operability assessment.

Sea-state parameter coverage: Figure 8.2 displays the coverage of significant wave height $H_{s,\mathrm{sea}}$ and peak period $T_{p,\mathrm{sea}}$. The selected subset follows the shape of the full distribution and includes all important regions, including the short-period range where nonlinear responses are strongest.

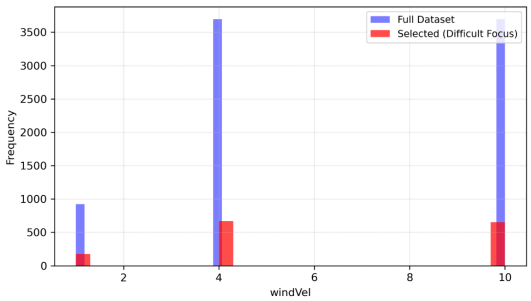
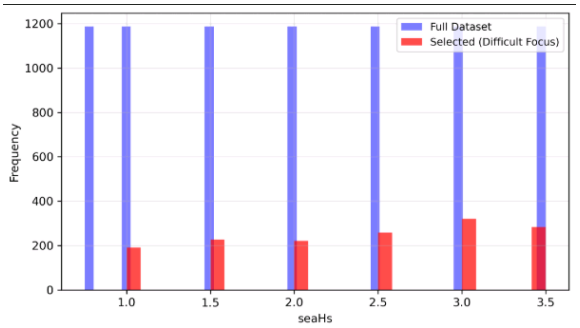
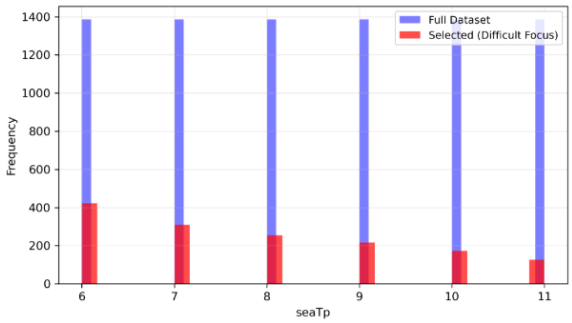


Figure 8.1: Coverage of significant wave height V_{wind} (m) in the selected dataset (red) relative to the full dataset (blue).



(a) Coverage of significant wave height $H_{s,\mathrm{sea}}$ (m) in the selected dataset (red) relative to the full dataset (blue).



(b) Coverage of peak period $T_{p,\mathrm{sea}}$ (s) in the selected subset (red) relative to the full dataset (blue).

Figure 8.2: Displayed coverage of key sea-state parameters in the selected training subset.

Mean offset coverage (surge and sway): Mean surge and sway offsets strongly influence relative motions and mooring loads. The plots in Figure 8.3 and Figure 8.4 show that the selected subset spans almost the full range of offsets observed for both the FLNG and LNGC.

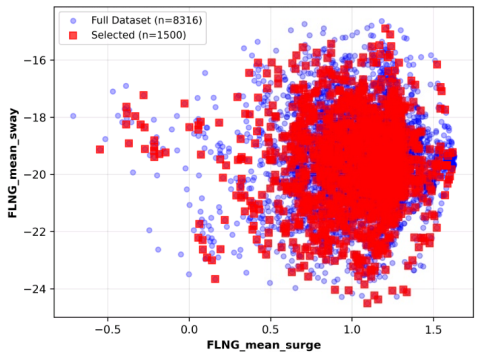


Figure 8.3: Coverage of mean surge vs. mean sway for the FLNG.

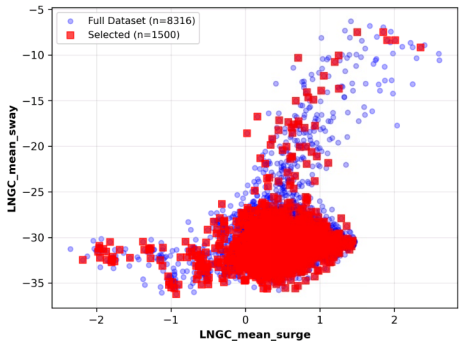


Figure 8.4: Coverage of mean surge vs. mean sway for the LNGC.

Coverage of response extremes: Finally, Figure 8.5 and Figure 8.6 show the coverage of two operability-critical responses: maximum line tension and FLNG roll. The selected subset displays good representation of the upper range of both responses, ensuring that extreme behaviour is sufficiently present in the training data.

Table 8.2: Coverage of selected vs. full range for critical parameters

Metric	FLNG roll		LNGC roll		Rel. manifold surge		Rel. manifold sway		Line tension	
	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max
Coverage [%]	8.3	94.9	7.2	90.1	15.2	96.4	4.5	100.0	1.8	84.8

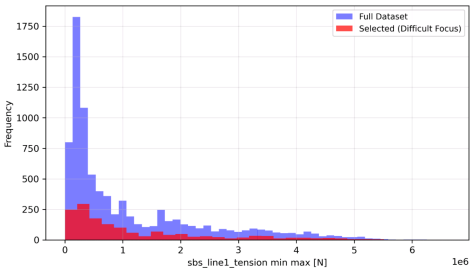


Figure 8.5: Coverage of maximum mooring line tension in the selected subset (red) relative to the full dataset (blue).

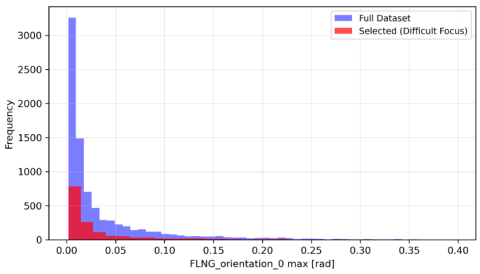


Figure 8.6: Coverage of FLNG roll in the selected subset (red) compared to the full dataset (blue).

Resulting minimal set: Combining the empirical rules ($\sim 20\%$ sufficiency, targeted coverage of $T_{p,\mathrm{sea}}$ and θ_{sea}, and explicit response-range coverage) with the two-stage selection procedure yields a compact training set of

$$N_{\mathrm{train}} = 1500 \text{ simulations},$$

which balances data efficiency with robust coverage of both input space and operability-relevant response regimes.

8.2. SSBL Operability Assessment

To assess whether the SSBL model can support operability decisions, the key operability drivers including roll, relative surge and sway, line tension, and fender loads are predicted and compared against their MARIN-computed values. The main prediction metrics are listed in the Appendix in Table F.1 and the operability classification matrix is presented below.

Overall, the SSBL model shows strong predictive performance. Roll and relative motions achieve R^2 values above 0.94 with low scaled RMSE. Line tension prediction is similarly accurate ($R^2 = 0.963$), ensuring correct identification of limiting mooring conditions. Fender force prediction, although more nonlinear, remains conservative near the threshold, avoiding unsafe underestimation.

For classification, predicted quantities are compared to their MARIN-defined limits. The confusion matrix in Figure 8.7 shows an overall accuracy of 98.22%, with a very low false-safe rate (1.81%), meaning almost all non-operable cases are correctly flagged. The false-unsafe rate (1.73%) is also small, indicating that unnecessary downtime predictions remain limited.

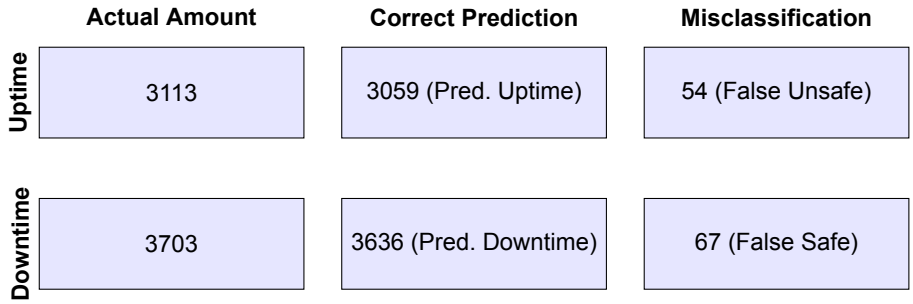


Figure 8.7: Confusion-matrix summary of SSBL operability classification for all 6816 cases (Accuracy: 98.22%; False Safe: 1.81%; False Unsafe: 1.73%).

In summary, the SSBL model provides accurate predictions for all key operability quantities and delivers

a reliable classification with a very low false-safe rate. This demonstrates that SSBL is a suitable and efficient tool for operability screening.

Per-criterion breakdown: Table 8.3 displays classification performance by criterion. Line tension is violated most frequently (54.0 %), yet carries a lower per-criterion false-safe rate (1.04 %) than manifold surge (2.90 %) or fender force (2.54 %). Roll has the lowest per-criterion false-safe rate (0.89 %), reflecting its near-linear dependence on $H_{s,sea}$ and θ_{sea}. None of the 61 roll per-criterion false safes escalate to system level, meaning another criterion was correctly flagged in each case. All 67 system-level false safes involve line tension as the sole active criterion.

Table 8.3: Per-criterion violation and false-safe rates on the 6,816 test cases. System-level false-safe rate is 1.81 % (67 cases), all involving line tension only.

Criterion	Violations	Violation rate	Per-criterion FS rate
Line tension (≤ 672 kN)	3,678	53.96 %	1.04 %
Manifold surge (± 2.0 m)	2,531	37.13 %	2.90 %
Manifold sway (9.26–13.26 m)	2,166	31.78 %	2.00 %
Fender force ($\leq 2{,}997$ kN)	2,010	29.49 %	2.54 %
Roll ($\pm 3.0°$)	1,699	24.93 %	0.89 %

8.2.1. False-Safe Environmental Combinations

All 67 system-level false-safe cases involve a missed line-tension violation. Two physically distinct regimes are present, summarised in Table 8.4.

Table 8.4: False-safe regimes identified in the SSBL test set. All cases have $V_{cur} = 0.24$ m/s.

Regime	$H_{s,sea}$ (m)	$T_{p,sea}$ (s)	Cases	Mechanism
Wind- and current-driven	1.5–2.5	6–11	45	Wind and current loading amplifies LF surge drift
Long-period	0.75–1.0	10–11	20	Long-period wave groups enhance difference-frequency drift forcing

Wind- and current-driven regime: The majority of false-safe cases (45 of 67) involve $H_{s,sea} = 1.5$–2.5 m with $T_{p,sea} = 8$–11 s and $V_{wind} = 4$ or 10 m/s. These cases concentrate in near-head seas ($\theta_{sea} = 160°$–$200°$, 29 of 45 cases), where the fixed heading of the spread-moored FLNG means wave-drift forces act directly along the vessel's longitudinal axis, maximising surge excursions against the SBS mooring system. Combined wind, wave, and current loading drives low-frequency drift forces that push mooring-line tensions above the 50 % MBL limit, even at moderate wave heights [75, 77, 51].

Long-period regime: The remaining 20 cases involve $H_{s,sea} \leq 1.0$ m with $T_{p,sea} = 10$–11 s, predominantly at near-beam headings ($\theta_{sea} = 90°$–$120°$ and $255°$–$270°$, 16 of 20 cases). Because the FLNG cannot weathervane, beam seas persistently excite transverse sway drift against the mooring system regardless of the environmental conditions, and the long-period JONSWAP spectrum produces wave groups whose difference-frequency content $\omega^- = |\omega_m - \omega_n|$ more efficiently excites the low-frequency

mooring modes [51]. The SSBL model encodes $H_{s,sea}$ and $T_{p,sea}$ as scalar inputs and has no access to the spectral shape driving this effect.

8.2.2. Computational Burden of the SSBL Operability Assessment

The TSBL operability models were trained and evaluated on a workstation equipped with an Intel Xeon (Model 85) Skylake-SP processor with 8 physical cores and 16 logical threads at a base frequency of 2.59 GHz. Training is computationally negligible: both TSBL models completed training in a combined 3.86 s. Inference is similarly efficient. Predicting operability for the remaining 6,816 sea states required only 0.03 s in total, corresponding to an inference throughput of approximately 236,000 predictions per second. Scaling this result to the full set of 8,317 sea states yields an equivalent inference cost of approximately 0.04 s.

Table 8.5: Essential CPU specifications used for SSBL operability model training and inference

CPU Model	Architecture	Cores / Threads	Base Frequency	Note
Intel Xeon Model 85	Skylake-SP	8 / 16	2.59 GHz	Workstation CPU

Summary of Computational Savings: The TSBL operability surrogate requires only 3.86 s of CPU time for training and approximately 0.04 s of inference to evaluate all 8,317 sea states. Expressed in the same units as the numerical baseline, this corresponds to fewer than 4 CPU-seconds of compute, compared with more than 1,400 core-hours for the original time-domain simulations. Although the TSBL approach still depends on an initial difficult-case extraction from the numerical dataset, its computational burden is effectively negligible, enabling full operability evaluations on a single workstation in real time and without the need for multi-node HPC resources.

8.3. TSBL Operability Assessment

The TSBL model evaluates operability using full time-series inputs, enabling it to capture time-dependent dynamics that are not available to the SSBL model. For the uptime/downtime task, predictions are compared directly against MARIN classifications.

The TSBL confusion matrix in Figure 8.8 shows that the model is able to identify most downtime events accurately. While the false-negative count (missed downtime) is higher than in the SSBL case, the precision for predicted downtime remains high. TSBL could therefore be well-suited for detailed dynamic assessments, complementing the faster SSBL approach.

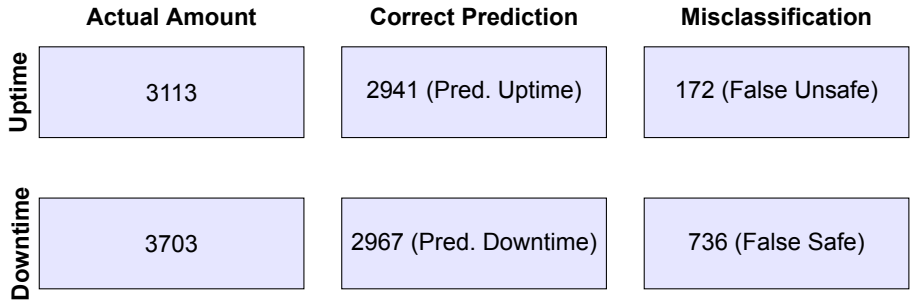


Figure 8.8: Confusion-matrix summary of TSBL uptime/downtime classification for 6,816 simulations (Accuracy: 86.7%; False Safe: 19.9%; False Unsafe: 5.5%).

The TSBL confusion matrix in Figure 8.8 shows that the model correctly classifies 5908 out of 6816 evaluated simulations, corresponding to an overall classification accuracy of 86.7%. Of the 3703 downtime cases in the dataset, 2967 are correctly identified, while 736 are classified as operable, resulting in a false-safe rate of 19.9% relative to all downtime cases. For uptime conditions, 2941 out of 3113 cases are classified correctly, with 172 cases incorrectly flagged as downtime, corresponding to a false-unsafe rate of 5.5%.

While the false-negative rate (missed downtime) is higher than for the SSBL approach, the precision of predicted downtime remains high at 94.5%. These results indicate that TSBL is capable of identifying the majority of downtime events, while incurring a moderate number of false-safe classifications. TSBL is therefore best suited for detailed dynamic assessment of selected conditions, complementing the computationally efficient SSBL approach used for large-scale operability screening.

8.3.1. False-Safe Environmental Combinations

A total of 736 out of 6,816 cases (10.8 %) result in false-safe predictions, all involving at least one line-tension violation. The dominant environmental combinations are listed in Table 8.6, with $V_{cur} = 0.24$ m/s appearing in all top combinations.

Table 8.6: Most frequent V_{wind}–$H_{s,sea}$–$T_{p,sea}$ combinations in the 736 TSBL false-safe cases. All entries have $V_{cur} = 0.24$ m/s.

V_{wind} (m/s)	$H_{s,sea}$ (m)	$T_{p,sea}$ (s)	Count
10	3.0	7	46
4	3.0	7	44
10	3.0	9	39
4	3.0	9	35
4	1.5	8	26
10	2.0	7	24

Dominant pattern and model limitations: $H_{s,sea} = 3.0$ m with $T_{p,sea} = 7$–10 s and maximum wind and current loading is the primary driver, though $H_{s,sea} = 1.5$ m and 2.0 m cases are also present, indicating the model underestimates mooring loads across a range of severe to moderately severe sea states. For the dominant $H_{s,sea} = 3.0$ m subset, false-safe cases concentrate in near-head and quartering-head seas ($\theta_{sea} = 140°$–$200°$), with the most critical directional combination being $\theta_{wind} = \theta_{cur} = 180°$ (108 cases), where wind, current, and waves act collinearly from ahead, maximising the steady longitudinal drift load [51]. At these conditions, second-order drift forces scale with the square of wave amplitude and the nonlinear mooring response becomes strongly amplitude-dependent. The 1,500-case training set provides limited coverage of the most extreme response combinations, and the LF model cannot reliably extrapolate beyond its training distribution. More fundamentally, the LF model must infer the steady wind and current drift contribution indirectly from scalar environmental inputs, without access to the force time histories that drive the mean offset. This leads to systematic underestimation of the steady bias component, reducing the predicted excursion margin and causing the model to misclassify borderline high-severity cases as operable. This differs from the SSBL failure mode: the TSBL misses the most severe conditions, whereas the SSBL misses conditions that appear mild but carry hidden hydrodynamic hazard.

Line geometry: Lines 3 and 4 (front breast lines at the LNGC bow) are exceeded in 81.5 % and 75.8 % of false-safe cases, lines 15 and 16 (quarter breast lines at the stern) follow at 60.2 % and 57.7 %, and lines 7 and 8 (after bow springs) produce zero violations. This is consistent with surge-driven loading in near-head seas, where longitudinal drift loads the transverse breast lines most heavily while the longitudinally oriented springs remain lightly loaded. Together with the directional clustering and collinear wind-current alignment, this confirms that the TSBL model underestimates specifically the peak surge-driven breast line response under combined high-severity head-sea forcing.

8.3.2. Training and Inference Efficiency of the LF and WF Models

The LF and WF models were trained on 1,500 simulations and evaluated on 6,816 test simulations. Both architectures converged within approximately 9 hours of training on a single NVIDIA Quadro P4000 GPU, corresponding to a per-simulation training cost of roughly 3.6 s. Inference of the 6,816 test simulations required 1.3–1.5 h in total, equivalent to 0.7–0.8 s per simulation.

Table 8.7: Essential GPU specifications used for LF and WF model training and inference

GPU Model	CUDA Cores	Memory Bandwidth	Available Memory	Graphics Clock
Quadro P4000	1,792	243 GB/s	204 GB	1,202 MHz

GPU Hardware Overview: Training and Inference Times:

- **LF model:** 9.94 h training; 1.51 h inference for 6,816 simulations.
- **WF model:** 8.97 h training; 1.34 h inference for 6,816 simulations.
- **Combined:** 18.91 h total training and 2.85 h total inference.

Workflow Cost: The surrogate workflow requires an initial training cost of 18.91 GPU-hours and an inference cost of approximately 3.5 GPU-hours to evaluate all 8,317 sea states. Expressed in computational work, this corresponds to a total of:

$$21.8 \text{ GPU-hours.}$$

For subsequent operability studies, only the inference stage is required, reducing the repeated computational cost from 1,428 core-hours to just 2.85 GPU-hours. This shift from repeated core-hour expenditure to a one-time training effort enables rapid and resource-efficient environmental sweeps on a single workstation and eliminates the need for multi-node HPC resources during repeated operability evaluations.

9

Discussion

This chapter provides a critical interpretation of the results and modelling choices described in this thesis. The aim is not to restate findings, but to analyse the methodological decisions, the structure of the dataset, and the relevant hydrodynamic mechanisms that shaped the performance of the ML models. By following the order of the preceding chapters, this discussion evaluates (i) the nature of the data used, (ii) the behaviour of the SSBL and TSBL models, and (iii) the hydrodynamic explanations for the observed model limitations. The chapter concludes with a broader reflection on what the combined findings suggest for the application of ML models to multibody offshore operability assessment more generally.

9.1. Input limitations and implications

In the present dataset, the incoming wave elevation signal $\zeta(t)$ for each sea state was generated using linear wave theory based on a JONSWAP spectrum. For every $(H_{s,\mathrm{sea}}, T_{p,\mathrm{sea}})$ combination a *single* random seed was used. This resulted in all wave realisations sharing the same temporal pattern whenever they had the same peak period. The only difference between these signals was the amplitude scaling associated with changes in H_s. The use of a single-point elevation time series rather than a directional wave field amplified this limitation, because no directional spreading or spatial coherence of the wave field was represented.

These characteristics influenced the behaviour and generalisation capability of the ML models. The models were exposed to only one specific wave-phase realisation per sea state. As a consequence, they learned to predict responses accurately only for sea states constructed in the same manner as the MARIN simulations. The resulting models perform reliably for new H_s values at fixed T_p because the temporal pattern of the incoming wave remains unchanged. They are not able to anticipate different phase realisations, varied wave-group structures or alternative random-phase combinations. Furthermore, due to the absence of directional spreading, the models cannot account for short-crestedness or changes in encounter geometry. These assumptions also imply that training on measured field data would require additional preprocessing, since real wave records typically contain noise, nonstationarity, and uncertainty in directional content compared to the idealised simulation inputs used here.

The reliance on linear wave theory with a JONSWAP spectrum introduced additional limitations. The input signal contained no nonlinear bound harmonics and no possibility for extreme or freak wave

events. As a result, the models were not exposed to the full range of phenomena that can occur in real seas. When applied outside the training conditions, the models cannot be expected to respond correctly to unusual or extreme wave conditions.

The impact of these constraints can be summarised as follows:

- the models generalise only to sea states sharing the same wave-generation procedure;
- no variability in wave-group structure was present which prevented the TSBL models from learning relationships between group modulation and motion response;
- the single-point elevation input contained no directional spreading which removed spatial effects such as phase differences along the hulls;
- nonlinear wave amplification and extreme wave events could not occur which limits the applicability of the models to operational conditions.

These limitations are especially relevant when considering the potential use of the segment-based learning method for near-future motion forecasting in real operations. Operational wave fields are short-crested, directional and often exhibit nonlinear behaviour. Therefore, a model trained solely on single-seed, single-point elevation signals generated using linear JONSWAP waves cannot be expected to anticipate the variability present in real sea states.

If the models developed in this thesis are to be applied in operational settings, the training dataset would need substantial expansion. Multiple random seeds should be introduced to capture realistic wave-group variability. Directional wave fields should replace point-elevation signals to provide spatial coherence and directional richness. Finally, validation using actual wave and motion measurements would be essential because real sea states may contain nonlinear features or extreme events that lie outside the assumptions of linear wave theory. Only with such extensions can model accuracy, generalisation limits, and failure modes be attributed to physical system behaviour rather than artefacts of idealised input data. Without these extensions, the present models cannot reliably interpret or forecast vessel motions in conditions involving larger wave variability, nonlinear wave amplification, or extreme wave events.

9.2. Evaluation of the Modelling Approach

This section evaluates the modelling approaches developed in this thesis by relating their performance and limitations to the underlying physical processes and data representations.

9.2.1. Summary-Statistics-Based Learning

The SSBL models demonstrate strong predictive skill for most operability relevant quantities. This performance follows from the behaviour of the FLNG and LNGC system in the MARIN dataset. Both the low frequency and the wave frequency maxima show clear and consistent patterns across the environmental matrix. Their dominant variation aligns strongly with broad spectral properties of the sea state, particularly the overall energy level and the dominant wave period. Within this structured dataset these relationships are sufficiently regular that the maximum responses can be captured reliably from their summary descriptors. This helps explain why the SSBL approach is able to reproduce the response envelopes with high accuracy using only summary statistics and why it is effective for rapid environmental screening.

At the same time, the false-safe cases in the results reveal important limitations of using summary descriptors alone. These misclassifications arise in conditions where the limiting responses are influ-

enced by physical mechanisms that are not encoded in the environmental input vector. In the present dataset, the false-safe cases do not cluster in a single region of wave height or peak period. Instead, they occur primarily in environmental headings around 90° and 270°, where the MARIN simulations show that gap hydrodynamics, multibody interaction effects and coupled sway–yaw modes are most active. Several false-safe classifications correspond to cases in which the reference simulations exhibit amplified low-frequency surge or sway offsets or elevated line tensions due to these interaction-driven mechanisms, even though the associated wave statistics alone do not indicate critical behaviour.

9.2.2. Time-Series-Based Learning: Segment-Based Prediction

The segment-based learning method predicts a future segment from a fixed-length input window of past motions and wave elevation. Its behaviour is best understood by comparing the chosen modelling parameters in chapter 6 with the physical autocorrelation characteristics of the responses.

Wave-frequency component. The WF model used an input window of $350\ \mathrm{s}$, while the autocorrelation memory of the WF motions was only 8–$15\ \mathrm{s}$. This means that more than 90% of the input window contains information that does not contribute to WF predictability. The long window remained computationally efficient because it avoided creating a larger number of short overlapping segments during training, and the LSTM encoder was able to extract the short-memory structure directly from the full sequence.

Low-frequency component. The LF model used a $900\ \mathrm{s}$ input window, consistent with LF memory times that exceeded several hundred seconds. Because slow drift evolves gradually, the LF state was fully encoded in the input window, enabling accurate LF forecasts. This explains why the segment-based LF predictions were significantly better than the LF performance in the full reconstruction model.

The key technical observations are:

- WF memory (8–$15\ \mathrm{s}$) is much shorter than the window ($350\ \mathrm{s}$); the encoder filters redundant information, but the design is not memory-optimal.
- LF memory ($> 500\ \mathrm{s}$) matches the window ($900\ \mathrm{s}$), allowing the LF state to be inferred directly from the history.
- Segment-based LF forecasting works well because the LF state is observed; this is not the case in the full reconstruction model, where the LF state must be inferred from wave elevation alone.

9.2.3. Time-Series-Based Learning: Full Reconstruction

In the full reconstruction task, the model was required to generate the complete three-hour response from the point wave elevation signal and the environmental parameters. Unlike the segment-based model, no motion history was provided. As a consequence, the model had no direct information about the existing LF state at the start of the prediction. The slow-drift offset and ongoing LF trends had to be inferred solely from the wave elevation and static descriptors. This setup fundamentally limits LF prediction accuracy because the physical drivers of slow drift are not present in the inputs.

This limitation is reinforced by differences in the amount of learnable structure in the data. Over a three-hour simulation, the WF component contains approximately 800–$1\,000$ cycles (periods range from 8 - 15 seconds), while the LF component contains roughly 70–100 cycles for drift periods in the range of 100–$150\,\mathrm{s}$. The WF model therefore benefits from substantially more repeated patterns per simulation. WF motions also depend strongly on the instantaneous wave elevation, which is available to the model at every timestep. By contrast, LF drift evolves slowly, contains fewer distinct variations within each

record and depends on hydrodynamic forcing that was not included in the feature set. This combination makes LF reconstruction significantly more challenging than WF reconstruction.

The absence of LF information also impacts the mooring-line tension predictions. Mooring-line tension is dominated by the LF offset and its slow variations. Errors in LF reconstruction therefore appear directly in the tension predictions, meaning the tension model inherits the same limitations as the LF model.

A further constraint is the lack of explicit LF forcing in the inputs. According to the hydrodynamic formulation summarised in section 4.7 and Appendix B, LF drift loads are driven by second-order difference-frequency wave interactions. These forcing terms were not included in the input for the models. The environmental inputs also lacked directional wind and current loads, which contribute to LF drift in operational conditions. Without these physical drivers, the model cannot reproduce LF behaviour that originates from processes not observable in the wave elevation signal.

Overall, the full reconstruction results demonstrate that LF prediction cannot be performed reliably from the wave elevation and static descriptors alone, with limited training coverage. The model can learn WF behaviour effectively due to the large number of WF cycles and their strong dependence on the wave input, but LF reconstruction remains underdetermined unless explicit LF excitation and directional environmental forcing are provided.

9.2.4. Low-Frequency Model Training and Optimisation

The learning curve in Figure 6.6 reveals two features of the LF model training. First, the validation loss stabilises around epoch 40–50 while the training loss continues to decrease steadily, producing a persistent training-validation gap. Second, the best validation loss is reached at epoch 110, after which overfitting sets in. Together, these observations show that the model fits the training simulations increasingly well but stops generalising to unseen sea states early in training. The underlying cause is the low information density of LF responses: each three-hour simulation contains only approximately 70–100 low-frequency cycles, compared to roughly 1000 wave-frequency cycles for the WF model, providing far fewer distinct pattern repetitions from which the Transformer can extract generalisable slow-drift dynamics. This is consistent with the segment-based learning curve in Figure 7.19, which showed that stable LF performance required approximately 60% of the available simulations against 40% for the WF model. Notably, the head-wave operability assessment in Figure 7.2.2 demonstrated that the LF model trained on only 300 of 756 head-sea simulations performed well, suggesting that the data requirement is driven by the complexity of the environmental space rather than by the absolute number of simulations. For the full multi-directional dataset, where slow-drift forcing varies strongly with wave direction and peak period, the 1,500 simulation training set is insufficient to fully constrain the LF Transformer.

The hyperparameter optimisation (Appendix D.2.3) reinforces this conclusion. The Optuna TPE search, conducted on only 20% of the training data, produced a best validation loss of 5.124 against a manual baseline of 3.994, so the baseline configuration was retained. Bayesian optimisation requires a stable validation signal to reliably rank candidate configurations, a condition that cannot be met when the training set covers only a fraction of the relevant LF response regimes. Meaningful LF hyperparameter search is therefore identified as a priority for future work, requiring a larger share of the full dataset and a greater number of trials.

9.2.5. Incorporating Extracted Low-Frequency Drift Forcing as Model Input

As outlined in Appendix B and in the classical formulation by Pinkster [111], the low-frequency (LF) drift force originates from the quadratic interaction of the wave-induced horizontal fluid velocity. For deep-water conditions this velocity can be approximated directly from the surface elevation by

$$u(t) \approx \omega_{\mathrm{peak}}\, \zeta(t), \qquad u^2(t) \propto \zeta^2(t), \tag{9.1}$$

where ω_{peak} is the dominant angular frequency of the sea state. This relation shows that the LF drift excitation is not contained in $\eta(t)$ itself, but in the low-frequency content of $\zeta^2(t)$. To understand why $\zeta^2(t)$ contains the physically relevant LF excitation, consider the elevation expanded as a sum of spectral components,

$$\zeta(t) = \sum_n A_n \cos(\omega_n t + \phi_n), \tag{9.2}$$

where A_n are the wave component amplitudes, ω_n the angular frequencies, and ϕ_n the corresponding random phases. Expanding $\eta^2(t)$ yields three characteristic terms:

$$\text{Mean (DC) contribution:} \qquad \frac{1}{2}\sum_n A_n^2, \tag{9.3}$$

$$\text{Difference-frequency (slow drift):} \qquad \sum_{m\neq n} A_m A_n \cos[(\omega_m - \omega_n)t + (\phi_m - \phi_n)], \tag{9.4}$$

$$\text{Sum-frequency (HF contribution):} \qquad \sum_{m,n} A_m A_n \cos[(\omega_m + \omega_n)t + (\phi_m + \phi_n)]. \tag{9.5}$$

Pinkster's insight is that only the first two components ontribute to the slowly varying drift force: the DC component determines the mean-drift level, while the difference-frequency components encode the wave-group modulation that drives slow drift. These contributions are entirely absent from $\zeta(t)$ itself but present in $\zeta^2(t)$.

From this decomposition, the associated drift forcing can be expressed compactly by applying the physical scaling used in Appendix B. Approximating the horizontal fluid velocity through $u(t) \approx \omega_{\mathrm{peak}}\zeta(t)$, the quadratic interaction yields

$$F_{\mathrm{drift}}(t) = \frac{1}{2}\rho g\, u^2(t) \;\approx\; \frac{1}{2}\rho g\, \omega_{\mathrm{peak}}^2\, \zeta^2(t), \tag{9.6}$$

which contains the constant mean drift and the slowly varying difference-frequency contribution identified by Pinkster. To isolate only the low-frequency band that drives the vessel's slow-drift motion, a fourth-order Butterworth low-pass filter is applied to the quadratic term:

$$F_{\mathrm{LF}}(t) = \mathrm{BW}_4[F_{\mathrm{drift}}(t),\, f_c = 0.03\ \mathrm{Hz}]\,. \tag{9.7}$$

The resulting signal $F_{\mathrm{LF}}(t)$ represents the physically relevant slow-drift excitation that governs the LF vessel response. Since this forcing is not contained in $\eta(t)$ itself but only emerges after the quadratic interaction and filtering, it is absent from the current TSBL model inputs. Including $F_{\mathrm{LF}}(t)$ as an additional input channel would therefore provide the missing LF excitation mechanism required for accurate

slow-drift and LF-dominated tension prediction. Figure 9.1 shows this decomposition, where peaks in the wave-group envelope align with peaks in the LF drift force.

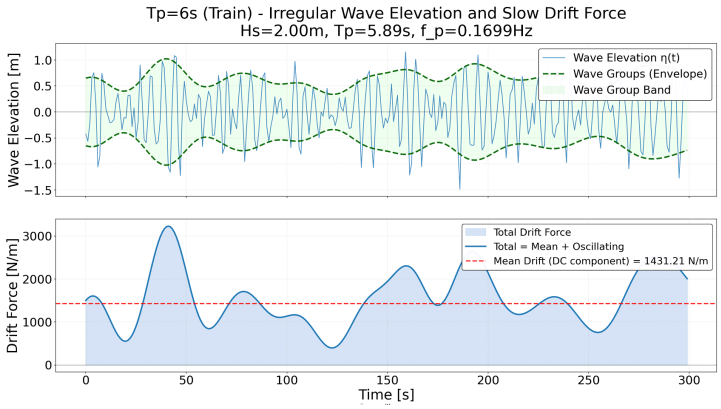


Figure 9.1: LF drift forcing extracted from the wave elevation. The upper panel shows $\eta(t)$ and its group envelope. The lower panel shows the corresponding LF drift force consisting of a mean component and a slowly varying oscillatory contribution.

This formulation is directly relevant for the ML models used in this thesis. In the full reconstruction setup, the model receives only $\zeta(t)$ and environmental parameters. However, $\zeta(t)$ contains none of the difference-frequency terms that physically generate slow drift. The model is therefore forced to extract LF excitation from an input signal that does not directly contain the relevant information. This explains the systematic LF underprediction observed in chapter 7 and the associated errors in tension prediction, which is dominated by the LF offset.

Providing the extracted LF forcing signal $F_{\mathrm{LF}}(t)$ as an additional input channel would supply the model with the correct physical excitation mechanism driving slow drift. This would close a key information gap in the present feature set and is expected to significantly improve LF reconstruction and LF-sensitive quantities such as mooring-line tension.

9.2.6. Wind and Current Forcing

The environmental pattern analysis in subsection 8.2.1 (SSBL) and subsection 8.3.1 (TSBL) shows that many false-safe classifications occur for conditions where wind, waves, and current are closely aligned. In these cases the vessels experience significant steady horizontal loads that act on the same low-frequency time scale as the second-order wave drift forces, yet these contributions are not provided to the TSBL models.

According to MARIN's load formulations (Appendix B.2.6–B.2.7), the steady wind and current forces acting on a ship-shaped hull are

$$F_{\mathrm{wind}} = \tfrac{1}{2}\rho_{\mathrm{air}}\, V_{\mathrm{wr}}^2\, C_w(\theta_{\mathrm{wr}})\, A, \qquad F_{\mathrm{current}} = \tfrac{1}{2}\rho_{\mathrm{water}}\, V_{\mathrm{cr}}^2\, C_c(\theta_{\mathrm{cr}})\, A, \tag{9.8}$$

where V_{wr} and V_{cr} are the relative wind and current velocities, θ_{wr} and θ_{cr} their ship-fixed directions, and C_w, C_c the associated dimensionless coefficients. These loads provide a persistent LF excitation in surge, sway, and yaw.

When wind and current are aligned with the wave direction, these steady loads reinforce the slowly varying drift excitation extracted from $\zeta^2(t)$. Because the present TSBL models only receive the instantaneous elevation $\zeta(t)$ and static descriptors, they do not have access to these steady LF forcing components. As a result, the LF model is required to infer the combined slow-drift environment without

receiving all the physical drivers that shape the LF response. Including the physically meaningful LF inputs

$$F_{\mathrm{wind}}(t), \quad F_{\mathrm{current}}(t), \quad F_{\mathrm{LF}}(t),$$

would increase the amount of dynamically relevant LF excitation presented to the model. This enriches the LF input space with the steady and slowly varying horizontal loads that are fundamental to surge–sway–yaw drift in spread-moored systems, enabling the TSBL architecture to form a more complete and physically consistent representation of the forces governing slow-drift behaviour, particularly in the environmentally aligned conditions where most false-safe classifications occur.

9.2.7. Practical Meaning of Model Accuracy

The reported accuracy of the ML models must be interpreted in terms of their effect on operability decisions. Operability is determined by strict physical limits for roll, relative motions and mooring and fender loads. Near these limits even a small prediction error can change an operable sea state into an inoperable one and this may lead to unnecessary downtime or to unsafe calls for offloading. This sensitivity explains why a numerical accuracy value cannot be evaluated without considering its practical consequences.

The Summary Statistics Based Learning model achieves an accuracy of approximately 98% and the associated false safe rate remains very small. This indicates that the model is sufficiently conservative near the operational limits. For this reason it is well suited as a fast screening tool for large environmental datasets. Its predictions remain reliable as long as the sea state is not close to any threshold.

The Time Series Based Learning model achieves an accuracy of approximately 86%. This accuracy is not uniform across all physical regimes. The model predicts wave frequency motions well but shows reduced reliability in low frequency drift dominated conditions. These conditions govern the most critical operability limits. A small error in this regime can lead to incorrect classification. For that reason the Time Series Based Learning model cannot replace high fidelity time domain simulation in decision critical workflows.

A model may be considered operationally acceptable when its errors near the physical limits remain conservative and when false safe predictions remain very rare. Under these criteria the Summary Statistics Based Learning model is adequate for screening while the Time Series Based Learning model requires further development before it can be used for stand alone operability assessment.

9.3. Minimal Dataset

A central question of this thesis was whether a reduced dataset can still support accurate prediction modelling. The empirical findings in Chapters 7 and 8 show that the notion of a minimal dataset is not universal. It depends on the modelling approach and on the type of information the model needs to represent. The summary statistics based learning and the time series based learning respond in different ways to dataset reduction because they extract different structures from the data.

Summary-statistics-based learning: SSBL models learn a relatively smooth relationship between environmental descriptors and operability-limiting response maxima. Training on approximately 1 500 simulations yields almost the same predictive accuracy as training on the full dataset, reflecting the dense and well-structured nature of the MARIN environmental matrix. Once the dominant dimensions of significant wave height, peak period and wave direction are sufficiently covered, additional simulations contribute little new information. The observed saturation at approximately 20% of the dataset is specific to this uniformly sampled matrix and should not be taken as a general rule for ML-based

operability assessment. It does suggest, however, that covering approximately 20% of the relevant environmental ranges, well distributed across the key dimensions, is sufficient for reliable SSBL performance.

Time-series-based learning: TSBL models require exposure to the full range of temporal patterns arising from combinations of wave period, wave direction and multibody hydrodynamics. The three-hour simulation duration plays an important role here: a single simulation contains hundreds of WF cycles but only 70–100 LF cycles, meaning that each simulation provides far more learnable structure for WF prediction than for LF prediction. This is reflected directly in the training data requirements found in subsection 7.2.1: the WF model reaches stable performance using approximately 40% of the three-hour record, while the LF model continues to improve up to approximately 60%. LF drift also depends on multibody coupling effects that vary strongly with wave direction and peak period, requiring broader environmental coverage to represent well. The 1 500-simulation training set used in this thesis is therefore not sufficient for reliable TSBL, particularly for LF prediction across the full directional matrix.

Relevance for project preparation: Industry experts from Shell pointed out that in early project stages, simulation budgets are limited. When discrete levels of parameters such as current speed or direction must be included, a considerable portion of the available budget is spent before wave conditions are varied, leaving only a small selection of wave height, period and direction combinations to be examined. The findings of this thesis suggest that simulation campaigns should not be designed based on a fixed percentage of the environmental matrix, but on targeted coverage of the conditions that matter most for prediction accuracy. ML models can themselves help identify which conditions are poorly covered and would add the most value if simulated, progressively focusing effort where it is needed most.

9.4. Alternative Approach: Linearised Frequency Domain Modelling

The methodology developed in this thesis relies on ML models trained on MARIN's three hour time domain simulations. An implicit assumption in this approach is that time domain simulations are the only viable source of high fidelity reference data and that any reduction in computational burden must therefore originate from reducing the number of required simulations or replacing them with ML models. An alternative strategy used in offshore engineering is to replace selected time domain simulations with a linearised frequency domain model, especially for sea states where the system behaviour remains predominantly linear.

This raises the methodological question of whether a hybrid workflow that applies frequency domain analysis to the more linear conditions and time domain simulation only to the nonlinear cases could reduce computational effort more effectively than the ML-based strategy explored in this thesis.

Wave Frequency Reconstruction as a Starting Point

A first indication of the potential of such an approach is provided by the RAO based reconstruction of the wave frequency roll motions. The wave frequency component of each vessel's motion is reconstructed through the linear superposition

$$\eta_{\mathrm{WF}}(t) = \sum_{i=1}^{N} \mathrm{RAO}(\omega_i, \mu_i)\, a_i\, \mathrm{e}^{\mathrm{i}\omega_i t}, \tag{9.9}$$

where a_i are the complex JONSWAP amplitudes, ω_i the component frequencies and μ_i the correspond-

ing headings.

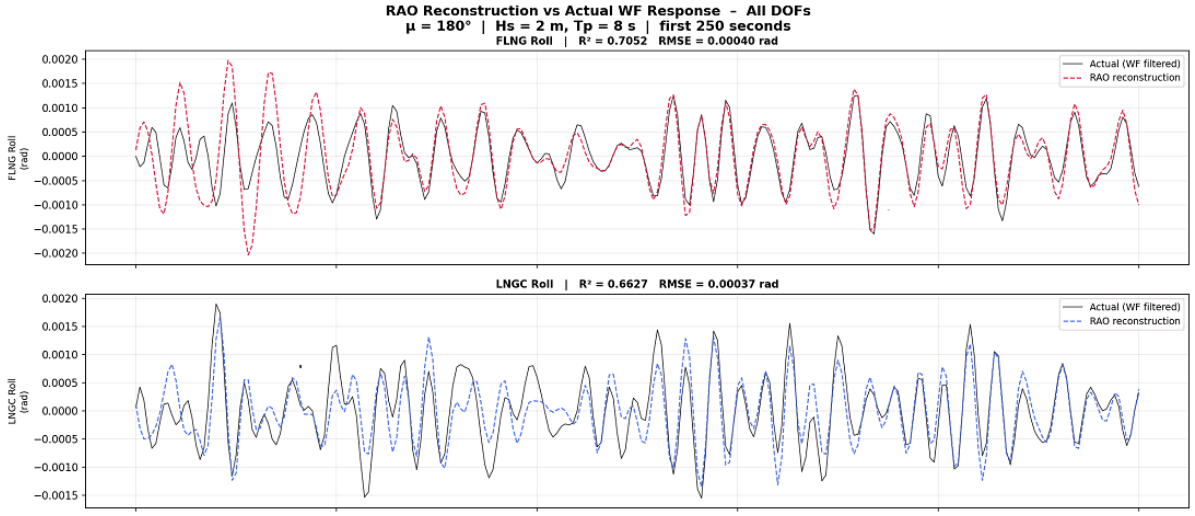


Figure 9.2: Comparison between RAO based reconstruction and the wave frequency filtered MARIN simulation for roll motion of the FLNG (top) and LNGC (bottom). The dashed coloured lines show the linear reconstruction and the solid black lines show the MARIN wave frequency response.

The reconstruction shows strong phase alignment and accurate reproduction of the dominant oscillatory structure, with coefficient of determination (R^2) values on the order of 0.65 to 0.70. This indicates that for this sea state, with head sea, the wave frequency roll response behaves largely linearly. In these regimes a frequency domain computation could replace a full three hour nonlinear simulation at much lower computational cost.

Incorporating Mechanical Coupling: Linearisation of Moorings and Fenders

A complete frequency domain representation of the side by side system must also incorporate the mechanical coupling between the vessels. The SBS mooring lines, the pneumatic fenders and the viscous damping introduce nonlinear behaviour. These effects can be represented in an equivalent linear form using statistical linearisation, as commonly introduced in the context of random vibration theory and discussed in Metrikine's lecture notes on stochastic dynamics [112]. For a nonlinear restoring law $f(x)$ the equivalent linear stiffness is given by

$$k_{\mathrm{eq}} = \frac{\mathbb{E}[x f(x)]}{\mathbb{E}[x^2]}. \tag{9.10}$$

Here, $x(t)$ denotes the scalar deformation variable associated with a nonlinear mechanical component, such as fender compression or mooring-line elongation in the SBS system, treated as a random process, which allows a linearised representation of

- fender compression,
- mooring line restoring behaviour around the pretensioned position,
- viscous contributions modelled as equivalent linear damping.

Using this approach, an equivalent linear stiffness matrix K_{eq} can be constructed. Combined with the hydrodynamic coefficients from DIFFRAC, this leads to the frequency domain equations for the coupled twelve degree of freedom system,

$$\big(-\omega^2 M_\infty + \mathrm{i}\omega\, C_\mathrm{R}(\omega) + K_\mathrm{eq} \big)\, \widehat{\eta}(\omega) = \widehat{F}_\mathrm{exc}(\omega), \tag{9.11}$$

from which the response spectrum follows as

$$S_\eta(\omega) = |H(\omega)|^2\, S_{\eta\eta}(\omega). \tag{9.12}$$

Extending the Approach to Low Frequency Response Using QTFs

While the wave frequency component can be reconstructed directly from first order RAOs, the low frequency response is driven by second order drift forces. These forces can be represented in the frequency domain using quadratic transfer functions. For each pair of wave components the difference frequency contribution is

$$F_k^\mathrm{LF}(\omega^-) = \sum_{m=1}^{M} \sum_{n=1}^{M} \mathrm{QTF}_k(\omega_m, \omega_n, \mu_m, \mu_n)\, a_m\, a_n,$$

which provides the slowly varying surge, sway and yaw forcing required to reconstruct the low frequency response. Once a drift forcing spectrum has been obtained, the same equivalent stiffness and damping representation can be used to generate an approximate low frequency response within the frequency domain formulation. This offers a potential route to replacing selected low frequency dominated time domain simulations when the drift behaviour remains within the region where linearisation is appropriate.

The combination of RAO based reconstruction for the wave frequency component, statistical linearisation for the mechanical coupling and quadratic transfer functions for the low frequency drift forcing shows that an extended frequency domain model may serve as a computationally efficient alternative for selected sea states. A complete replacement of time domain analysis is not feasible for all conditions, but a hybrid strategy that uses frequency domain modelling for the more linear conditions and time domain simulation for strongly nonlinear regimes may offer a useful balance between accuracy and computational effort.

9.5. Integrated Interpretation

The preceding sections have evaluated the performance and limitations of the SSBL and TSBL approaches in detail. This section steps back to reflect on what the combined picture suggests, both for the FLNG–LNGC case studied here and for the broader class of coupled floating-body operations that motivated this work.

A recurring theme across the results is that ML models can reduce the number of time-domain simulations required for operability assessment, but that the extent of this reduction depends on what is asked of the models. SSBL suggests that much of the information needed for go/no-go decisions is already contained in environmental summary statistics, while TSBL raises the question of whether the temporal detail that conventional workflows discard can be exploited to provide richer insight near operability boundaries. The results suggest it can, but only when training data adequately represent the relevant spectral and directional regimes. This points to a practical question this thesis does not fully resolve: how should simulation campaigns be designed when ML models are part of the intended workflow? A natural direction is to use the ML models themselves to identify which conditions are poorly represented and would add the most value if simulated. Such uncertainty-guided selection of new simulations could

make the workflow self-improving, progressively concentrating simulation effort on the conditions that matter most for operability prediction.

What the results also make clear is that a dependence on high-fidelity time-domain simulations is not eliminated. ML model accuracy degrades in the nonlinear regimes that are often most critical for operability, including conditions driven by gap resonance, strong LF drift, and collinear wind-current-wave loading. The role of ML may therefore be better understood as a means of concentrating simulation effort where it is needed, rather than replacing it.

Some findings are likely specific to this configuration. The false-safe regimes, the sensitivity of LNGC roll to gap resonance, and the SSBL saturation behaviour are tied to this spread-moored system and the structured MARIN environmental matrix. In contrast, the distinction between WF and LF predictability, the dependence of TSBL generalisation on response-range coverage, and the need for explicit LF forcing information are grounded in physical mechanisms not unique to this configuration, and are worth investigating beyond the FLNG–LNGC case.

From the helicopter view that motivated this thesis, the results indicate where ML methods add value, where they require careful data design, and where time-domain simulation remains indispensable. How far these findings extend to other coupled floating-body operations is a question the results can inform but not fully answer.

10

Conclusions & Recommendations

This thesis investigated the application of machine learning (ML) models to simulation-driven operability assessment for multibody offshore operations, using a Floating Liquefied Natural Gas and LNG Carrier (FLNG–LNGC) side-by-side offloading configuration as a case study. The work developed and evaluated two complementary modelling approaches: summary-statistics-based learning (SSBL), which predicts operability-limiting response extrema directly from environmental descriptors, and time-series-based learning (TSBL), which reconstructs full response histories from wave elevation and environmental inputs. Both approaches were trained on high-fidelity time-domain simulation data produced by the Maritime Research Institute the Netherlands (MARIN) for Shell. Together, they were used to systematically investigate what ML methods can and cannot offer for this class of problems, from rapid environmental screening to full response reconstruction.

This chapter presents the conclusions of the thesis in Section 10.1, structured around the main research question and its subquestions. Section 10.2 then presents recommendations for future work.

10.1. Conclusion

This section presents the conclusions of the thesis by systematically answering the main research question and its subquestions based on the findings discussed in the preceding chapters.

How can machine learning-based models support and accelerate operability assessment for side-by-side offloading?

ML models can support and accelerate operability assessment for FLNG–LNGC side-by-side offloading by substantially reducing the number of time-domain simulations required for a full environmental sweep. SSBL models trained on environmental summary statistics achieve 97% operability classification accuracy across more than 62,000 simulations when trained on the full sea and swell wave dataset. When trained on a reduced set of 1,500 sea-wave-only simulations, classification accuracy reaches 98.2% for 6,816 unseen sea states, with prediction for all 8,317 sea-wave conditions taking less than a second. TSBL models reconstruct complete three-hour response histories within the trained spectral and directional regime. Prediction for 6,816 test simulations required 2.85 GPU-hours, scaling to approximately 3.5 GPU-hours for all 8,317 sea-wave conditions, after a one-time training investment of 18.9 GPU-hours. This compares to approximately 1,428 core-hours for the original MARIN

workflow, representing a reduction in repeated assessment cost of several orders of magnitude and enabling complete environmental sweeps on a single workstation without relying on multi-node high-performance computing resources.

In both cases, a dependence on high-fidelity time-domain simulations is not eliminated. Both approaches require simulation data for training, and sufficient coverage of the relevant physical regimes in the training set is critical. This is confirmed by the head-sea operability assessment, where TSBL trained on only 300 out of 756 head-sea simulations achieved 99.7% classification accuracy, demonstrating that when training data adequately represents the governing response regimes, ML models can perform reliably even in conditions involving nonlinear mooring behaviour and LF drift. The role of ML is therefore not to replace time-domain simulation but to concentrate it on the conditions that matter most.

However, the ML models developed in this thesis cannot yet identify in advance which conditions require simulation, and cannot predict responses driven by physical mechanisms absent from the training inputs. For SSBL, the main limitations arise in beam and near-beam headings and crossed sea-swell conditions, where gap hydrodynamics and coupled sway–yaw modes produce responses that cannot be inferred from summary statistics alone. For TSBL, model accuracy degrades in severe sea states combined with maximum current velocity and aligned wind and current forcing, where steady drift loads are systematically underestimated because wind and current are not included as explicit model inputs. In these regimes, LNGC roll is especially challenging due to its weak damping and sensitivity to gap hydrodynamics.

To what extent can operability-relevant responses be predicted using only environmental summary statistics?

SSBL models predict most operability-relevant response extrema with high accuracy using only environmental summary statistics: significant wave height, peak wave period and mean wave direction of the sea and swell components, together with wind and current velocity and direction. Drift-dominated responses such as surge, line tension and roll in most headings are reproduced reliably, as their peak values follow consistent and predictable trends across the environmental parameter space. The MLP achieves 97% operability classification accuracy across more than 62,000 simulations, demonstrating that summary statistics are sufficient for fast, reliable screening across the majority of the environmental space. Performance saturates at approximately 20% of the dataset for this uniformly sampled environmental matrix, indicating that covering the relevant spectral and directional regimes matters more than the total number of simulations. While this saturation behaviour is specific to the structured MARIN dataset, it suggests that covering approximately 20% of the relevant environmental regimes, well distributed across wave height, period and direction, is likely to be a reasonable starting point for SSBL performance in similarly structured datasets.

The following lessons are likely to extend beyond the FLNG–LNGC configuration:

- SSBL is effective when response extrema follow consistent and predictable trends with the environmental input parameters. Responses governed by drift forces and restoring behaviour are well suited to this approach.
- SSBL can less effectively capture maximum responses that depend on wave grouping, phase timing or nonlinear multibody interaction. These require time-resolved inputs or direct simulation.
- For structured environmental matrices, covering approximately 20% of the relevant environmental regimes, well distributed across wave height, period and direction, is a useful starting point for

reliable SSBL performance, though the exact fraction depends on the density and uniformity of the dataset.

- SSBL cannot identify in advance which sea states require simulation. It reduces the cost of evaluation but not the need for a representative set of simulations from which to train.

Can multibody response behaviours be inferred from partial time histories?

Individual motions of the FLNG–LNGC system can be reliably forecast from partial time histories, but the degree of predictability depends on the physical nature of each response component. LF surge and sway remain predictable over extended horizons, with past–future correlation persisting over approximately 500–700 s. This reflects the slowly varying nature of drift forcing, which evolves gradually through wave-group modulation and quadratic drift interactions. As a result, the LF response structure is sufficiently regular that a partial motion history provides enough information to support forward extrapolation, though the decaying autocorrelation indicates this predictability is not unlimited and likely explains only part of the observed performance. WF motions such as heave and roll are predictable only when a compact representation of the future wave field is supplied to the model, as they respond directly to the instantaneous wave elevation rather than to past motion history. LNGC roll is the least predictable response due to its weak damping, sensitivity to multibody interaction and dependence on phase-specific excitation.

The following lessons are likely to extend beyond the FLNG–LNGC configuration:

- LF responses retain strong temporal memory and are naturally suited for multi-step forecasting once an initial segment is observed. Stable LF performance required approximately 60% of the three-hour record for training, reflecting the limited number of LF cycles available per simulation. For systems with longer drift periods, a larger training fraction or longer simulations may be needed.
- WF motions require explicit future-wave information to remain predictable. Stable WF performance was achieved using only approximately 40% of the three-hour record, as each simulation contains hundreds of WF cycles. This indicates that WF models are considerably more data-efficient than LF models within a single simulation.
- The physical nature of each response component determines its forecastability. Slowly varying responses governed by drift forcing are more tractable from partial histories, while rapidly oscillating responses depend on external forcing information to maintain accuracy.
- Segment-based TSBL cannot replace simulations for new environmental conditions. Because forecasting requires an initial motion history, it can shorten individual simulations but cannot generate peak responses for unsimulated sea states.

Under which environmental variations can full three-hour multibody responses be reconstructed from wave elevation and environmental inputs when predicting unseen wave conditions?

Full three-hour FLNG–LNGC response histories can be reconstructed accurately only when the environmental conditions remain within the regimes represented in the training data. Chapter 7 shows a clear hierarchy of generalisability across environmental parameters. Variations in significant wave

height $H_{s,\mathrm{sea}}$ are handled reliably, as amplitude scaling preserves the underlying spectral shape and wave-group structure. In contrast, changes in peak wave period $T_{p,\mathrm{sea}}$ introduce new spectral regimes that alter LF drift forcing, leading to rapid loss of reconstruction accuracy. Changes in wave direction θ_{sea} modify the vessel response amplitude operators through direction-dependent multibody interaction and gap-flow amplification, placing the response outside the mapping learned during training. LNGC roll is the most sensitive response in all cases, due to its weak damping and strong dependence on gap hydrodynamics and radiation–diffraction coupling.

Importantly, performance improves significantly when training data coverage is extended. The head-sea operability assessment in Chapter 7 demonstrated that TSBL trained on only 300 out of 756 head-sea simulations achieved 99.7% classification accuracy, and the final results in Chapter 8 confirm that a well-chosen set of 1,500 simulations spanning the key spectral and directional regimes supports reliable operability prediction. This shows that the generalisation limitations observed for unseen $T_{p,\mathrm{sea}}$ and θ_{sea} are not fundamental barriers, but reflect insufficient coverage of those regimes in the training set.

The following lessons are likely to extend beyond the FLNG–LNGC configuration:

- Spectral coverage of the training set governs reconstructability. Amplitude variations within a known spectral regime are handled reliably, while changes in spectral shape or wave direction require those regimes to be represented in training.
- LF drift responses are the most demanding to generalise, as they depend on wave-group time scales that change fundamentally with peak wave period.
- WF motions generalise across wave height variations but are sensitive to changes in wave direction through direction-dependent RAO variation and multibody coupling effects.
- Generalisation limitations are not fixed. Extending training coverage to include the relevant $T_{p,\mathrm{sea}}$ and θ_{sea} regimes substantially improves performance, as demonstrated by both the head-sea assessment and the final results chapter.

How can we construct a minimal dataset which captures all conditions and system dynamics?

A minimal dataset must retain the environmental and dynamical regimes that meaningfully influence operability while omitting cases that do not add new information. This thesis does not definitively identify the minimal dataset for FLNG–LNGC side-by-side operability assessment. Rather, it demonstrates that a substantially reduced dataset can support reliable operability prediction, and provides a set of principles for how such a dataset should be constructed.

The 1,500-simulation training set used in the final results chapter was constructed using a two-stage selection procedure combining stratified environmental sampling with response-space clustering, designed to ensure coverage of both the key spectral and directional regimes and the full range of operability-relevant response outputs. This approach achieved 98.2% SSBL classification accuracy and 86.7% TSBL classification accuracy on unseen sea states, demonstrating that a well-chosen compact dataset can support reliable operability prediction. However, it is not demonstrated that 1,500 simulations is the minimum required. Further reduction may be possible, particularly for SSBL, where performance was shown to saturate at approximately 20% of the full dataset. For TSBL, the current training set is likely still larger than strictly necessary for some regimes and insufficient for others, particularly for LF prediction across the full directional matrix.

The false-safe cases from Chapter 8 further define which conditions cannot be omitted. Beam and

near-beam headings, where gap resonance and coupled sway–yaw modes are most strongly excited, and short peak period conditions ($T_{p,\mathrm{sea}} = 6$–$7\,\mathrm{s}$) combined with high sea states, where LF drift forces and mooring tensions are highest, represent the boundaries of the operability envelope and must be included in any reduced training set.

The following lessons are likely to extend beyond the FLNG–LNGC configuration:

- Coverage of the relevant spectral and directional regimes outweighs the total number of simulations. Representative sampling of the main physical regimes is more important than increasing dataset size.
- A two-stage selection combining stratified environmental sampling with response-space clustering is an effective approach for constructing compact but representative training sets, ensuring both input-space coverage and output-range diversity.
- Critical regimes define the lower bound of the dataset. Conditions that introduce large variations in LF drift, roll and relative sway must be included, as removing them risks unsafe generalisation.
- Wave height can be sampled sparsely for fixed spectral period and heading, as changes in $H_{s,\mathrm{sea}}$ primarily scale the response amplitude without altering the underlying wave structure. Note that this holds for the present single-seed dataset and should be verified when multiple wave seeds are used.
- The minimal dataset is specific to the modelling objective. SSBL requires broad coverage of the environmental descriptors, while TSBL requires diverse dynamic responses with sufficient variation in spectral content and wave direction.

10.1.1. Integrated Interpretation

As this thesis aims to provide a helicopter view on the application of ML models to multibody offshore operability assessment, the FLNG–LNGC configuration was selected as a representative case study rather than an end in itself. The case shares the defining characteristics of a broader class of coupled floating-body operations, including offshore hydrocarbon offloading, CO_2 ship-to-ship transfer, and floating wind installation: response-based operability limits, nonlinear interaction between vessels, and dependence on costly time-domain simulations. The subquestion answers presented above provide the detailed evidence base. This subsection draws together the two central findings that emerge from thi thesis.

The first central finding is that, for a given number of simulated sea states, the additional information contained in full time-domain response histories does not automatically lead to improved operability classification. SSBL models trained on summary statistics achieved equal or better classification performance than TSBL models using the same underlying simulations, indicating that the information required for go/no-go decisions is already largely contained in response extrema. The added value of TSBL lies not in improved classification but in the temporal detail it provides: full response time series allow transient limit exceedances to be distinguished from sustained violations and support engineering judgement near operability boundaries.

The second central finding, and a novel contribution of this thesis, is that complete three-hour response time series of a multibody offshore system can be reconstructed using ML models. To the authors knowledge, this is the first demonstration that data-driven models can reproduce complete motion and load histories for a coupled multibody configuration. The conditions under which reconstruction is reliable are clearly defined: training data must span the key spectral and directional regimes of the system, and explicit LF forcing information beyond the wave elevation signal is needed for accurate

slow-drift prediction. Reconstruction accuracy is governed primarily by training data coverage rather than by network architecture.

At the same time, ML models cannot yet reliably predict responses in regimes not adequately covered in training, cannot identify in advance which conditions require simulation, and cannot replace physics-based modelling where responses depend on mechanisms not represented in the training inputs, most critically explicit LF forcing from wind, current and second-order drift. The role of ML is therefore not to replace time-domain simulation but to concentrate it where it is genuinely needed.

Whether these findings hold for other coupled floating-body systems has not been verified in this thesis.

10.2. Recommendations

The findings of this thesis demonstrate that full three hour multibody responses can be reconstructed and predicted with good accuracy for the selected FLNG and LNGC system. The Discussion also identified several structural limitations and opportunities for extending the framework. This section presents recommendations for future work. The recommendations are grouped into five areas. These are generalisation across different system configurations, improvement of low-frequency prediction through environmental inputs, extension of the input space to operational and directional data, enhancement of model reliability and data efficiency, and further exploration of the role of frequency domain modelling.

10.2.1. Application with Extended Input Spaces and Operational Data

The current framework relies on single-seed, single-peaked JONSWAP wave inputs, which limits its applicability to real operational conditions. The following extensions are needed to bring the framework closer to operational deployment.

- **Include multiple wave seeds and directional wave information:** Multiple seeds capture the distribution of possible wave-group realisations and are a prerequisite for meaningful uncertainty quantification. Directional spreading parameters provide a richer description of the sea state and allow models to capture the directional sensitivity of the multibody gap response.
- **Validate against operational measurements:** Operational data include directional complexity, steep and grouped waves, and extreme events absent from the synthetic training data. Validation against such data is the essential step to establish whether model limitations reflect physical system behaviour or artefacts of the idealised simulation inputs.

10.2.2. Provide Physically Complete Model Inputs

The most critical limitation identified in this thesis is that the TSBL models do not receive the physical forcing that drives the LF response. Wind and current loads contribute significantly to slow-drift motion and mooring tensions, yet they are not included as explicit model inputs. The extracted low-frequency drift forcing signal $F_{\mathrm{LF}}(t)$, derived from the quadratic wave elevation term as demonstrated in Section 9.2.5, provides a compact and physically meaningful representation of the difference-frequency excitation that governs slow drift. Including $F_{\mathrm{LF}}(t)$ together with explicit wind and current forcing as physics-informed inputs would directly address the systematic LF underprediction observed in strong wind and current cases. These are precisely the regimes where false-safe classifications are most concentrated, making this a valuable direction for future work.

- **Include the extracted LF drift forcing signal $F_{\mathrm{LF}}(t)$ as a physics-informed input to the LF model:** Model performance should then be compared with the current wave-only input setup while progressively reducing the training-set size. This would directly test whether explicitly providing

the physically relevant LF excitation allows the TSBL model to learn slow-drift behaviour more data-efficiently and narrows the observed performance gap between LF and WF prediction.

- **Include steady wind and current forcing as physics-informed inputs:** Steady wind and surface current velocity exert significant influence on drift motion. Incorporating these quantities directly provides the model with the physical drivers of mean offset and slow drift that are currently inferred indirectly from scalar environmental descriptors.

10.2.3. Improve Model Reliability and Data Efficiency

The discussion highlighted opportunities to improve model reliability and the efficiency of future simulation campaigns. These improvements can support the development of models that are transparent, robust and computationally efficient. The recommendations below build directly on the findings of this thesis.

- **Introduce uncertainty quantification:** Operability assessment is sensitive to prediction errors near threshold values. ML models should be able to report not just a predicted value but also how confident that prediction is. Approaches such as Bayesian neural networks, deep ensembles or Gaussian Process emulators can provide this, indicating where the model is extrapolating and where additional simulations are needed.
- **Use active learning to target new simulations:** Once uncertainty estimates are available, they can guide which new simulations to run. Rather than expanding the simulation matrix uniformly, active learning selects the conditions that reduce model uncertainty most effectively. For SSBL this means targeting poorly constrained response regimes, for TSBL it means targeting conditions with complex LF drift behaviour. This makes simulation campaigns more efficient and could effectively lead to a minimal data set to capture all dynamics.

These recommendations align with recent work that unifies active learning and Bayesian optimisation as goal-driven adaptive sampling strategies, in which uncertainty-aware models guide the efficient acquisition of new data [113].

10.2.4. Towards a General Operability Screening Model

The low data requirements and fast prediction of SSBL make it well suited for training across multiple system configurations, locations and environmental regimes. A model trained on a sufficiently diverse dataset, spanning different hull forms, mooring arrangements, water depths and site conditions, could learn patterns that generalise across coupled floating-body operations, moving towards a general-purpose operability screening tool that does not require retraining for each new project. This direction is most tractable for SSBL given its low computational cost and data efficiency, and represents a natural next step.

10.2.5. Exploring the Role of Frequency Domain Modelling

The discussion in Section 9.4 raised the question of whether linearised frequency domain modelling could serve as a different strategy for reducing simulation dependence. While frequency domain models cannot capture all multibody effects, they are orders of magnitude faster then time domain simulations.

- **Construct and validate a linearised frequency domain model for the coupled FLNG–LNGC system across headings and sea states:** This would establish where frequency domain predictions are sufficiently accurate and where nonlinear time-domain simulations remain necessary,

while also providing a reference to benchmark against the ML model results and assess where each approach is most appropriate.

- **Use frequency domain models for rapid pre-screening:** Frequency domain predictions can quickly filter out sea states that are clearly within operability limits, reducing the number of time-domain simulations needed for both training and evaluation.
- **Investigate frequency domain generated datasets as a low-cost source of training data:** Frequency domain models make it possible to generate responses with very low computational effort, allowing the creation of large datasets that span many system configurations and site conditions, including different hull forms, mooring and fender arrangements and water depths. Training ML models on such extended datasets may support earlier application of ML-based operability assessment in the design process, before time-domain simulation campaigns are available.

References

[1] "DNV-RP-H103: Modelling and Analysis of Marine Operations". In: (2011).

[2] Finlay McPhail et al. "Prelude Side by Side Offloading – Quantifying the Qualitative to Assess Economics and Operational Capability". In: *Offshore Technology Conference*. OTC-30876-MS. OTC. Houston, TX: Offshore Technology Conference, 2020.

[3] Andrzej Grzybowski, Katarzyna Pawlikowska–Łagód, and W. Clark Lambert. "A History of Artificial Intelligence". In: *Clinics in Dermatology* 42 (2024), pp. 221–229. DOI: 10.1016/j.clindermatol.2023.12.016.

[4] Iqbal H. Sarker. "Machine Learning: Algorithms, Real-World Applications and Research Directions". In: *SN Computer Science* 2.160 (2021). DOI: 10.1007/s42979-021-00592-x. URL: https://link.springer.com/article/10.1007/s42979-021-00592-x.

[5] Ioannis Antonopoulos et al. "Artificial Intelligence and Machine Learning Approaches to Energy Demand-Side Response: A Systematic Review". In: *Renewable and Sustainable Energy Reviews* 130 (2020), p. 109899. DOI: 10.1016/j.rser.2020.109899.

[6] Jiawei Han, Micheline Kamber, and Jian Pei. *Data Mining: Concepts and Techniques*. Elsevier, 2011.

[7] M. Mohammed, M.B. Khan, and M.E. Bashier. *Machine Learning: Algorithms and Applications*. CRC Press, 2016.

[8] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. "Learning representations by back-propagating errors". In: *Nature* 323 (1986), pp. 533–536. DOI: 10.1038/323533a0.

[9] Yann LeCun et al. "Gradient-based learning applied to document recognition". In: *Proceedings of the IEEE* 86.11 (1998), pp. 2278–2324. DOI: 10.1109/5.726791.

[10] Zachary C. Lipton, John Berkowitz, and Charles Elkan. "A critical review of recurrent neural networks for sequence learning". In: *arXiv preprint arXiv:1506.00019* (2015). URL: https://arxiv.org/abs/1506.00019.

[11] Sepp Hochreiter and Jürgen Schmidhuber. "Long Short-Term Memory". In: *Neural Computation* 9 (Nov. 1997), pp. 1735–1780. DOI: 10.1162/neco.1997.9.8.1735.

[12] Sandeep Singh Sengar et al. "Generative Artificial Intelligence: A Systematic Review and Applications". In: *Multimedia Tools and Applications* 84 (2024), pp. 23661–23700. DOI: 10.1007/s11042-024-20016-1. URL: https://link.springer.com/article/10.1007/s11042-024-20016-1.

[13] Mourad Gridach et al. "Agentic AI for Scientific Discovery: A Survey of Progress, Challenges, and Future Directions". In: *arXiv preprint arXiv:2503.08979* (2025). DOI: 10.48550/arXiv.2503.08979. URL: https://arxiv.org/abs/2503.08979.

[14] Zane Durante et al. "Agent AI: Surveying the Horizons of Multimodal Interaction". In: *arXiv preprint arXiv:2401.03568* (2024). DOI: 10.48550/arXiv.2401.03568. URL: https://arxiv.org/abs/2401.03568.

[15] Kamal Hassan et al. “Application of Artificial Intelligence in Aerospace Engineering and Its Future Directions: A Systematic Quantitative Literature Review”. In: *Archives of Computational Methods in Engineering* 31.7 (Sept. 2024), pp. 4031–4086. ISSN: 1886-1784. DOI: 10.1007/s11831-024-10105-7. (Visited on 07/28/2025).

[16] Ahmed A. Abdou El-Abbasy. “Artificial Intelligence-Driven Predictive Modeling in Civil Engineering: A Comprehensive Review”. In: *Journal of Umm Al-Qura University for Engineering and Architecture* (July 2025). ISSN: 1658-8150. DOI: 10.1007/s43995-025-00166-5. (Visited on 07/28/2025).

[17] General Electric. *GE Digital Predictive Maintenance Solutions*. https://www.gevernova.com/digital/solutions/predictive-maintenance. Accessed: 2025-08-07. 2023.

[18] Waymo Research Team. “Waymo Open Dataset: Advancing Autonomous Driving with Real-World Data”. In: *Waymo Research* (2020). Accessed: 2025-08-07. URL: https://waymo.com/research/open-dataset.

[19] Airbus and Autodesk. *Generative Design in Aerospace: Airbus Case Study*. https://www.autodesk.com/customer-stories/airbus-generative-design. Accessed: 2025-08-07. 2019.

[20] Google DeepMind. *Machine Learning for Data Center Energy Efficiency*. https://deepmind.google/impact/energy-efficiency/. Accessed: 2025-08-07. 2016.

[21] Autodesk. *Autodesk and Airbus show the future of aerospace design and manufacture in pioneering generatively designed 3D printed partition*. Accessed: 2025-08-07. 2015. URL: https://adsknews.autodesk.com/en/news/autodesk-and-airbus-show-the-future-of-aerospace-design-and-manufacture-in-pioneering-generatively-d.

[22] *Enhance Oil & Gas Recovery and Safety: How Meridium APM and SmartSignal AI/ML Predictive Analytics Drive Operational Excellence*. Tech. rep. Whitepaper. USA: GE Vernova, 2025. URL: https://www.gevernova.com/software.

[23] Zeeshan Tariq et al. “A Systematic Review of Data Science and Machine Learning Applications to the Oil and Gas Industry”. In: *Journal of Petroleum Exploration and Production Technology* 11.12 (Dec. 1, 2021), pp. 4339–4374. ISSN: 2190-0566. DOI: 10.1007/s13202-021-01302-2. URL: https://doi.org/10.1007/s13202-021-01302-2 (visited on 07/28/2025).

[24] Alain Gomez-Cabrera and Ponciano Jorge Escamilla-Ambrosio. “Review of Machine-Learning Techniques Applied to Structural Health Monitoring Systems for Building and Bridge Structures”. In: *Applied Sciences* 12.21 (2022), p. 10754. DOI: 10.3390/app122110754. URL: https://www.mdpi.com/2076-3417/12/21/10754.

[25] Isaac Kofi Nti et al. “Applications of artificial intelligence in engineering and manufacturing: a systematic review”. In: *Journal of Intelligent Manufacturing* 33 (2022), pp. 1581–1601. DOI: 10.1007/s10845-021-01771-6.

[26] Eric VanDerHorn and Sankaran Mahadevan. “Digital Twin: Generalization, Characterization and Implementation”. In: *Decision Support Systems* 145 (2021), p. 113524. DOI: 10.1016/j.dss.2021.113524.

[27] Amil Merchant et al. “Scaling Deep Learning for Materials Discovery”. In: *Nature* 624.7990 (Dec. 1, 2023), pp. 80–85. ISSN: 1476-4687. DOI: 10.1038/s41586-023-06735-9. URL: https://doi.org/10.1038/s41586-023-06735-9.

[28] Claudio Rodríguez Castillo et al. “A Critical Review of Challenges and Opportunities for Offshore Structures Supporting Renewable Hydrogen”. In: *Wind Energy Science* 9 (2024), pp. 533–558. DOI: 10.5194/wes-9-533-2024. URL: https://wes.copernicus.org/articles/9/533/2024/.

[29] Nima Amiri et al. “A comprehensive review on design, monitoring, and failure in offshore platforms”. In: *Journal of Marine Science and Engineering* 9.12 (2021), p. 1349. DOI: 10.3390/jmse9121349. URL: https://doi.org/10.3390/jmse9121349.

[30] W. Guachamin Acero, Z. Gao, and T. Moan. “Methodology for assessment of the allowable sea states during installation of an offshore wind turbine transition piece structure onto a monopile foundation”. In: *Journal of Offshore Mechanics and Arctic Engineering* 138 (2016).

[31] Wenhua Zhao et al. “Current practice and research directions in hydrodynamics for FLNG-side-by-side offloading”. In: *Ocean Engineering* 158 (2018), pp. 99–110. DOI: 10.1016/j.oceaneng.2018.03.076.

[32] John V. Sharp and John Wintle. *Assessment of Materials in Ageing Offshore Structures*. eBook ISBN: 978-3-031-84489-8. Cham, Switzerland: Springer Nature Switzerland AG, 2025. ISBN: 978-3-031-84488-1. DOI: 10.1007/978-3-031-84489-8.

[33] Torgeir Moan. “Life cycle structural integrity management of offshore structures”. In: *Structure and Infrastructure Engineering* 14.7 (2018), pp. 911–927. DOI: 10.1080/15732479.2018.1438478. URL: https://www.tandfonline.com/doi/pdf/10.1080/15732479.2018.1438478.

[34] Hadi Pezeshki et al. “State of the art in structural health monitoring of offshore and marine structures”. In: *Proceedings of the Institution of Civil Engineers - Maritime Engineering* 176.2 (2023), pp. 89–108. DOI: 10.1680/jmaen.2022.027.

[35] Jeremy T. Claisse et al. “Impacts from Partial Removal of Decommissioned Oil and Gas Platforms on Fish Biomass and Production”. In: *Proceedings of the National Academy of Sciences* 111.43 (2014), pp. 15462–15467. DOI: 10.1073/pnas.1422274112.

[36] Ashley M. Fowler et al. “A multi-criteria decision approach to decommissioning of offshore oil and gas infrastructure”. In: *Ocean & Coastal Management* 87 (2014), pp. 20–29. DOI: 10.1016/j.ocecoaman.2013.10.019.

[37] Climate and Pollution Agency. *Decommissioning of Offshore Installations*. Tech. rep. Ministry of the Environment, Norway, 2011. URL: https://www.osti.gov/etdeweb/biblio/1024425.

[38] Brigitte Sommer et al. “Decommissioning of offshore oil and gas structures – Environmental opportunities and challenges”. In: *Science of The Total Environment* 658 (2019), pp. 973–981. DOI: 10.1016/j.scitotenv.2018.12.193. URL: https://www.sciencedirect.com/science/article/pii/S0048969718350551.

[39] Kristina Shea, Robert Aish, and Marina Gourtovaia. “Towards Integrated Performance-Driven Generative Design Tools”. In: *Automation in Construction*. Education and Research in Computer Aided Architectural Design in Europe (eCAADe 2003), Digital Design 14.2 (Mar. 1, 2005), pp. 253–264. ISSN: 0926-5805. DOI: 10.1016/j.autcon.2004.07.002. URL: https://www.sciencedirect.com/science/article/pii/S0926580504000809 (visited on 02/23/2026).

[40] V. Singh and N. Gu. “Towards an integrated generative design framework”. In: *Design Studies* 33.2 (2012), pp. 185–207. DOI: 10.1016/j.destud.2011.06.001.

[41] S. Krish. “A practical generative design method”. In: *Computer-Aided Design* 43.1 (2011), pp. 88–100. DOI: 10.1016/j.cad.2010.09.009.

[42] Meysam Karimi et al. “A multi-objective design optimization approach for floating offshore wind turbine support structures”. In: *Journal of Ocean Engineering and Marine Energy* 3.1 (2017), pp. 69–87. DOI: 10.1007/s40722-016-0072-4. URL: https://doi.org/10.1007/s40722-016-0072-4.

[43] Seowoo Jang, Soyoung Yoo, and Namwoo Kang. “Generative Design by Reinforcement Learning: Enhancing the Diversity of Topology Optimization Designs”. In: *Computer-Aided Design* 146 (2022), p. 103225. DOI: 10.1016/j.cad.2022.103225.

[44] Qinshuo Shen et al. “Metamodel-Based Generative Design of Wind Turbine Foundations”. In: *Automation in Construction* 138 (June 1, 2022), p. 104233. ISSN: 0926-5805. DOI: 10.1016/j.autcon.2022.104233. URL: https://www.sciencedirect.com/science/article/pii/S0926580522001066 (visited on 07/29/2025).

[45] Nouman Saeed et al. “A novel design of multifunctional offshore floating platform structure based on topology optimization”. In: *Engineering Structures* 306 (2024), p. 117782. DOI: 10.1016/j.engstruct.2024.117782. URL: https://doi.org/10.1016/j.engstruct.2024.117782.

[46] Hongbo Sun and Ling Ma. “Structural Topology Optimization by Combining BESO with Reinforcement Learning”. In: *Journal of Harbin Institute of Technology* 53.1 (2021), pp. 1–10. DOI: 10.11916/j.issn.1005-9113.19078. URL: https://hit.alljournals.cn/jhit_cn/ch/reader/view_abstract.aspx?file_no=20210109&flag=1.

[47] *DNV-OS-H101: Marine Operations, General*. Tech. rep. Offshore Standard. Høvik, Norway: Det Norske Veritas (DNV), 2011. URL: https://studylib.net/doc/18091064/dnv-os-h101--marine-operations--general.

[48] P. S. Indrevoll et al. “Optimizing Offshore Operations: Integrating Machine Learning with Vessel Motion Based Criteria for Accurate and Efficient Operability Assessments”. In: Offshore Technology Conference. May 5, 2025, D022S059R003. DOI: 10.4043/35726-MS. URL: https://doi.org/10.4043/35726-MS (visited on 03/01/2026).

[49] G. Clauss and T. Riekert. “Operational limitations of offshore crane vessels”. In: *Offshore Technology Conference*. OTC 6217. Houston, Texas, 1990.

[50] G. Clauss and T. Riekert. “Seakeeping Characteristics and Related Workability of Offshore Crane Vessels”. In: *Schiffstechnik* 37 (1990).

[51] J.J. de Wilde et al. “Direct Time Domain Downtime Assessment for LNG Operations using Computer Cluster”. In: *Proceedings of the Nineteenth International Offshore and Polar Engineering Conference (ISOPE)*. Osaka, Japan: The International Society of Offshore and Polar Engineers, 2009, pp. 127–134. URL: https://onepetro.org/ISOPEIOPEC/proceedings/ISOPE09/ISOPE09/ISOPE-I-09-170/7546.

[52] L. Li et al. “Assessment of allowable sea states during installation of offshore wind turbine monopiles with shallow penetration in the seabed”. In: *Journal of Offshore Mechanics and Arctic Engineering* 138 (2016).

[53] Shiliang Duan et al. “A LSTM Deep Learning Model for Deterministic Ship Motions Estimation Using Wave-Excitation Inputs”. In: *Proceedings of the 29th International Ocean and Polar Engineering Conference*. International Society of Offshore and Polar Engineers. Honolulu, Hawaii, USA, 2019.

[54] B. Mak and T.P. Scholcz. “Ship as a Wave Buoy: Estimating Relative Wave Direction From In-Service Ship Motion Measurements Using Deep Learning”. In: *Proceedings of the 38th International Conference on Ocean, Offshore and Arctic Engineering (OMAE)*. ASME. 2019. DOI: 10.1115/OMAE2019-96145.

[55] Xiaoxian Guo et al. “Probabilistic prediction of the heave motions of a semi-submersible by a deep learning problem model”. In: *Ocean Engineering* 245 (2022), p. 110504. ISSN: 0029-8018. DOI: 10.1016/j.oceaneng.2022.110504. URL: https://www.sciencedirect.com/science/article/pii/S0029801822000506.

[56] Faisal Khan and Mahmoud Haddara. “Risk-based maintenance (RBM): a new approach for maintenance decision-making”. In: *Journal of Loss Prevention in the Process Industries* 41 (2016), pp. 59–71. DOI: 10.1016/j.jlp.2016.02.009.

[57] Eric J. Tuegel et al. “Reengineering Aircraft Structural Life Prediction Using a Digital Twin”. In: *International Journal of Aerospace Engineering* 2011 (2011), p. 154798. DOI: 10.1155/2011/154798.

[58] Ruhan Liu et al. “The Paradigm of Digital Twin Communications”. In: *arXiv preprint* arXiv:2105.07182 (2021). URL: https://arxiv.org/abs/2105.07182.

[59] Jincheng Sha et al. “Research Progress in Predictive Maintenance of Offshore Platform Structures Based on Digital Twin Technology”. In: *Journal of Marine Science and Application* (2025). DOI: 10.1007/s11804-025-00649-w.

[60] Thyago P. Carvalho et al. “A systematic literature review of machine learning methods applied to predictive maintenance”. In: *Computers & Industrial Engineering* 137 (2019), p. 106024. DOI: 10.1016/j.cie.2019.106024.

[61] B.Q. Chen, C. Guedes Soares, and P.M. Videiro. “Review of Digital Twin of Ships and Offshore Structures”. In: *Developments in Maritime Technology and Engineering*. CRC Press, 2021, pp. 445–451. DOI: 10.1201/9781003216582-50.

[62] Concetta Semeraro et al. “Data-driven invariant modelling patterns for digital twin design”. In: *Journal of Industrial Information Integration* 31 (2023), p. 100424. DOI: 10.1016/j.jii.2022.100424.

[63] Mahadi Hasan Imran et al. “A critical analysis of machine learning in ship, offshore, and oil & gas corrosion research, part I: Corrosion detection and classification”. In: *Ocean Engineering* 313 (2024), p. 119600. DOI: 10.1016/j.oceaneng.2024.119600.

[64] Abdulaziz Sindi et al. “Advancing Digital Healthcare Engineering for Aging Ships and Offshore Structures: An In-Depth Review and Feasibility Analysis”. In: *Data-Centric Engineering* 5 (2024), e18. DOI: 10.1017/dce.2024.14. URL: https://doi.org/10.1017/dce.2024.14.

[65] Zhen Gao et al. “Numerical simulation of marine operations and prediction of operability using response-based criteria with an application to installation of offshore wind turbine support structures”. In: Sept. 2016.

[66] Filip Neele et al. “CO2 Transport by Ship: The Way Forward in Europe”. In: *Energy Procedia* 114 (2017), pp. 6824–6834. ISSN: 1876-6102. DOI: 10.1016/j.egypro.2017.03.1813. URL: https://www.sciencedirect.com/science/article/pii/S1876610217320155.

[67] Mun Sung Kim, Mun Keun Ha, and Byung Woo Kim. “Relative motions between LNG-FPSO and side-by-side positioned LNG carrier in waves”. In: *Proceedings of the 13th International Offshore and Polar Engineering Conference*. ISOPE. 2003.

[68] A. Ashok, A. Nap, and J. J. de Wilde. *Time Domain Simulations for the ROTOR Tool: Hydrodynamic Numerical Analysis, Volume I – Text Report (Final)*. Tech. rep. 30311-1-PO. Ordered by Shell Global Solutions International B.V. Wageningen, The Netherlands: MARIN (Maritime Research Institute Netherlands), June 2017.

[69] *Guidelines for Offshore Tanker Operations (GOTO)*. Tech. rep. First Edition. London, United Kingdom: Oil Companies International Marine Forum (OCIMF), 2018.

[70] A. Ashok, A. Nap, and J. J. de Wilde. *Extension to the ROTOR Database: Hydrodynamic Numerical Analysis, Volume I – Text Report (Final)*. Tech. rep. 31189-1-PO. Ordered by Shell Global Solutions International B.V. Wageningen, The Netherlands: MARIN (Maritime Research Institute Netherlands), Oct. 2020.

[71] Wenhua Zhao et al. “Hydrodynamics of an FLNG system in tandem offloading operation”. In: *Ocean Engineering* 57 (2013), pp. 150–162. DOI: 10.1016/j.oceaneng.2012.09.015.

[72] O.M. Faltinsen. *Sea Loads on Ships and Offshore Structures*. No DOI. See: https://books.google.com/books/about/Sea_Loads_on_Ships_and_Offshore_Structur.html?id=qZq4Rs2DZXoC. Cambridge University Press, 1990. ISBN: 978-0521458702.

[73] L. Sun, R. Eatock Taylor, and P.H. Taylor. “First- and second-order analysis of resonant waves between adjacent barges”. In: *Journal of Fluids and Structures* 26.6 (2010), pp. 954–978. DOI: 10.1016/j.jfluidstructs.2010.03.002.

[74] M. Perić and C. Swan. “An experimental study of the wave excitation in the gap between two closely spaced bodies, with implications for LNG offloading”. In: *Applied Ocean Research* 51 (2015), pp. 320–330. DOI: 10.1016/j.apor.2015.01.010.

[75] R.H.M. Huijsmans, J.A. Pinkster, and J.J. de Wilde. “Diffraction and radiation of waves around side-by-side moored vessels”. In: *Proceedings of the International Offshore and Polar Engineering Conference (ISOPE)*. Stavanger, Norway, 2001.

[76] B. Buchner, A. W. V. van Dijk, and J. J. de Wilde. “Numerical Multiple-Body Simulations of Side-by-Side Mooring to an FPSO”. In: *Proceedings of the Eleventh International Offshore and Polar Engineering Conference (ISOPE 2001)*. Demonstrates importance of full off-diagonal added-mass and retardation matrices; discusses gap-flow treatment. Stavanger, Norway, 2001.

[77] Mamoun Naciri, Olaf Waals, and Jaap de Wilde. “Time Domain Simulations of Side-by-Side Moored Vessels: Lessons Learnt From a Benchmark Test”. In: *Proceedings of the 26th International Conference on Offshore Mechanics and Arctic Engineering (OMAE 2007)*. Paper OMAE2007-29756. San Diego, California, USA: ASME, June 2007, OMAE2007–29756. DOI: 10.1115/OMAE2007-29756.

[78] J.R. Fournier, M. Naciri, and X.-B. Chen. “Hydrodynamics of two side-by-side vessels – Experiments and numerical simulations”. In: *Proceedings of the International Offshore and Polar Engineering Conference (ISOPE)*. ISOPE06-PT-02. San Francisco, USA, 2006.

[79] W.H. Pauw, R.H.M. Huijsmans, and A.J. Voogt. “Advances in the hydrodynamics of side-by-side moored vessels”. In: *26th International Conference on Offshore Mechanics and Arctic Engineering (OMAE)*. OMAE2007-29374. San Diego, California, USA, 2007.

[80] J.T.M. van Doorn and D. ten Hove. “FPSO Offloading Concepts, their Nautical Feasibility and Safety Assessment”. In: *Offshore Technology Conference*. OTC-14296. Houston, USA, 2002.

[81] R. Grin, J.J. de Wilde, and J.T.M. van Doorn. “Assessment of LNG transport chains using weather-based voyage simulations”. In: *SNAME 2005 ETT* (2005). Paper 003.

[82] L. Poldervaart, H. Oomen, and J. Ellis. “Offshore LNG transfer: a worldwide review of offloading availability”. In: *Offshore Technology Conference*. OTC-18026. Houston, USA, 2006.

[83] Maritime Research Institute Netherlands (MARIN). *Side-by-side offloading operability analysis of spread moored FLNG. - Summary Report*. Tech. rep. Wageningen, The Netherlands: MARIN, 2025.

[84] Xiaoxian Guo et al. “Predicting heave and surge motions of a semi-submersible with neural networks”. In: *Applied Ocean Research* 112 (2021), p. 102708. DOI: 10.1016/j.apor.2021.102708.

[85] Tao Zhang, Xiao-Qing Zheng, and Ming-Xin Liu. “Multiscale attention-based LSTM for ship motion prediction”. In: *Ocean Engineering* 230 (2021), p. 109066. DOI: 10.1016/j.oceaneng.2021.109066.

[86] Sepp Hochreiter and Jürgen Schmidhuber. “Long Short-Term Memory”. In: *Neural Computation* 9.8 (1997), pp. 1735–1780. DOI: 10.1162/neco.1997.9.8.1735.

[87] Jae-Hoon Lee et al. “Prediction of wave-induced ship motions based on integrated neural network system and spatiotemporal wave-field data”. In: *Physics of Fluids* 35 (2023), p. 097127. DOI: 10.1063/5.0163795.

[88] Ashish Vaswani et al. “Attention is All You Need”. In: *Advances in Neural Information Processing Systems* 30 (2017), pp. 5998–6008.

[89] Mingyang Zhang et al. “A deep learning method for the prediction of 6-DoF ship motions in real conditions”. In: *Proceedings of the Institution of Mechanical Engineers, Part M: Journal of Engineering for the Maritime Environment* 237.4 (2023), pp. 887–905. DOI: 10.1177/14750902231157852.

[90] Haoyi Zhou et al. “Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 11677–11688.

[91] Wenhe Shen et al. “A pre-trained multi-step prediction informer for ship motion prediction with a mechanism-data dual-driven framework”. In: *Engineering Applications of Artificial Intelligence* 139 (2025), p. 109523. DOI: 10.1016/j.engappai.2025.109523.

[92] Yucheng Liu et al. “The input vector space optimization for LSTM deep learning model in real-time prediction of ship motions”. In: *Ocean Engineering* 213 (2020), p. 107681. DOI: 10.1016/j.oceaneng.2020.107681.

[93] Yu Yang, Tao Peng, and Shijun Liao. “Predicting future mooring line tension of floating structure by machine learning”. In: *Ocean Engineering* 269 (2023), p. 113470. DOI: 10.1016/j.oceaneng.2023.113470.

[94] Lang-Wei Zhong et al. “From Classical Models to Neural Operators: A Study on Mooring Tension Prediction for Floating Offshore Wind Turbine”. In: *Advanced Steel Construction* 21.4 (2025), pp. 357–366. DOI: 10.18057/IJASC.2025.21.4.8.

[95] Erwan Auburtin, Thiago Miliante, and Diego Lima. “Prediction of Coral South FLNG Side-by-Side Offloading Operability Using Radial Basis Function Artificial Neural Networks”. In: *Proceedings of the Offshore Technology Conference (OTC)*. OTC. Houston, Texas, 2019.

[96] Ulrik D. Nielsen, Astrid H. Brodtkorb, and Jørgen J. Jensen. “Response predictions using the observed autocorrelation function”. In: *Marine Structures* 58 (2018), pp. 31–52. DOI: 10.1016/j.marstruc.2017.10.012.

[97] MARIN. *AnySim Software – MARIN*. https://www.marin.nl/en/about/facilities-and-tools/software/anysim. n.d. URL: https://www.marin.nl/en/about/facilities-and-tools/software/anysim.

[98] J.M.J. Journée and W.W. Massie. *Offshore Hydromechanics*. 2nd ed. Based on original lecture notes (2000). 2008.

[99] Tin Kam Ho. "Random Decision Forests". In: *Proceedings of the Third International Conference on Document Analysis and Recognition*. IEEE. 1995, pp. 278–282. DOI: 10.1109/ICDAR.1995.598994.

[100] Leo Breiman. "Random Forests". In: *Machine Learning* 45.1 (2001), pp. 5–32. DOI: 10.1023/A:1010933404324.

[101] Gopinath Rebala, Ajay Ravi, and Sanjay Churiwala. *An Introduction to Machine Learning*. Springer, 2019. DOI: 10.1007/978-3-030-15729-6. URL: https://link.springer.com/book/10.1007/978-3-030-15729-6.

[102] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. "Learning representations by back-propagating errors". In: *Nature* 323.6088 (1986), pp. 533–536.

[103] George Box. "Box and Jenkins: Time Series Analysis, Forecasting and Control". In: *A Very British Affair: Six Britons and the Development of Time Series Analysis During the Twentieth Century*. Ed. by Terence C. Mills. Palgrave Advanced Texts in Econometrics. London: Palgrave Macmillan, 2013, pp. 161–215. DOI: 10.1057/9781137291264_6. URL: https://link.springer.com/chapter/10.1057/9781137291264_6.

[104] Karl Pearson. "Mathematical Contributions to the Theory of Evolution. III. Regression, Heredity, and Panmixia". In: *Philosophical Transactions of the Royal Society of London A* 187 (1896), pp. 253–318.

[105] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd ed. Springer Series in Statistics. New York: Springer, 2009. DOI: 10.1007/978-0-387-84858-7.

[106] S. Lloyd. "Least Squares Quantization in PCM". In: *IEEE Transactions on Information Theory* 28.2 (Mar. 1982), pp. 129–137. ISSN: 1557-9654. DOI: 10.1109/TIT.1982.1056489. URL: https://ieeexplore.ieee.org/document/1056489/authors (visited on 04/09/2026).

[107] Stephen Butterworth. "On the Theory of Filter Amplifiers". In: *Wireless Engineer* 7 (1930). Reprinted online; original publication from 1930, pp. 536–541. URL: https://www.changpuak.ch/electronics/downloads/On_the_Theory_of_Filter_Amplifiers.pdf.

[108] Optuna Contributors. *Optuna: A Hyperparameter Optimization Framework*. https://github.com/optuna/optuna. Accessed: 2026-03-03. 2026.

[109] K. Hasselmann et al. *Measurements of Wind–Wave Growth and Swell Decay during the Joint North Sea Wave Project (JONSWAP)*. Tech. rep. Deutsches Hydrographisches Institut, 1973.

[110] D. E. Cartwright and M. S. Longuet-Higgins. "The statistical distribution of the maxima of a random function". In: *Proceedings of the Royal Society of London A* 237 (1956), pp. 212–232.

[111] J. A. Pinkster. "Low Frequency Second Order Wave Exciting Forces on Floating Structures". PhD thesis. TU Delft, 1980.

[112] Andrei Metrikine. *Random Vibrations: Lecture 6 – Dynamic Response of a Nonlinear SDOF to Random Excitation*. PDF lecture slides. Lecture notes, course CIEM5220 – Applied Dynamics of Structures, Delft University of Technology. 2025.

[113] Francesco Di Fiore, Michela Nardelli, and Laura Mainini. “Active Learning and Bayesian Optimization: A Unified Perspective to Learn with a Goal”. In: *Archives of Computational Methods in Engineering* 31 (2024), pp. 2985–3013. DOI: 10.1007/s11831-024-10064-z. URL: https://link.springer.com/article/10.1007/s11831-024-10064-z.

[114] David Silver et al. “Mastering the game of Go without human knowledge”. In: *Nature* 550 (2017), pp. 354–359. DOI: 10.1038/nature24270.

[115] Pierre Baldi. “Autoencoders, Unsupervised Learning, and Deep Architectures”. In: *Proceedings of the ICML Workshop on Unsupervised and Transfer Learning*. Ed. by Isabelle Guyon et al. Vol. 27. JMLR Workshop and Conference Proceedings. Brookline, MA: JMLR, 2012, pp. 37–49. URL: http://proceedings.mlr.press/v27/baldi12a.html.

[116] Geoffrey E. Hinton. “A Practical Guide to Training Restricted Boltzmann Machines”. In: *Neural Networks: Tricks of the Trade*. Ed. by Grégoire Montavon, Geneviève B. Orr, and Klaus-Robert Müller. Vol. 7700. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2012, pp. 599–619. DOI: 10.1007/978-3-642-35289-8_32. URL: https://doi.org/10.1007/978-3-642-35289-8_32.

[117] Ian J. Goodfellow et al. *Generative Adversarial Networks*. 2014. arXiv: 1406.2661 [stat.ML]. URL: https://arxiv.org/abs/1406.2661.

[118] Diederik P. Kingma and Max Welling. “Auto-Encoding Variational Bayes”. In: *arXiv preprint arXiv:1312.6114* (2013). URL: https://arxiv.org/abs/1312.6114.

[119] Ashish Vaswani et al. “Attention Is All You Need”. In: *Advances in Neural Information Processing Systems* 30 (2017). URL: https://arxiv.org/abs/1706.03762.

[120] Tim Salimans and Jonathan Ho. “Progressive Distillation for Fast Sampling of Diffusion Models”. In: *arXiv preprint arXiv:2202.00512* (2022). DOI: 10.48550/arXiv.2202.00512. URL: https://arxiv.org/abs/2202.00512.

[121] W. E. Cummins. “The Impulse Response Function and Ship Motions”. In: *Symposium on Ship Theory, Institut für Schiffbau, Universität Hamburg*. DTMB Report 1661. 1962.

[122] T. F. Ogilvie. “Recent Progress Toward the Understanding and Prediction of Ship Motions”. In: *Proceedings of the Fifth Symposium on Naval Hydrodynamics*. Foundational transform relations linking frequency-domain hydrodynamic coefficients to time-domain retardation functions. Bergen, Norway, 1964.

[123] Michael A. Nielsen. *Neural Networks and Deep Learning*. Online book. Determination Press, 2015. URL: http://neuralnetworksanddeeplearning.com/.

A

General Review of of Technology

Table A.1: Detailed Overview of AI Methods by Subfield

Subfield	Method	Capabilities	Limitations
ML	Supervised Learning	Models learn from labeled data, used for classification and regression.	Requires a large amount of labeled data, sensitive to noise in data.
	Unsupervised Learning	Identifies hidden patterns or intrinsic structures in unlabeled data, useful for clustering and anomaly detection.	Difficult to evaluate results, lacks clear outputs.
	Reinforcement Learning (RL) [114]	Learns optimal actions through trial and error in dynamic environments. Used in robotics, gaming, and autonomous systems.	Requires large amounts of data, can be sample inefficient, complex reward tuning.
DL	CNN [9]	Specializes in image recognition and spatial data processing.	Needs large datasets, computationally expensive, and memory-intensive.
	RNN [10]	Sequence modeling, useful for time-series or natural language processing.	Struggles with long-term dependencies, suffers from vanishing gradient problems.
	LSTM [11]	A type of RNN that addresses the long-term dependency problem, improving sequence modeling.	Complex structure, slow training time.
	Autoencoder [115]	Unsupervised learning, useful for feature extraction and data compression.	Can overfit, sensitive to architecture tuning, not always generalizable.
	RBM [116]	Probabilistic modeling and dimensionality reduction, often used in deep learning models.	Difficult to train, less popular now in deep learning research.
	DRL [114]	Decision-making in dynamic environments, integrates deep learning with reinforcement learning.	Sample inefficiency, hard to tune reward functions.
Gen AI	GAN [117]	High-quality image and data generation, adversarial training.	Training instability, mode collapse, hard to tune.
	VAE [118]	Generates diverse data samples, with a structured latent space.	Blurry outputs, less sharp than GANs.
	Transformer (TRM) [119]	Text and code generation, excels in contextual understanding.	Requires significant computational power, sensitive to prompt phrasing.
	Diffusion Model [120]	Stable, high-quality sample generation, works well for images and videos.	Slow sampling process, computationally heavy.
Agentic AI	Agentic AI [13, 14]	Autonomous task execution, reasoning, tool use, decision-making.	Limited memory, lacks long-term planning and causal reasoning.

B

Theoretical Framework aNySIM

Side-by-side offloading places two large floating bodies in close proximity under the influence of the same metocean environment. Their relative motions are constrained by manifolds, fenders and mooring lines, creating a strongly coupled hydromechanical system. The narrow gap between the hulls modifies the local hydrodynamic field. Incident waves diffract around the first hull and radiated waves from both bodies interact within the gap, producing amplified pressures, velocities and free-surface elevations. These effects may excite gap resonances and lead to enhanced sway, roll and yaw motions. Reliable operability assessments therefore require a theoretical foundation that captures:

i) Linear wave kinematics and spectral representations of realistic sea states;

ii) First-order (wave-frequency) and second-order (difference- and sum-frequency) hydrodynamic loads;

iii) Governing equations of motion in both frequency and time domains, including hydrostatics, added mass, and radiation damping with memory effects;

iv) Multi-body diffraction and radiation interactions that dominate the narrow-gap behaviour;

v) Nonlinear mechanical behaviour of moorings and fenders that couple the two bodies.

Scope and assumptions: The theoretical framework in this chapter is based on MARIN's aNySIM Theory Documentation [97] and the formulations of Journée and Massie [98]. To avoid repetition, these two sources will not be cited again in each subsection; additional citations are introduced only when new theory, data or assumptions beyond these baselines are used. Unless stated otherwise, the flow is assumed inviscid, incompressible and irrotational, enabling potential-flow modelling of wave–body interactions together with linear wave theory.

Structure of this chapter: This chapter introduces the environmental models for waves, wind and current, the hydromechanic load components,including hydrostatics, first- and second-order wave loads, radiation with memory effects, and viscous damping. Followed by the nonlinear mechanical models for mooring lines and fenders. These components are assembled into the single- and multi-body equations of motion used in aNySIM, concluding with the distinction between wave-frequency and low-frequency motions relevant for side-by-side operability.

Relevance for operability analysis: The theoretical formulations in this chapter serve two essential purposes: (i) they define the physics resolved in MARIN's aNySIM simulations that produce the training, validation and test data for the predictive models developed in this thesis; and (ii) they provide the interpretative framework needed to understand and assess the predicted motions, relative responses and limiting parameters. Together, this chapter establishes the physical and mathematical foundations on which both the simulation-based analyses and the data-driven modelling approach in the remainder of the thesis rely.

B.1. Environmental Modelling

This section outlines the environmental inputs driving the hydrodynamic response of the FLNG–LNGC system. Linear wave kinematics and the JONSWAP spectral representation are introduced, followed by the reconstruction of irregular wave time series for use in diffraction–radiation and time-domain simulations. Steady wind and current fields are then defined, completing the environmental inputs applied in aNySIM.

B.1.1. Linear Wave Modelling

Irregular seas are represented as linear superpositions of monochromatic plane waves. Each wave component is described at a field point $x = (x, y, z)$ and time t by amplitude a_i, wave vector $k_i = (k_{ix}, k_{iy})$, angular frequency ω_i, and phase α_i. Introducing the dispersion relation, which becomes

$$\omega_i^2 = g\, k_i\, \tanh(k_i h), \tag{B.1}$$

and applies to deep and intermediate water depths.

The free-surface elevation associated with a component is

$$\zeta_i(x,t) = a_i \cos\phi_i(x,t), \tag{B.2}$$

and the corresponding velocity potential Φ_i, pressure p_i, aorbital velocity v_i, a_i follow directly from linear theory:

$$\Phi_i(x,t) = -\frac{g a_i}{\omega_i}\frac{\cosh[k_i(z+h)]}{\cosh(k_i h)}\sin\phi_i(x,t), \tag{B.3}$$

$$p_i(x,t) = \rho_w \frac{\partial \Phi_i}{\partial t} = \rho_w g a_i \frac{\cosh[k_i(z+h)]}{\cosh(k_i h)}\cos\phi_i(x,t), \tag{B.4}$$

$$v_i(x,t) = \nabla\Phi_i(x,t), \qquad a_i(x,t) = \frac{\partial}{\partial t}\nabla\Phi_i(x,t). \tag{B.5}$$

These expressions capture the linear orbital motion beneath waves and constitute the basis for calculating first-order wave excitation forces. Nonlinear steepness effects, wave breaking, and shallow-water nonlinearity are outside the applicability of this model.

B.1.1.1. Spectral description and discretisation

In the simulation workflow used by MARIN, irregular sea states are represented through the JONSWAP spectrum. It is established by Hassemanl et al. and it provides a realistic description of developing, fetch-limited seas due to its pronounced and narrow spectral peak spectrum [109]. The characteristic steep spectral peak results from the peak-enhancement factor γ, which amplifies the dominant frequency band and concentrates more energy around the peak period than equilibrium spectra.

The JONSWAP spectral density function is given by

$$S(\omega) = \alpha g^2\, \omega^{-5} \exp\left[-\frac{5}{4}\left(\frac{\omega_p}{\omega}\right)^4\right] \gamma^{\exp\left[-\frac{(\omega-\omega_p)^2}{2\sigma^2\omega_p^2}\right]}, \tag{B.6}$$

where $S(\omega)$ is the energy density at angular frequency ω, ω_p is the peak frequency, α is a scaling coefficient, σ is the spectral width parameter (0.07 for $\omega < \omega_p$ and 0.09 for $\omega > \omega_p$), and γ (typically 3.3) controls the sharpness of the peak. The mean JONSWAP spectrum obtained by Hasselman et al. [109] is displayed in Figure B.1, where the characteristic steep spectral peak observed.

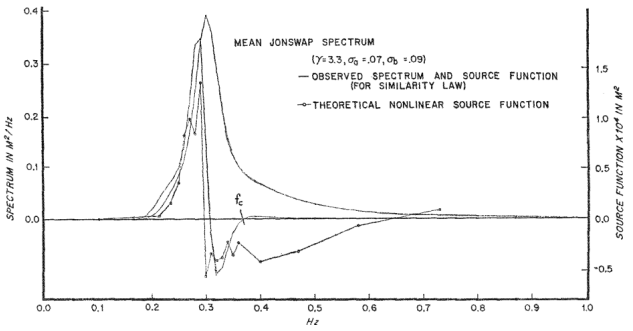


Figure B.1: Mean JONSWAP spectrum [109]

The first part of the expression describes the general decay of energy with frequency, while the second part introduces the peak enhancement essential for modelling fetch-limited conditions.

This spectral formulation is discretised into a finite set of wave components $\{\omega_i, a_i, \alpha_i, \mu_i\}$ such that the target significant wave height is preserved:

$$H_s^2 = 8 \sum_{i=1}^{n} a_i^2. \tag{B.7}$$

The resulting component set forms the input for MARIN's diffraction–radiation and time-domain hydrodynamic models.

B.1.1.2. Superposition and time series

By linear superposition, the irregular free-surface elevation and associated kinematics are obtained by summing all discretised wave components:

$$\zeta(x,t) = \sum_{i=1}^{n} \zeta_i(x,t), \qquad p(x,t) = \sum_{i=1}^{n} p_i(x,t), \tag{B.8}$$

$$v(x,t) = \sum_{i=1}^{n} v_i(x,t), \qquad a(x,t) = \sum_{i=1}^{n} a_i(x,t). \tag{B.9}$$

This reconstruction defines the irregular wave field used as input for diffraction–radiation and time-domain hydrodynamic models.

For time-domain applications, a deterministic but statistically equivalent wave record is generated from the spectrum. The surface elevation is formed as

$$\zeta(t) = \sum_{i=1}^{n} a_i \cos(k_i x - \omega_i t + \varepsilon_i)\,, \tag{B.10}$$

where amplitudes follow

$$a_i^2 = 2\, S_\zeta(\omega_i)\, \Delta\omega, \tag{B.11}$$

ensuring that the regenerated time series reproduces the target energy density spectrum. Random phases $\varepsilon_i \in [0, 2\pi)$ are selected independently, creating a new realisation with identical statistical properties. Because the frequency spacing $\Delta\omega$ is constant, the signal repeats after a period $2\pi/\Delta\omega$.

This approach provides a statistically consistent time-domain wave field suitable for use in nonlinear and multi-body simulations. Directional spreading can also be introduced by subdividing each frequency component into several directional components, although this option is not applied in the present MARIN & Shell workflow. The overall procedure, from obtaining the energy density spectrum, assigning amplitudes based on spectral energy, selecting random phases, and reconstructing a new irregular wave record, is illustrated in Figure B.2.

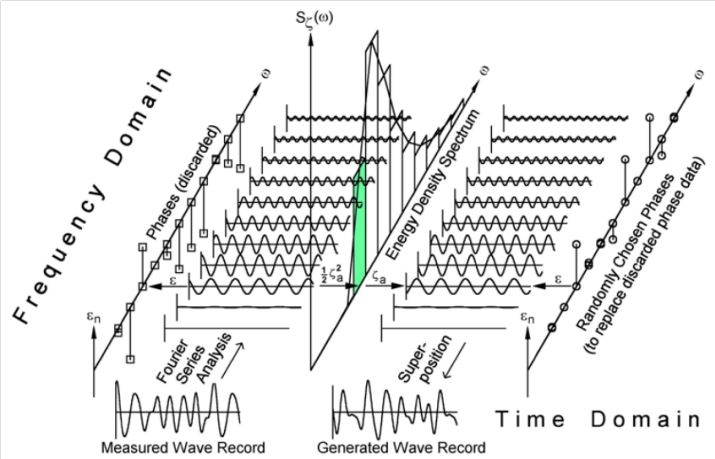
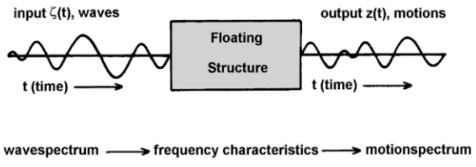
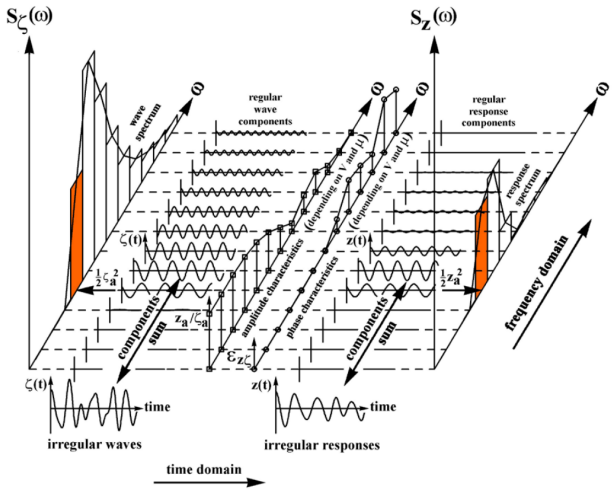


Figure B.2: Wave record analysis and regeneration: transformation between frequency-domain spectrum and time-domain realisation [98].

Linear input–output behaviour: Because the governing equations and boundary conditions are linearised, the hydrodynamic system behaves as a linear filter. Each monochromatic wave component produces a response of the same frequency, and the total response is the superposition of all components. This behaviour is illustrated in Figure B.3, where the wave elevation $\zeta(t)$ is the input and the vessel motions $z(t)$ form the output.



(a) Conceptual linear input–output relationship.



(b) Spectral decomposition of waves and motion response.

Figure B.3: Two complementary representations of linear wave-to-motion behaviour [98]

B.1.2. Wind and current fields

In the aNySIM modelling, environmental loading includes wind and current fields in addition to waves. These fields are treated independently from the wave model and act as steady external forces in the global coordinate system. Although aNySIM supports more advanced representations, these options are not applied in the present MARIN & Shell workflow. Instead, both wind and current are prescribed as constant, uniform fields, providing steady aerodynamic and hydrodynamic load contributions. These load contributions will be described in subsection B.2.6 and subsection B.2.7.

B.2. Hydromechanic Load Models

Side-by-side moored floating structures, such as FLNG–LNGC systems, experience complex hydrodynamic interactions driven by waves, radiation effects, hydrostatics, viscous damping, and mooring–fender forces. In aNySIM, these mechanisms are represented through modular linear and nonlinear load components, assembled in the global equations of motion for all floating bodies. This section summarises the underlying theory relevant for the hydromechanic loading in side-by-side offloading simulations.

B.2.1. Linear Hydrostatics

Hydrostatic restoring forces follow from the balance between weight and buoyancy. In equilibrium, the gravitational force acts through the centre of gravity G, while the buoyant force, equal to the displaced water weight, acts through the centre of buoyancy B. his shift generates restoring forces and moments, which may be linearised for sufficiently small motions.

Assumptions: The linear hydrostatics model assumes small heave, roll, and pitch motions; negligible influence of waves on hydrostatic properties; a constant waterplane shape; and that the hull intersects the waterline approximately vertically. All rotations are taken about the centroid of the waterplane area to avoid introducing artificial heave coupling.

Heave restoring: A small vertical displacement z modifies the submerged volume by

$$\mathrm{d}V = A_{\mathrm{wp}}\, z,$$

where A_{wp} is the waterplane area. The change in buoyancy force is

$$F_z = -\rho_w g A_{\mathrm{wp}}\, z,$$

giving the linear heave stiffness

$$k_z = \rho_w g A_{\mathrm{wp}}. \tag{B.12}$$

This remains accurate as long as A_{wp} is approximately constant over the motion amplitude.

Roll and pitch restoring: A small rotation about the waterplane centroid produces two geometric changes: (i) a horizontal shift of the original buoyancy centre, and (ii) an additional shift because emerged and immersed wedges are generated on opposite sides of the waterplane. Together with the corresponding horizontal shift of G, these effects create a righting arm that is linear in the rotation angle. Figure B.4 illustrates the formation of the wedges, the shift of B to B_ϕ, and the definition of the metacentre.

For a small roll angle ϕ, the wedge-induced buoyancy shift is

$$\Delta x = \frac{I_{xx}}{V}\,\phi, \qquad \Delta y = \frac{I_{yy}}{V}\,\phi,$$

where I_{xx}, I_{yy} are the waterplane second moments of area and V is the displaced volume. Combining all contributions yields the linear restoring stiffnesses

$$k_{xx} = \rho_w g V \, GM_t, \tag{B.13}$$

$$k_{yy} = \rho_w g V \, GM_l, \tag{B.14}$$

with GM_t and GM_l the transverse and longitudinal metacentric heights.

Metacentric height: The metacentric height follows from

$$GM = BM - KG, \qquad BM = \frac{I_{\mathrm{wp}}}{V},$$

where I_{wp} is the waterplane second moment area and KG is the vertical coordinate of the centre of gravity. For small angles (typically $\phi \lesssim 10°$), the metacentric approximation is accurate; at larger angles, the vertical motion of B becomes relevant, and roll/pitch no longer decouple.

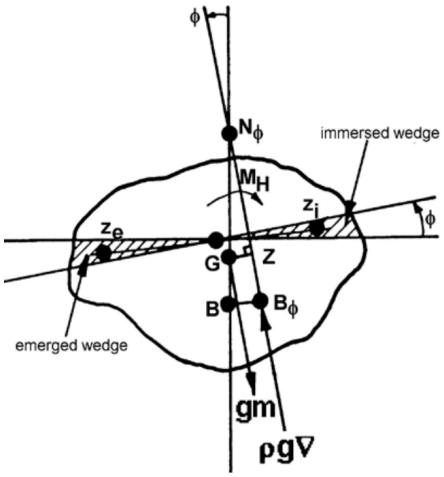


Figure B.4: Shift of the centre of buoyancy due to heel: formation of emerged and immersed wedges, movement of B to B_ϕ, and definition of the metacentre N_ϕ [98].

Restoring matrix and reference point: The hydrostatic stiffnesses derived above (k_z, k_{xx}, k_{yy}) are defined with respect to the centre of floatation, where the waterplane centroid acts as the natural reference point for small-amplitude hydrostatics. In this location the restoring matrix $\mathbf{K}_c$ is diagonal in the heave, roll, and pitch modes, because a pure vertical displacement or a pure small rotation does not induce additional geometric couplings.

For the equations of motion, however, hydrostatic forces and moments must act at the centre of gravity (CoG). Translating the restoring matrix from the waterplane centroid to the CoG requires a standard rigid-body transformation,

$$\mathbf{K}_g = \mathbf{T}\,\mathbf{K}_c\,\mathbf{T}^*,$$

where $\mathbf{T}$ contains the linearised geometric shifts $(\Delta x, \Delta y, \Delta z)$ between the two reference points. This translation introduces additional off-diagonal terms: a heave force at the waterline generates roll and pitch moments at the CoG if the CoG is not located exactly above the waterplane centroid, and conversely small rotations create vertical offsets that modify the effective restoring in heave.

The resulting coupling terms (proportional to $\Delta x\, k_z$ and $\Delta y\, k_z$) are purely geometric. They do not change the physical restoring mechanism, but they do affect how the hydrostatic forces appear in the global equations of motion.

B.2.2. First Order Wave Loads

Linear wave excitation represents the first-order hydrodynamic forces generated by incident waves acting on the hull. In linear potential-flow theory, the fluid is assumed inviscid, incompressible, and irrotational, such that the velocity field can be expressed as the gradient of a scalar potential $\Phi(x, y, z, t)$. This potential satisfies the Laplace equation,

$$\nabla^2\Phi = \frac{\partial^2\Phi}{\partial x^2} + \frac{\partial^2\Phi}{\partial y^2} + \frac{\partial^2\Phi}{\partial z^2} = 0, \tag{B.15}$$

subject to boundary conditions on the seabed, the free surface, and the vessel hull. The linearised free-surface condition couples the dynamic and kinematic effects of surface elevation,

$$\frac{\partial^2\Phi}{\partial t^2} + g\frac{\partial\Phi}{\partial z} = 0 \qquad \text{at } z = 0, \tag{B.16}$$

while the body boundary condition prescribes that the normal fluid velocity equals the normal velocity of the hull. Figure B.5 summarises the boundary-value problem for the present side-by-side setting.

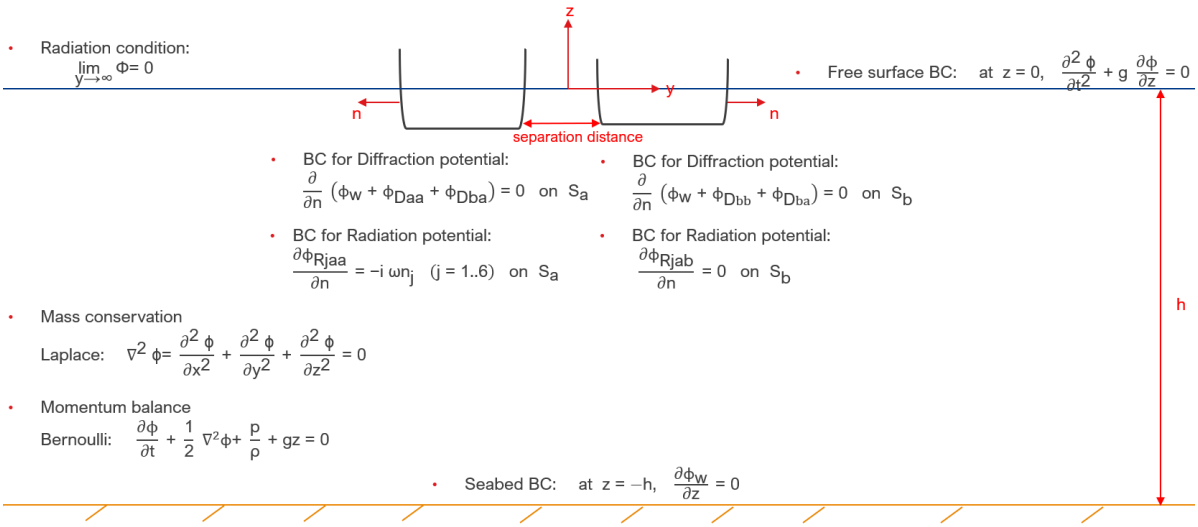


Figure B.5: Linear potential-flow boundary conditions for two interacting floating bodies.

Decomposition of the velocity potential: In linear potential flow theory, the total velocity potential is expressed as the superposition of the incident wave field, the diffraction field generated by the presence of the body, and the radiation fields associated with the six rigid-body motions:

$$\Phi = \sum_{j=1}^{6} \Phi_j + \Phi_w + \Phi_d. \tag{B.17}$$

Here, Φ_w denotes the incident wave potential, Φ_d the diffraction potential, and Φ_j the radiation potential due to motion in degree of freedom j.

This decomposition corresponds to two complementary load problems. First (i), with the body fixed in waves, the wave excitation loads consist of the Froude–Krylov contribution from the incident field Φ_w and the diffraction loads arising from the scattered field Φ_d. In irregular waves, the incident field is represented as a linear superposition of many wave components of different frequencies and directions, and the corresponding excitation loads are obtained through the vessel's frequency–directional response operators.

Second (ii), in calm water ($\Phi_w = 0$), forced oscillations of the body generate the radiation potentials Φ_j, which lead to added-mass forces (fluid inertia), radiation damping associated with wave energy radiated away from the vessel, and the hydrostatic restoring forces due to weight and buoyancy. In irregular seas, these radiation effects are applied consistently across all frequencies and enter the equations of motion through the frequency- dependent added-mass and damping coefficients or their equivalent time-domain convolution representation.

For two vessels a and b in close proximity, the decomposition on each hull contains both self and cross contributions, which is the source of hydrodynamic interaction:

$$\Phi_a = \Phi_w + \Phi_{d,aa} + \Phi_{d,ba} + \sum_{j=1}^{6} (\Phi_{j,aa} + \Phi_{j,ba}) \,, \tag{B.18}$$

$$\Phi_b = \Phi_w + \Phi_{d,bb} + \Phi_{d,ab} + \sum_{j=1}^{6} (\Phi_{j,bb} + \Phi_{j,ab}) \,. \tag{B.19}$$

Equations (B.18)–(B.19) show that each vessel is subject to both its own and the neighbour's diffraction fields, as well as radiation generated by the motions of both bodies, leading to off–diagonal hydrodynamic coupling.

First-order wave loads: The first-order hydrodynamic pressure follows from the linearised Bernoulli equation,

$$p = -\rho_w \frac{\partial \Phi}{\partial t} - \rho_w g z. \tag{B.20}$$

Integrating this pressure over the instantaneous wetted surface S yields the first-order force,

$$F^{(1)}(t) = -\rho_w \iint_S \frac{\partial \Phi}{\partial t} \, n \, \mathrm{d}S. \tag{B.21}$$

This force consists of two physically meaningful contributions:

- **Froude–Krylov force** — the pressure from the undisturbed incoming wave field acting directly on the hull,

$$F_{\mathrm{FK}} = -\rho_w \iint_S \frac{\partial \Phi_w}{\partial t} \, n \, \mathrm{d}S, \tag{B.22}$$

With the full incident wave potential for irregular waves being:

$$\Phi_w(x, y, z, t) = \sum_n \frac{a_n g}{\omega_n} \frac{\cosh[k_n(z+h)]}{\cosh(k_n h)} \cos(k_n(x \cos\beta_n + y \sin\beta_n) - \omega_n t + \varepsilon_n) \,. \tag{B.23}$$

- **Diffraction force** — the additional pressure generated by the scattered wave field,

$$F_{\mathrm{diff}} = -\rho_w \iint_S \frac{\partial \Phi_d}{\partial t} \, n \, \mathrm{d}S. \tag{B.24}$$

The diffraction potential Φ_d represents the wave field scattered by a fixed body when it is exposed to an incident harmonic wave of frequency ω. It is defined as

$$\Phi_d(\mathbf{x}, t) = \Re\{ \, \phi_d(\mathbf{x}) \, e^{-i\omega t} \} \,, \tag{B.25}$$

where $\phi_d(\mathbf{x})$ is the complex spatial diffraction potential satisfying the Laplace equation together with the body boundary condition $\partial(\phi_w + \phi_d)/\partial n = 0$, the linear free-surface condition, and the outgoing-wave radiation condition.

The total first-order excitation force is therefore

$$F^{(1)}(t) = F_{\mathrm{FK}}(t) + F_{\mathrm{diff}}(t). \tag{B.26}$$

B.2.3. Radiation Forces and Added Mass

When a floating structure oscillates in waves, its motion generates radiated waves that propagate away from the hull. These radiated waves carry energy and produce reaction loads collectively known as radiation forces. Radiation forces appear in two components:

- an added-mass force, proportional to the body acceleration, arising from the inertia of the water that must be accelerated together with the hull;
- a radiation-damping force, proportional to the velocity, representing the energy carried away by the radiated wave field.

In linear potential-flow theory, these forces are expressed using the Cummins equation [121], which separates instantaneous inertial effects from memory-dependent damping:

$$F_k^{\mathrm{rad}}(t) = -\sum_j A_{kj}^{\infty}\,\ddot{x}_j(t) - \sum_j \int_0^t K_{kj}(\tau)\,\dot{x}_j(t-\tau)\,\mathrm{d}\tau, \tag{B.27}$$

where A_{kj}^{∞} is the infinite-frequency added mass and $K_{kj}(t)$ is the retardation kernel, wich is related to the frequency-dependent radiation damping $B_{kj}(\omega)$ through the cosine transform [122]:

$$K_{kj}(\tau) = \frac{2}{\pi}\int_0^{\infty} B_{kj}(\omega)\cos(\omega\tau)\,\mathrm{d}\omega, \qquad \tau \geq 0, \tag{B.28}$$

while A_{kj}^{∞} follows from the high-frequency limit of the added-mass spectrum.

Multi-body radiation coupling: For side-by-side moored vessels, the radiation forces must be written in block-matrix form because the motion of one hull generates radiated waves that act on the other. The added-mass and retardation operators therefore become

$$\mathbf{A}^{\infty} = \begin{bmatrix} \mathbf{A}_{11}^{\infty} & \mathbf{A}_{12}^{\infty} \\ \mathbf{A}_{21}^{\infty} & \mathbf{A}_{22}^{\infty} \end{bmatrix}, \qquad \mathbf{K}(t) = \begin{bmatrix} \mathbf{K}_{11}(t) & \mathbf{K}_{12}(t) \\ \mathbf{K}_{21}(t) & \mathbf{K}_{22}(t) \end{bmatrix},$$

where the off-diagonal blocks represent radiation cross-coupling between the two bodies.
Each vessel has six degrees of freedom, so every block $\mathbf{A}_{ij}^{\infty}$ and $\mathbf{K}_{ij}(t)$ is a 6×6 matrix containing 36 hydrodynamic coefficients.
For a two-body system, the full matrices $\mathbf{A}^{\infty}$ and $\mathbf{K}(t)$ therefore contain 12×12 entries, of which 72 belong to the cross-coupling blocks $\mathbf{A}_{12}^{\infty}$ and $\mathbf{A}_{21}^{\infty}$ (and similarly for $\mathbf{K}_{12}(t)$ and $\mathbf{K}_{21}(t)$).
These off-diagonal terms become significant in narrow-gap configurations and must be retained for accurate motion prediction. Model-test results show that leaving out these coupling terms causes the wave-frequency heave and pitch motions to be overestimated [76].

Hydrodynamic modelling considerations: In MARIN's workflow, the frequency-dependent hydrodynamic coefficients $A_{kj}(\omega)$ and $B_{kj}(\omega)$ are obtained from a multi-body radiation–diffraction computation in DIFFRAC, with both hulls included in the boundary-value problem [75]. This ensures that wave shielding, interference, and gap-flow effects are captured in the hydrodynamic matrices.
The damping matrix is then transformed into the memory kernel using Equation B.28, while A_{kj}^{∞} is taken from the high-frequency limit of the added-mass spectrum. In the gap between the two vessels, potential-flow solvers may predict unrealistic oscillations and drift forces; therefore, a free-surface lid technique is applied to stabilize the diffraction solution and produce physically consistent hydrodynamic coefficients [75], which form the basis of the time-domain radiation model used in Eq. (B.27).

B.2.4. Viscous Damping

Linear potential-flow radiation accounts for wave-making energy losses but does not represent viscous mechanisms. In side-by-side arrangements, the narrow gap between the vessels introduces dissipative effects that are absent in isolated-vessel hydrodynamics [76, 75]:

- flow separation at the gap entrance,
- vortex shedding within the gap,
- oscillatory jet-like exchange driven by pressure differences,
- nonlinear damping under large sway and yaw amplitudes.

These processes produce substantially higher damping—particularly in sway and yaw—than predicted by radiation alone. Model tests confirm that the presence of a neighbouring hull increases horizontal-plane damping significantly [76].

Two-body viscous representation: For two vessels, viscous effects are introduced through additional linear or quadratic damping acting on each hull and on their relative motion. A compact form is

$$\mathbf{F}_d^{(1)} = -\mathbf{B}_{11}\,\mathbf{v}^{(1)} - \mathbf{B}_{\mathrm{rel}}\big(\mathbf{v}^{(1)} - \mathbf{v}^{(2)}\big), \qquad \mathbf{F}_d^{(2)} = -\mathbf{B}_{22}\,\mathbf{v}^{(2)} - \mathbf{B}_{\mathrm{rel}}\big(\mathbf{v}^{(2)} - \mathbf{v}^{(1)}\big),$$

where $\mathbf{v}^{(i)}$ is the 6-DOF velocity of body i, $\mathbf{B}_{11}$ and $\mathbf{B}_{22}$ represent single-body viscous damping, and $\mathbf{B}_{\mathrm{rel}}$ models gap-induced losses. The dominant entries of $\mathbf{B}_{\mathrm{rel}}$ lie in sway and yaw, consistent with decay-test observations [76]. For larger motions, quadratic damping terms capture vortex shedding, where the gap flow breaks into swirling vortices that add extra energy loss.

Identification from experiments: The relative damping coefficients are determined experimentally by MARIN by comparing decay tests of a single vessel with tests where the neighbouring hull is held fixed. The difference in damping between the two scenarios gives the effective gap-induced coefficients in sway and yaw [76]. These empirically derived values are then introduced in the viscous-damping matrices of the simulation model.

B.2.5. Second-Order (Drift) Wave Loads

Second-order (drift) wave loads are the main cause of slow surge, sway, and yaw motions in moored or side-by-side floating structures. These slow-drift forces come from the nonlinear interaction of pairs of regular wave components, which creates difference-frequency loads that align with the natural low-frequency motions of the system. Because mooring systems usually have low stiffness and only limited damping at these frequencies, these slowly changing forces tend to dominate the low-frequency behaviour.

Figure B.6 illustrates this behaviour. The upper curve shows the irregular wave elevation, while the shaded lower curve displays the corresponding second-order drift force. A clear decomposition into a constant (mean) component and a slowly varying component is visible, with the latter following the modulation of the wave groups.

Origin of Second-Order Loads: The formulation of second-order forces begins with the complete Bernoulli equation,

$$p = -\rho\frac{\partial\Phi}{\partial t} - \frac{1}{2}\rho|\nabla\Phi|^2 - \rho g z, \tag{B.29}$$

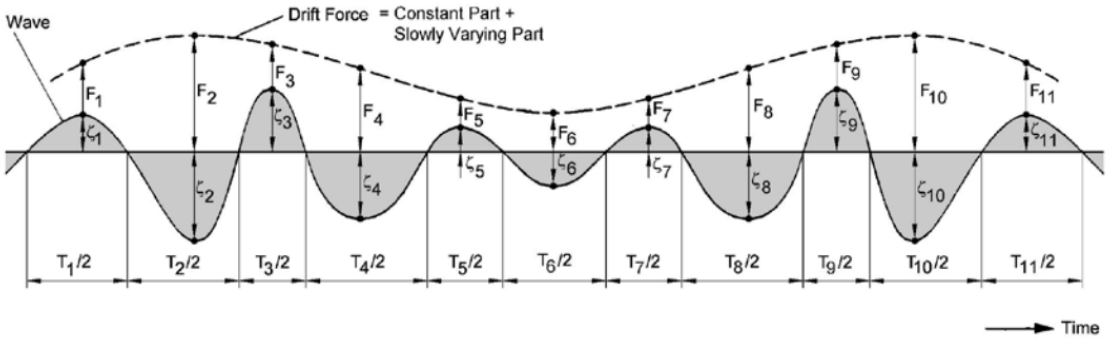


Figure B.6: Time-domain representation of an irregular wave elevation (bottom) and the associated second-order drift force (top). The drift force consists of a constant and a slowly varying part that follows wave groupiness.

which is expanded using a second-order approximation about the mean free surface following the approach of Pinkster [111]. The quadratic velocity term produces mean, sum-frequency, and difference-frequency components.

For two interacting wave components m and n, the horizontal velocity can be expressed as

$$U_x^{(m,n)}(t) = \zeta_m \cos(\omega_m t + \epsilon_m) + \zeta_n \cos(\omega_n t + \epsilon_n), \tag{B.30}$$

leading to the quadratic velocity contribution

$$\begin{aligned}\frac{\rho}{2}U_x^2 = \frac{\rho}{2}\Big[&p;\zeta_m^2 \cos(2\omega_m t + 2\epsilon_m) + \zeta_n^2 \cos(2\omega_n t + 2\epsilon_n) \\ &p;+2\zeta_m\zeta_n \cos\big((\omega_m - \omega_n)t + (\epsilon_m - \epsilon_n)\big) \\ &p;+2\zeta_m\zeta_n \cos\big((\omega_m + \omega_n)t + (\epsilon_m + \epsilon_n)\big)\Big].\end{aligned} \tag{B.31}$$

Equation (B.31) reveals the fundamental structure of the second-order loads:

- a mean drift component arising from the time-averaged quadratic velocity;
- a difference-frequency component $|\omega_m - \omega_n|$, responsible for slow-drift motions;
- a sum-frequency component $(\omega_m + \omega_n)$, primarily relevant for high-frequency responses.

The slowly varying drift force in Fig. B.6 is directly associated with the difference-frequency term.

Quadratic Transfer Function (QTF) Formulation: The second-order drift forces can be compactly represented in the frequency domain using Quadratic Transfer Functions (QTFs). The second-order force in degree of freedom k is written as

$$F_k^{(2)} = \sum_{m=1}^{M}\sum_{n=1}^{M} QTF_k(\omega_m, \omega_n, \mu_m, \mu_n)\, \bar{a}_m \bar{a}_n, \tag{B.32}$$

where $\bar{a}_m$ and $\bar{a}_n$ are the complex wave amplitudes and QTF_k is the quadratic transfer function obtained from the second-order boundary-value problem. The QTF contains both sum- and difference-frequency contributions:

$$\omega_+ = \omega_m + \omega_n, \tag{B.33}$$

$$\omega_- = |\omega_m - \omega_n|. \tag{B.34}$$

Relation to the Near-Field Method: The near-field method provides a practical way to evaluate the second-order pressure acting on a floating body by integrating the nonlinear Bernoulli pressure over the mean wetted surface, the oscillating wetted strip, and the waterline. In the formulation presented by Pinkster [111], this decomposition yields the complete second-order load, including both the mean drift component and the low-frequency part associated with difference-frequency interactions.

DIFFRAC employs this near-field framework by computing the nonlinear pressure distribution on the hull and integrating it to obtain the second-order excitation forces. The resulting frequency dependence of these loads is then expressed in the form of a QTF, corresponding to the components in eq. (B.32).

Following Pinkster's formulation, the second-order force can be written as

$$\mathbf{F}^{(2)} = \iint_{S_0} p^{(2)} \mathbf{n} \, \mathrm{d}S + \iint_{S_1} p^{(1)} \mathbf{n} \, \mathrm{d}S + \oint_{\Gamma_1} \frac{1}{2} \rho g \zeta^2 \, \mathbf{n} \, \mathrm{d}\ell, \tag{B.35}$$

where the three terms represent the contributions from the mean wetted surface, the oscillating wetted strip, and the waterline, respectively.

Connection to the Time-Domain Representation of Drift Forces: The developments in eqs. (B.29)–(B.31) establish the physical origin of the second-order loads: the quadratic velocity product generates mean, sum-frequency, and difference-frequency contributions. The QTF formulation in eq. (B.32) provides a compact and efficient frequency-domain representation of these loads, while the near-field method offers a physically transparent evaluation of the corresponding nonlinear pressures.

To obtain the explicit time-domain forcing used in simulation models, the complex QTF contributions are decomposed into their in-phase and quadrature components. For each wave-pair (i, j),

$$QTF_k(\omega_i, \omega_j) \, \bar{a}_i \bar{a}_j = P_{ij} - i \, Q_{ij}, \tag{B.36}$$

such that P_{ij} and Q_{ij} represent the real (cosine) and imaginary (sine) parts of the associated difference-frequency force. Substituting this into the inverse representation of the second-order load yields

$$F_k^{(2)}(t) = \sum_{i=1}^{M} \sum_{j=1}^{M} \left[P_{ij} \, \cos\big((\omega_i - \omega_j)t\big) + Q_{ij} \, \sin\big((\omega_i - \omega_j)t\big) \right], \tag{B.37}$$

which corresponds to the slowly varying component illustrated in Fig. B.6. The frequencies $(\omega_i - \omega_j)$ are the same difference-frequency terms identified in eq. (B.31), while P_{ij} and Q_{ij} follow directly from the real and imaginary parts of the QTF.

In summary, the quadratic expansion from the Bernoulli equation establishes the origin of second-order loads; the near-field method provides the physical pressure distribution responsible for these loads; the QTF formulation captures the frequency dependence into a compact operator; and the P_{ij}–Q_{ij} representation gives the explicit time-domain forcing used in low-frequency simulations of moored and side-by-side floating systems.

B.2.6. Current Forces

Current forces are computed in aNySIM using a coefficient-based formulation in which the flow-induced loads depend on the relative current velocity and the angle of attack of the flow. The model assumes a horizontally uniform current. The load calculation uses the relative velocity between the vessel and the current.

The hydrodynamic current load components in surge, sway and yaw are expressed as

$$X_{\text{current}} = \tfrac{1}{2} \rho_w V_{\text{cr}}^2 A \, C_{cx}(\alpha_{\text{cr}}), \quad Y_{\text{current}} = \tfrac{1}{2} \rho_w V_{\text{cr}}^2 A \, C_{cy}(\alpha_{\text{cr}}), \quad N_{\text{current}} = \tfrac{1}{2} \rho_w V_{\text{cr}}^2 A L_{pp} \, C_{cn}(\alpha_{\text{cr}}).$$

The coefficients C_{cx}, C_{cy} and C_{cn} capture how the underwater hull generates forces for a given inflow direction. For small angles ($\alpha_{\mathrm{cr}} \approx 0^\circ$ or 180°), the longitudinal coefficient is dominated by the bow or stern shape and produces primarily surge drag. Around beam-on conditions ($\alpha_{\mathrm{cr}} \approx \pm 90^\circ$), the side area of the vessel becomes dominant, resulting in significant sway forces and yaw moments. The yaw coefficient reflects how asymmetric pressure and flow separation patterns along the hull generate a turning tendency.

B.2.7. Wind Forces

Wind loading is represented using a coefficient-based aerodynamic model in which the wind velocity at 10 m elevation is combined with windage areas and dimensionless wind coefficients. The wind field is assumed uniform in space, and the instantaneous relative wind direction determines which coefficient values are applied.

The wind-induced surge, sway and yaw components are given by

$$X_{\mathsf{wind}} = \tfrac{1}{2}\rho_a V_{\mathrm{wr}}^2 A\, C_{wx}(\alpha_{\mathrm{wr}}), \quad Y_{\mathsf{wind}} = \tfrac{1}{2}\rho_a V_{\mathrm{wr}}^2 A\, C_{wy}(\alpha_{\mathrm{wr}}), \quad N_{\mathsf{wind}} = \tfrac{1}{2}\rho_a V_{\mathrm{wr}}^2 A L_{pp}\, C_{wn}(\alpha_{\mathrm{wr}}).$$

The aerodynamic coefficients follow the same directional structure as the current coefficients but are governed by the vessel's above-water geometry. When the wind approaches from the bow or stern, C_{wx} is governed by the frontal or aft windage area of the hull and superstructure. Beam-on wind conditions produce the largest C_{wy} and C_{wn} values due to the large projected side area, accommodation block and funnel, which generate substantial lateral force and yawing moment.

In aNySIM the tabulated coefficients are interpolated based on the instantaneous wind angle α_{wr}, and the resulting forces and moments are applied at a representative aerodynamic reference point or distributed over segments when wind segmentation is enabled. For the configuration used in this study, wind acts as a steady environmental load influencing the horizontal-plane offset and heading alignment of the coupled system.

B.3. Mechanical Forces: Mooring Lines and Fenders

Thee physically different mechanical interaction influence the coupled FLNG–LNGC dynamics:

1. The anchoring mooring system of the FLNG (spread mooring), modelled as a linear springs.
2. The side-by-side mooring lines between the FLNG and the LNG carrier, modelled as fully nonlinear mechanical elements.
3. The fenders, which act as nonlinear contact elements providing normal reaction, damping, and friction forces when the vessels come into compression.

The mechanical interaction of all three are explained in more detail in the following sections.

B.3.1. FLNG Anchoring Mooring System (Linear Model)

The spread-mooring system of the FLNG is represented using a linear stiffness matrix. Its purpose is to provide low-frequency global restoring forces in surge, sway and yaw, without modelling the detailed behaviour of individual lines.

The restoring force is computed as

$$\mathbf{F}_{\text{anchor}} = \begin{bmatrix} k_{11} & 0 & 0 \\ 0 & k_{22} & 0 \\ 0 & 0 & k_{66} \end{bmatrix} \begin{bmatrix} x \\ y \\ \psi \end{bmatrix}, \tag{B.38}$$

where x and y are the horizontal offsets and ψ is the yaw angle. This formulation captures the anchoring influence on global station keeping but does not include nonlinear axial behaviour, or dynamic line mechanics.

The anchoring system therefore acts as:

- a linear spring in surge and sway,
- a linear yaw stiffness,

This global stiffness is added directly to the FLNG equations of motion.

Catenary Geometry

The static shape of the mooring line follows a catenary determined by weight and horizontal tension. For a distributed weight w and horizontal tension component H, the vertical displacement is

$$z(s) = \frac{H}{w}\left[\cosh\left(\frac{w}{H}s\right) - 1\right]. \tag{B.39}$$

This provides the geometric stiffness and the initial pretension before dynamic effects are considered.

B.3.2. Side-by-Side Mooring Lines (Nonlinear and Dynamic Model)

The SBS mooring system between the FLNG and LNGC is modelled in full detail. Each mooring line contributes forces through:

1. nonlinear axial stiffness from load–elongation tables,
2. geometric stiffness from catenary equilibrium,
3. optional dynamic behaviour based on lumped-mass physics.

Nonlinear Axial Behaviour

The line tension is determined from the elongation

$$\Delta L = L - L_0, \tag{B.40}$$

and the axial force is

$$F_{\text{line}} = f(\Delta L), \tag{B.41}$$

where $f(\cdot)$ is obtained directly from the nonlinear load–elongation curves.

The axial line force vector acting on the bodies is

$$\mathbf{F}_{\text{line}} = F_{\text{line}}\,\mathbf{t}, \tag{B.42}$$

with $\mathbf{t}$ the unit vector along the line.

B.3.3. Fender Mechanics

Fenders are represented as mechanical contact elements that generate forces only when the gap between the two vessels closes. The model detects compression along a prescribed normal direction. Let p_1 and p_2 be the earth-fixed positions of the two contact points, and n the outward normal at the contact on body 1. The projected gap distance is

$$d = (p_2 - p_1) \cdot n, \tag{B.43}$$

and the instantaneous fender compression is

$$\Delta h = h_0 - d, \qquad \text{contact if } \Delta h \geq 0, \tag{B.44}$$

where h_0 is the nominal stand-off distance (fender diameter or height).

The normal reaction force is obtained from non-linear compression curve:

$$F_n = f_{\text{fender}}(\Delta h). \tag{B.45}$$

A tangential friction force acts in the local contact plane. Let $\dot{p}_{\text{rel}}$ be the relative velocity of the contact points and let t be the unit tangent direction in the plane. The friction force is

$$F_t = -\,\mu\, F_n\ \text{sgn}(t \cdot \dot{p}_{\text{rel}}), \tag{B.46}$$

with μ a friction coefficient.

The total fender force acting on body 1 is

$$F_{\text{fender}} = F_n\, n + F_t\, t, \tag{B.47}$$

and body 2 receives the equal and opposite force.

This formulation captures the essential behaviour of pneumatic fenders in side-by-side operations: non-linear compression resistance in the normal direction combined with frictional resistance to relative sliding in the tangential direction.

B.4. Rigid-Body Dynamics Framework

The hydromechanical load components introduced in the previous sections act on each vessel through its six rigid-body degrees of freedom (6-DOF). This section summarises how these loads are assembled into the time-domain equations of motion used by `aNySIM` for multi-body simulation of side-by-side offloading operations. The formulation follows classical Newton–Euler rigid-body dynamics, extended with linear diffraction–radiation hydrodynamics and the retardation-function representation of memory effects.

B.4.1. Six-Degree-of-Freedom Vessel Dynamics

Each vessel is treated as a rigid body with configuration vector

$$\boldsymbol{\eta} = (x,\, y,\, z,\, \phi,\, \theta,\, \psi)^T,$$

expressed in the earth-fixed frame introduced in Section 3.2.2. The corresponding body-fixed velocity vector is

$$\mathbf{v} = (u,\, v,\, w,\, p,\, q,\, r)^T.$$

Inertia: The rigid-body inertia matrix $\mathbf{M}_{\mathrm{RB}}$ contains the vessel mass and its rotational inertia tensor. Hydrodynamic radiation causes additional inertia, captured by the infinite-frequency added-mass matrix $\mathbf{A}_\infty$ obtained from the multi-body DIFFRAC analysis. The total (inertial) mass matrix used in the equations of motion is therefore

$$\mathbf{M}_\infty = \mathbf{M}_{\mathrm{RB}} + \mathbf{A}_\infty.$$

This matrix reflects both the dry rigid-body inertia and the fluid inertia associated with accelerating the surrounding water.

Coriolis terms: The matrices $\mathbf{C}_{\mathrm{RB}}$ and $\mathbf{C}_A$ describe how rotation couples with translation when the equations are written in a rotating body-fixed frame. They represent gyroscopic effects: the vessel's velocity creates additional apparent forces when the vessel is turning. These terms are formed automatically inside aNySIM and do not require user input.

Radiation memory: Wave radiation leads to energy losses and added inertia that depend on the motion history. In the time domain this effect is represented by a convolution integral involving the radiation retardation kernel $\mathbf{R}(t)$, derived from the frequency-dependent radiation damping coefficients via an inverse cosine transform. This term accounts for the fact that a vessel moving at time t continues to experience hydrodynamic resistance due to waves radiated at earlier times.

Hydrostatics: Changes in buoyancy and weight resulting from small vessel displacements are represented by the linear hydrostatic stiffness matrix $\mathbf{K}$. This term provides the restoring forces and moments that act to return the vessel toward its equilibrium position.

External loads: All other contributions are collected in the total external load vector $\mathbf{F}_{\mathsf{glob}}(t)$, which consists of the first-order wave excitation forces $\mathbf{F}^{(1)}$, the second-order (slow-drift) loads $\mathbf{F}^{(2)}$, the viscous damping forces $\mathbf{F}^{\mathrm{visc}}$, the environmental loads from wind and current $\mathbf{F}^{\mathrm{env}}$, and the mechanical loads transmitted through the mooring lines and fenders $\mathbf{F}^{\mathrm{mech}}$.

Collecting all contributions yields the 6-DOF equation of motion for a single vessel:

$$\underbrace{\mathbf{M}_{\infty}\,\dot{\mathbf{v}}(t)}_{\text{inertia}} + \underbrace{\mathbf{C}_{\mathrm{RB}}(\mathbf{v})\,\mathbf{v}(t) + \mathbf{C}_{A}(\mathbf{v})\,\mathbf{v}(t)}_{\text{Coriolis}} + \underbrace{\int_{0}^{t} \mathbf{R}(t-\tau)\,\mathbf{v}(\tau)\,\mathrm{d}\tau}_{\text{radiation memory}} + \underbrace{\mathbf{K}\,\boldsymbol{\eta}(t)}_{\text{hydrostatic restoring}} = \underbrace{\mathbf{F}_{\text{glob}}(t)}_{\text{external loads}}\,. \tag{B.48}$$

Equation (B.48) is the core of the vessel-level dynamic model used in aNySIM. It describes how the vessel accelerates in response to the combined hydrodynamic, environmental and mechanical loads.

B.4.2. Coordinate Transformations and Kinematic Consistency

Although hydrodynamic loads are defined in the earth-fixed coordinate system, the equations of motion are solved in the vessel's body-fixed frame. The two frames are linked by a rotation matrix $\mathbf{R}(\boldsymbol{\eta})$ based on the yaw–pitch–roll representation, and a kinematic mapping matrix $\mathbf{J}(\boldsymbol{\eta})$:

$$\dot{\boldsymbol{\eta}} = \mathbf{J}(\boldsymbol{\eta})\,\mathbf{v}.$$

This ensures that positions, velocities and orientations remain consistent, and that forces defined in either frame are transformed correctly.

B.4.3. Coupled Multi-Body System (12-DOF)

When two vessels are moored side-by-side, their motions are coupled through both hydrodynamics and the mechanical connection (mooring lines, fenders). The global configuration and velocity vectors collect the 6-DOF states of both vessels and are therefore of dimension 12×1:

$$\mathbf{X} = \begin{bmatrix} \boldsymbol{\eta}_1 \\ \boldsymbol{\eta}_2 \end{bmatrix}, \qquad \mathbf{V} = \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix},$$

with each subvector $\boldsymbol{\eta}_i,\ \mathbf{v}_i \in \mathbb{R}^6$.

Global inertia and hydrodynamic coupling: The global inertia matrix has size 12×12 and is block-diagonal, since each vessel contributes its own 6×6 total inertia matrix $\mathbf{M}_{\infty,i}$:

$$\mathbf{M}_{\infty}^{\text{glob}} = \begin{bmatrix} \mathbf{M}_{\infty,1} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}_{\infty,2} \end{bmatrix}.$$

In contrast, the global radiation kernel $\mathbf{R}^{\text{glob}}(t)$ is also a 12×12 matrix but, it contains non-zero 6×6 off-diagonal blocks. These off-diagonal terms describe hydrodynamic interaction: motions of vessel 1 generate radiated waves that act on vessel 2, and vice versa, and are essential for representing the coupled surge–sway–yaw dynamics in narrow-gap configurations.

Mechanical coupling: Side-by-side mooring lines, fenders and the FLNG anchoring system also introduce cross-body forces. These appear in the mechanical load term $\mathbf{F}^{\text{glob}}(t)$ and result in additional off-diagonal contributions.

Global equation of motion: The coupled 12-DOF system is written as:

$$\mathbf{M}_{\infty}^{\text{glob}}\,\dot{\mathbf{V}}(t) + \int_{0}^{t} \mathbf{R}^{\text{glob}}(t-\tau)\,\mathbf{V}(\tau)\,\mathrm{d}\tau + \mathbf{K}^{\text{glob}}\,\mathbf{X}(t) = \mathbf{F}_{\text{glob}}(t), \tag{B.49}$$

where $\mathbf{K}^{\text{glob}}$ contains the hydrostatic restoring terms for each vessel, and $\mathbf{F}_{\text{glob}}$ collects all hydrodynamic, mechanical and environmental loads.

Equation (B.49) forms the dynamic backbone of the side-by-side simulation. It ensures that (i) wave radiation from one hull acts on the other, (ii) drift forces and gap dynamics influence the relative surge–sway–yaw behaviour, and (iii) nonlinear SBS mooring-line and fender loads are applied consistently at each time step.

B.4.4. Interpretation

The structure of the global equations highlights the physical mechanisms governing side-by-side operations:

- The inertia matrix reflects both rigid-body and hydrodynamic inertia.
- The convolution term captures memory effects from radiated waves.
- The stiffness term controls static stability and small-amplitude restoring behaviour.
- The right-hand side collects all excitation, damping, environmental and mechanical loads that drive relative motions and mooring-line forces.

This framework provides a description of the coupled dynamics of the FLNG and LNGC, and is consistent with both the hydrodynamic data generated by DIFFRAC and the nonlinear mechanical models used in aNySIM developed by MARIN.

B.5. Conclusion

This chapter established the theoretical framework used in the aNySIM time-domain simulations of side-by-side LNG operations. Starting from linear wave modelling, introduced the hydromechanic load components—hydrostatics, first-order (FK+diffraction) excitation, radiation with memory effects, viscous damping in close proximity, and second-order (drift) loads—and explained how multi-body interactions arise from diffraction shielding and radiation cross-coupling. section B.2 provides the frequency- and time-domain relations that are used consistently in the solver; section B.3 summarised nonlinear SBS lines and fenders; and section B.4 assembled these contributions into the per-vessel 6-DOF and coupled 12-DOF equations of motion.

For side-by-side configurations, two modeling aspects are especially important. First, the hydrodynamics in the narrow gap between the vessels must be handled carefully. Using a free-surface lid in the diffraction calculations helps prevent unrealistically large resonant motions in the gap, while additional viscous damping accounts for the energy losses that mainly affect sway and yaw. Second, the full off-diagonal radiation terms-both added mass and the retardation-function components—need to be included to reproduce the correct wave-frequency heave and pitch motions, as well as the low-frequency horizontal responses. These requirements are supported by validation studies for side-by-side systems and are incorporated in the hydrodynamic database and time-domain model used by MARIN for this study.

Finally, we linked the theory to operability evaluation: the WF-LF motion decomposition clarifies why vertical-plane motions are WF-dominated while horizontal-plane excursions are governed by low-frequency drift, wind, and current; and the manifold-based relative-motion evaluation defines the statistics used later in this thesis. The framework above therefore gives an understanding of the physics-based numerical simulations, which are approached in this thesis with methods explained in chapter 5.

C

Autocorrelation Function Plots

C.1. Wave Frequency ACF's

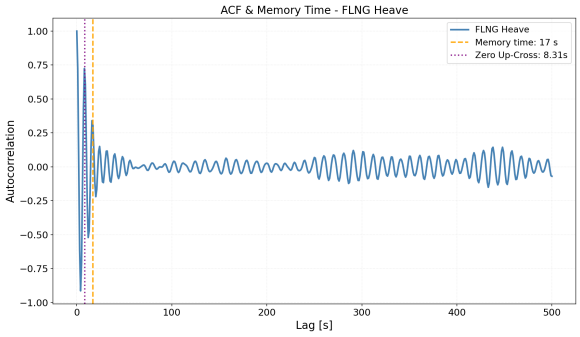


Figure C.1: ACF & memory time for FLNG Heave (WF). Memory time: 17 s; Dominant period: 8.31 s.

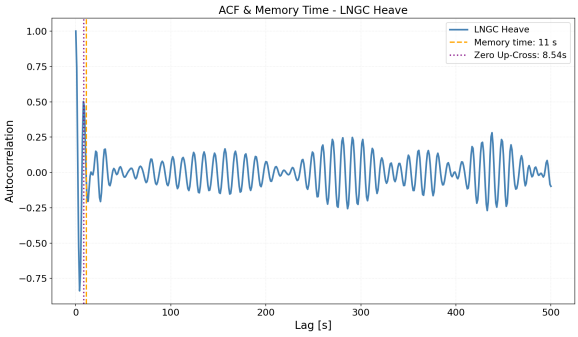


Figure C.2: ACF & memory time for LNGC Heave (WF). Memory time: 11 s; Dominant period: 8.54 s.

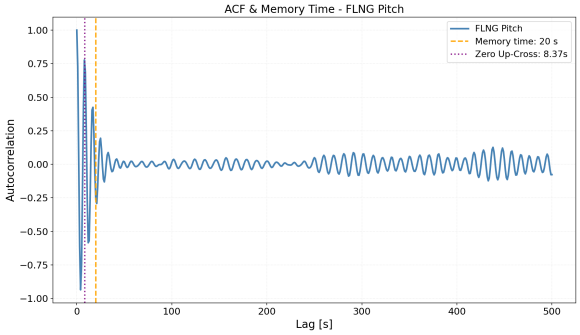


Figure C.3: ACF & memory time for FLNG Pitch (WF). Memory time: 20 s; Dominant period: 8.37 s.

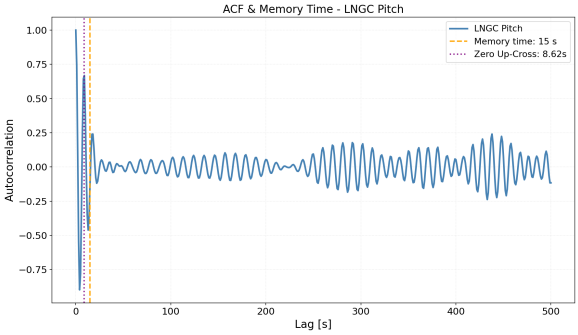


Figure C.4: ACF & memory time for LNGC Pitch (WF). Memory time: 15 s; Dominant period: 8.62 s.

C.2. Low Frequency ACFs

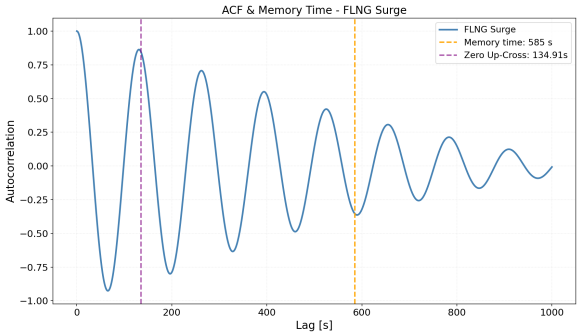


Figure C.5: ACF & memory time for FLNG Surge (LF). Memory time: 585 s; Zero up-cross: 134.91 s.

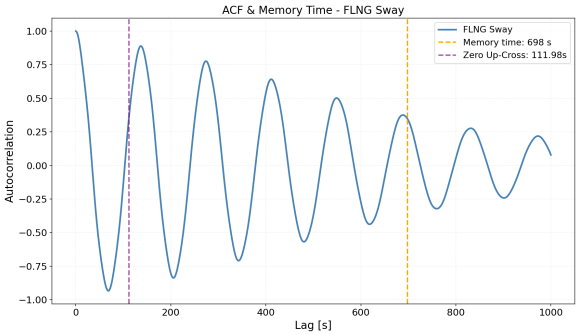


Figure C.6: ACF & memory time for FLNG Sway (LF). Memory time: 698 s; Zero up-cross: 111.98 s.

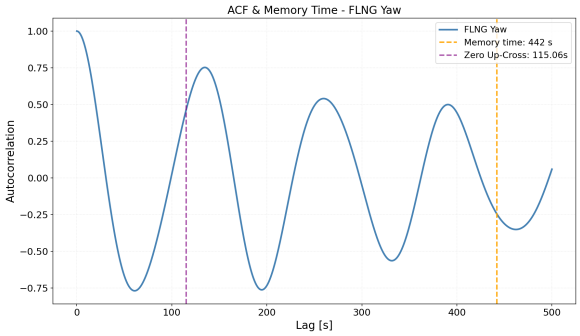


Figure C.7: ACF & memory time for FLNG Yaw (LF). Memory time: 442 s; Zero up-cross: 115.06 s.

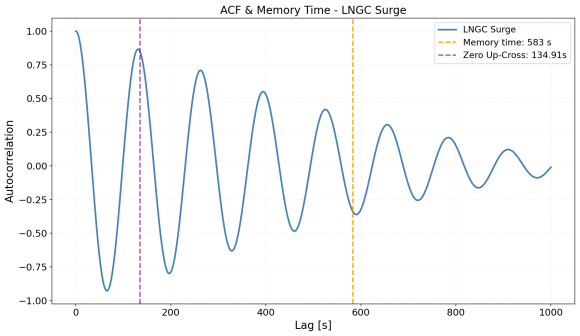


Figure C.8: ACF & memory time for LNGC Surge (LF). Memory time: 583 s; Zero up-cross: 134.91 s.

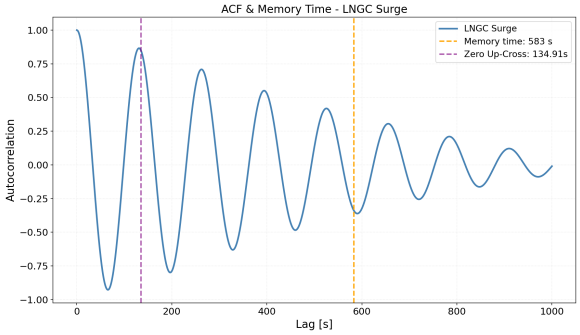


Figure C.9: ACF & memory time for LNGC Surge (LF). Memory time: 738 s; Zero up-cross: 120.91 s.

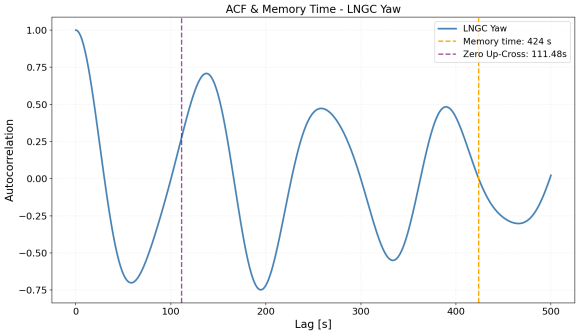


Figure C.10: ACF & memory time for LNGC Yaw (LF). Memory time: 424 s; Zero up-cross: 111.48 s.

D

Detailed Architectures of Models

Summary of Machine Learning Regression Theory

This appendix provides a concise overview of the machine learning concepts underlying the models used in this thesis, following Nielsen [123].

Forward Pass

A neuron computes a weighted sum of its inputs plus a bias term, passed through a nonlinear activation function:

$$a^{(l)} = \sigma\left(W^{(l)} a^{(l-1)} + b^{(l)}\right), \tag{D.1}$$

where $W^{(l)}$ and $b^{(l)}$ are the weights and bias of layer l, and σ is the activation function. This vectorised formulation enables efficient computation across all neurons in a layer simultaneously.

Cost Functions

The mean squared error (MSE) is used to measure the discrepancy between predicted and true values:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^{N} \left(\hat{y}_i - y_i\right)^2. \tag{D.2}$$

Custom loss functions can be defined to emphasise specific aspects of model performance, such as penalising errors near operability thresholds more heavily.

Gradient Descent and Backpropagation

Model weights are updated by gradient descent, moving in the direction that reduces the cost function:

$$W \leftarrow W - \eta \frac{\partial \mathcal{L}}{\partial W}, \tag{D.3}$$

where η is the learning rate. Backpropagation computes the required gradients efficiently by propagating the output error backwards through the network layer by layer, making training of deep networks computationally tractable.

Train-Validation-Test Split

The dataset is divided into three non-overlapping subsets. The training set is used to optimise the model weights. The validation set monitors performance during training and is used for hyperparameter tuning and early stopping. The test set provides a final unbiased estimate of generalisation performance on unseen data.

Overfitting and Regularisation

Overfitting occurs when a model performs well on training data but poorly on unseen data. Early stopping halts training when the validation loss no longer improves, preventing the model from memorising the training set. Two common regularisation techniques are used to improve generalisation:

- **Weight decay (L2 regularisation):** adds a penalty proportional to the squared magnitude of the weights, discouraging large weight values.
- **Dropout:** randomly disables a fraction of neurons during each training step, reducing co-adaptation between neurons.

Adaptive Learning Rate Scheduling

A fixed learning rate can cause training to overshoot minima or stall in flat regions of the loss surface. Learning rate scheduling reduces the learning rate when progress stalls, allowing the optimiser to refine convergence without overshooting.

D.1. Segment-Based Sequence-to-Sequence Learning: Model Specifications

This appendix provides the detailed architectural configurations of the segment-based sequence-to-sequence (seq2seq) learning models used for low-frequency (LF) and wave-frequency (WF) motion prediction. The main chapters discuss the modelling choices and the physical reasoning behind the architectures; here, the focus is strictly on the implementation-level details required for reproducibility.

All activation functions and normalisation layers referenced in the notes follow the standard implementations used in the main model code.

D.1.1. Low-Frequency (LF) Model

Table D.1: LF motion encoder specification

Layer	Input size	Hidden size	Notes
LSTM layer 1	13	192	12 motion DOFs + 1 elevation
LSTM layer 2	192	192	deeper temporal abstraction

Motion encoder:

Table D.2: LF wave CNN specification

Layer	Channels (in→out)	Kernel	Pool stride	Notes
Conv 1	$1 \rightarrow 32$	5	2	BN + ReLU
Conv 2	$32 \rightarrow 64$	5	2	BN + ReLU
Conv 3	$64 \rightarrow 64$	5	2	BN + ReLU
Conv 4	$64 \rightarrow 64$	5	2	BN + ReLU
Fully connected	–	–	–	produces 96-dim embedding

Wave encoder: A four-layer 1D CNN with batch normalisation, ReLU activation, and max pooling is used to compress the wave elevation segment into a compact 96-dimensional embedding.

Table D.3: LF motion decoder specification

Layer	Input size	Hidden size
LSTM layer 1	96	192
LSTM layer 2	192	192
Output fully connected	$192 \rightarrow 12$	–

Motion decoder and output: The LF decoder reconstructs the slow-varying future motion trajectory in all 12 degrees of freedom, using both wave and motion context.

D.1.2. Wave-Frequency (WF) Model

Dual-stream encoder. The WF model uses separate encoders for motion and wave signals, reflecting the higher sensitivity of WF dynamics to both instantaneous vessel kinematics and wave-phase information.

Table D.4: WF encoder specification

Stream	Layers	Input size	Hidden size	Notes
Motion LSTM	2	12	96	vessel inertia, damping, restoring
Wave LSTM	1	3	64	wave phase and grouping

Table D.5: WF wave CNN specification

Layer	Channels (in→out)	Kernel	Pool stride	Notes
Conv 1	$3 \rightarrow 32$	5	2	BN + ReLU
Conv 2	$32 \rightarrow 64$	5	2	BN + ReLU
Conv 3	$64 \rightarrow 64$	5	2	BN + ReLU
Conv 4	$64 \rightarrow 64$	5	2	BN + ReLU
Fully connected	–	–	–	produces 96-dim embedding

Wave CNN encoder:

Table D.6: WF motion decoder specification

Layer	Input size	Hidden size
LSTM layer 1	96	160
LSTM layer 2	160	160
Output fully connected	$160 \rightarrow 12$	–

Motion decoder: The WF decoder predicts the short-period oscillatory vessel motions by integrating embeddings from the dual encoders and CNN, yielding a 12-DOF wave-frequency motion reconstruction.

D.2. Full Time Series-Based Learning

This appendix provides the implementation-level details of all deep-learning components used in the full time series–based prediction framework. While the main chapters focus on the modelling rationale, physical interpretation, and comparative performance, the material collected here documents the exact architectural configurations required for full reproducibility.

LF wave CNN (1 channel: elevation).

Table D.7: LF wave CNN (1-channel elevation)

Layer	Output shape	Kernel size
Input (elevation)	$[1, T]$	–
Conv1	$[64, T]$	15
Conv2	$[128, T]$	11
Conv3	$[256, T]$	7
MaxPool (stride 2)	$[256, T/2]$	2
MaxPool (stride 2)	$[256, T/4]$	2
Projection (1×1)	192-dim	–

WF wave CNN (4 channels: elevation + derivatives). Derivatives emphasise changes in the wave signal (slopes, zero-crossings, frequency), which are correlated with wave-frequency vessel motions.

Table D.8: WF wave CNN (4 channels: elevation + derivatives)

Layer	Output shape	Kernel size
Input (4 channels)	$[4, T]$	–
Conv1	$[64, T]$	5
Conv2	$[128, T]$	5
Conv3	$[256, T]$	5
MaxPool (stride 2)	$[256, T/2]$	2
Conv4	$[512, T/2]$	5
MaxPool (stride 2)	$[512, T/4]$	2
Projection (1×1)	128-dim	–

Design rationale. A kernel size of 5 means that each convolution operation looks at a window of 5 consecutive time samples at once when extracting features. LF uses larger, decreasing kernels (15, 11, 7) to capture envelope and wave-group structures. WF uses uniform smaller kernels (5) to preserve high-frequency phase detail. Both downsample to $T/4 = 2{,}568$ steps to match Transformer inputs.

Transformer: After fusion, each branch uses its own Transformer encoder (multi-head self-attention, ReLU, dropout 0.2, post-activation layer normalization). The feed-forward hidden size is $4 \times D_{\text{fuse}}$ and the sequence length remains $2{,}568$.

Table D.9: Transformer encoder parameters (LF vs. WF)

Parameter	LF branch	WF branch
Fused embedding dimension (input)	270	192
Number of attention heads	6	4
Number of encoder layers	3	3
Output embedding dimension	270	192

Interpretation: Each timestep embedding attends to the full $2{,}568$-step context. The larger LF embedding (270) and more heads (6) reflect the added complexity and the smaller amount of training cycles; the WF branch uses a more compact representation (192, 4 heads) adequate for frequency-dependent dynamics.

D.2.1. LF Model Specifics

Table D.10: RF baseline (LF mean only)

Item	Setting / Result
Algorithm	Random Forest Regressor
Model size	200 trees, max depth 20, min split 5
Inputs	11 directional/encoded environmental features
Outputs	12 DOFs (two vessels) + 16 line means

Table D.11: LF oscillation decoders

Decoder	Input	Hidden	Output	Notes
Motions	270-dim	[384, 256, 128]	12 DOFs	ReLU, BatchNorm, residual
Tensions	270-dim	[384, 256, 128]	16 lines	ReLU, BatchNorm, residual

D.2.2. WF Model Specifics

The 12 predicted motions are first embedded into a 256-dim vector, fused with transformer features, and decoded by ReLU-activated layers. Sixteen small per-line heads then output the individual line tensions.

Table D.12: WF decoder architecture

Component	Input dim	Hidden dims	Output dim
Motion Decoder	192	[768, 384, 192, 96]	12
Motion Embedding	12	[256]	256
Tension Decoder	192 + 256	[512, 256, 128, 64]	16
Per-line heads (×16)	64	[32]	1 each

D.2.3. Hyperparameter Optimization

Hyperparameter optimization was performed using Optuna's Tree-structured Parzen Estimator (TPE), a Bayesian optimisation method suitable for moderate search spaces. Two neural network models were tuned for their respective frequency regimes (LF and WF), and two Random Forest baselines were optimised as reference models. Each trial used early stopping on a dedicated validation split.

Search Spaces: The hyperparameter ranges explored for the LF and WF models are summarised in Table D.13. These include learning-rate schedules, Transformer capacity, decoder dimensions, and CNN widths.

Table D.13: Search spaces for the LF and WF models.

(a) LF model search space

Hyperparameter	Range
Learning rate	5×10^{-4}–5×10^{-3}
Transformer heads	{4, 8}
Transformer layers	{2, 3}
Decoder hidden dim	{128, 256, 512}
Wave embedding dim	{64, 128, 256}
Peak-alignment weight	1.5–2.5
Spectral weight	0.8–1.2

(b) WF model search space

Hyperparameter	Range
Learning rate	10^{-5}–10^{-3}
Batch size	{4, 8, 16}
Wave embedding dim	{64, 128, 256}
Condition embedding dim	{32, 64, 128}
Fusion dimension	{128, 192, 256}
Transformer heads	{2, 4, 8}
Transformer layers	{2, 3, 4}
Decoder hidden dim	{256, 512, 768}
Dropout rate	{0.1, 0.2, 0.3}
Wave CNN base channels	{32, 64, 128}

D.2.3.1. Low-Frequency Model Optimization

The LF model was optimised for trials using only 20% of the available training data (300 of 1500 simulations). This limited sample size produced unstable validation behaviour, leading to overfitting within the search space.

Outcome: The optimised LF configuration did not outperform the manual baseline. The best validation loss (5.124) was higher than the baseline (3.994). The primary causes were (i) insufficient data for Bayesian optimisation, (ii) premature early stopping, and (iii) excessive model capacity relative to dataset size.

Recommended configuration: Future LF optimisation should use the full dataset, increased early-stopping patience, and a larger number of trials.

Table D.14: LF baseline configuration retained for deployment.

Learning rate = 10^{-3}	Transformer heads = 6	Transformer layers = 3
Decoder hidden dim = 384	Wave embedding dim = 192	

D.2.3.2. Wave-Frequency Model Optimization

The WF model was optimised using 35+ trials on the full dataset (1500 simulations). This allowed for a stable validation signal and reliable convergence.

Table D.15: Optimised WF configuration retained for deployment.

lr = 1.84×10^{-4}	batch = 4	wave_emb = 128	cond_emb = 128	fusion = 192
heads = 4	layers = 3	decoder_hidden = 768	dropout = 0.2	

Outcome: The WF optimisation succeeded. Key factors were the use of the full dataset, a small batch size, a low learning rate, and sufficiently long training horizons.

D.2.3.3. Random Forest Baselines

Random Forest models for position and tension prediction were optimised over 30 trials each. Their discrete hyperparameter space and ensemble stability resulted in consistent convergence.

Table D.16: Optimised Random Forest hyperparameters for position and tension prediction.

(a) RF Position Model

Hyperparameter	Value
n_estimators	250
max_depth	27
min_samples_split	3
min_samples_leaf	1
max_features	log2

(b) RF Tension Model

Hyperparameter	Value
n_estimators	300
max_depth	24
min_samples_split	3
min_samples_leaf	1
max_features	log2

D.2.3.4. Key Insights

- Full-dataset usage is essential for stable neural-network optimisation.
- Early stopping patience must scale with data volume to avoid premature convergence.
- Model capacity must be aligned with dataset size to prevent overfitting.
- Small batch sizes improved optimisation stability for the WF model.
- RF models provided robust baselines due to their structured hyperparameter spaces.

D.2.3.5. Conclusion

The optimisation study highlights the importance of dataset size, early stopping configuration, and appropriate model capacity. The WF model benefited from full-data optimisation and converged reliably, whereas the LF model failed due to subsampling and premature stopping. Random Forest models offered stable baseline performance and serve as reference models for evaluating neural network improvements.

E

Prediction Statistics

E.1. SSBL Full Dataset Predictions

Table E.1: Top-5 most frequent 3-parameter combinations in False Safe cases for FLNG roll (min/max).

Parameter 1	Parameter 2	Parameter 3	Count
seaDir = 270	seaHs = 2.5	seaTp = 8.0	164
seaHs = 2.5	seaTp = 8.0	swellHs = 1.5	78
seaHs = 2.5	seaTp = 8.0	swellTp = 12.0	76
seaHs = 2.5	seaTp = 8.0	swellHs = 1.0	70
seaHs = 2.5	seaTp = 8.0	swellTp = 20.0	64

Table E.2: Top-5 most frequent 3-parameter combinations in False Safe cases for LNGC roll (min/max).

Parameter 1	Parameter 2	Parameter 3	Count
seaDir = 90	seaHs = 2.5	seaTp = 8.0	141
seaDir = 120	seaHs = 1.5	seaTp = 10.0	98
seaHs = 2.5	seaTp = 8.0	swellHs = 1.0	59
seaDir = 90	seaHs = 2.5	swellHs = 1.0	57
seaHs = 2.5	seaTp = 8.0	swellTp = 20.0	57

Table E.3: Top-5 most frequent 3-parameter combinations in False Safe cases for Manifold Distance X.

Parameter 1	Parameter 2	Parameter 3	Count
seaDir = 90	seaHs = 2.5	seaTp = 8.0	102
seaHs = 2.5	seaTp = 8.0	swellTp = 20.0	81
seaDir = 0	seaHs = 0	seaTp = 10.0	79
seaHs = 2.5	swellHs = 1.5	swellTp = 20.0	71
seaHs = 2.5	seaTp = 8.0	swellHs = 1.5	63

Table E.4: Top-5 most frequent 3-parameter combinations in False Safe cases for Manifold Distance Y.

Parameter 1	Parameter 2	Parameter 3	Count
seaDir = 90	seaHs = 1.0	seaTp = 10.0	170
seaDir = 270	seaHs = 2.0	seaTp = 8.0	149
seaDir = 270	seaHs = 1.0	seaTp = 10.0	142
seaHs = 1.0	seaTp = 10.0	swellHs = 1.5	120
seaHs = 1.0	seaTp = 10.0	swellTp = 12.0	114

Table E.5: Line Tension violation summary.

Metric	Count
False Safe cases (dangerous)	0
False Unsafe cases	26,650

Table E.6: Fender Force violation summary.

Metric	Count
False Safe cases (dangerous)	0
False Unsafe cases	12,240

E.2. TSBL Head Waves Prediction

These prediction statistics were derived by analysing, for each operability parameter, the full distribution of model outcomes across the test set and identifying the environmental combinations of incorrect classifications. The False Safe (FS) and False Unsafe (FU) rates were computed within each individual parameter. The tables in this appendix list the most frequent 3-parameter combinations that produced false safe cases, together with summary counts for line-tension and fender-force violations. This provides a compact overview of the specific environmental regimes where the SSBL model struggles most and where operability margins are tightest.

Table E.7: Most difficult environmental condition ranges for each predicted parameter, ranked by False Safe (FS) rate.

Parameter	Condition	Range	FS Rate	n
Relative Manifold Distance (Surge)				
Rel. Manifold Surge	Sea height	Q4: 3.00–3.50 m	52.3%	107
Rel. Manifold Surge	Sea period	Q1: 6.00–7.00 s	19.3%	140
Rel. Manifold Surge	Sea period	Q2: 8.00–9.00 s	11.8%	170
Relative Manifold Distance (Sway)				
Rel. Manifold Sway	Sea period	Q1: 6.00–7.00 s	2.1%	140
Rel. Manifold Sway	Sea height	Q4: 3.00–3.50 m	1.9%	107
Rel. Manifold Sway	Sea height	Q3: 2.50–2.50 m	1.5%	65
Line Tensions (Any Line)				
Tension	Sea height	Q4: 3.00–3.50 m	31.8%	107
Tension	Sea height	Q3: 2.50–2.50 m	21.5%	65
Tension	Sea period	Q1: 6.00–7.00 s	20.7%	140

Table E.8: Most difficult environmental conditions per tension line (Lines 1–8 left, 9–16 right). FS Count column widened for readability.

Lines 1–8			Lines 9–16		
Line	**FS Count**	**Most Difficult Conditions**	**Line**	**FS Count**	**Most Difficult Conditions**
1	16 (3.5%)	Sea height [3.00–3.50] (7) Sea height [1.50–2.00] (9) Sea period [6.00–7.00] (9)	9	4 (0.9%)	Sea height [1.50–2.00] (4) Sea period [6.00–7.00] (4) Sea height [0.75–1.00] (0)
2	16 (3.5%)	Sea height [3.00–3.50] (7) Sea height [1.50–2.00] (9) Sea period [6.00–7.00] (9)	10	9 (2.0%)	Sea period [6.00–7.00] (9) Sea height [1.50–2.00] (8) Sea height [2.50–2.50] (1)
3	17 (3.7%)	Sea height [3.00–3.50] (17) Sea period [8.00–9.00] (17) Sea height [0.75–1.00] (0)	11	24 (5.3%)	Sea height [2.50–2.50] (13) Sea period [6.00–7.00] (17) Sea height [3.00–3.50] (11)
4	13 (2.9%)	Sea height [3.00–3.50] (13) Sea period [8.00–9.00] (13) Sea height [0.75–1.00] (0)	12	26 (5.7%)	Sea height [2.50–2.50] (13) Sea period [6.00–7.00] (19) Sea height [3.00–3.50] (13)
5	15 (3.3%)	Sea period [6.00–7.00] (12) Sea height [1.50–2.00] (9) Sea height [3.00–3.50] (6)	13	26 (5.7%)	Sea height [2.50–2.50] (13) Sea period [6.00–7.00] (19) Sea height [3.00–3.50] (13)
6	17 (3.7%)	Sea height [3.00–3.50] (8) Sea period [6.00–7.00] (10) Sea height [1.50–2.00] (9)	14	16 (3.5%)	Sea height [3.00–3.50] (16) Sea period [8.00–9.00] (14) Sea period [6.00–7.00] (2)
7	0 (0.0%)	—	15	16 (3.5%)	Sea height [3.00–3.50] (7) Sea height [1.50–2.00] (9) Sea period [6.00–7.00] (9)
8	0 (0.0%)	—			
			16	17 (3.7%)	Sea height [3.00–3.50] (8) Sea height [1.50–2.00] (9) Sea period [6.00–7.00] (9)

F

Final Results Plots & Prediction Statistics

F.1. SSBL Insights

Below, the individual prediciton results are plotted against the actual values in these parity plots.

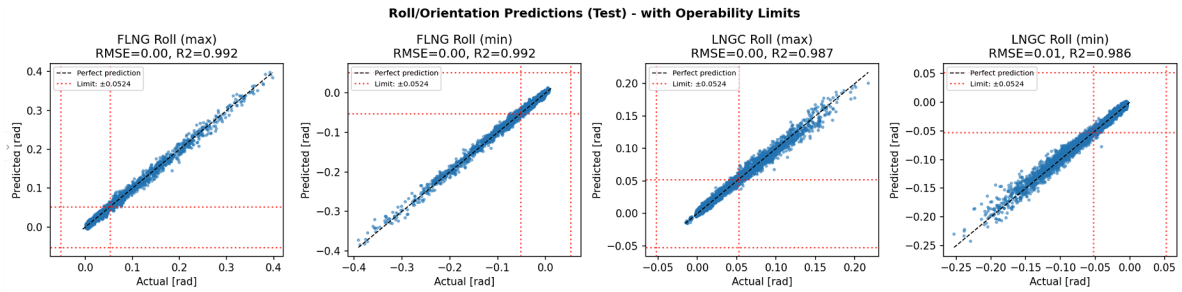


Figure F.1: Roll predictions (max/min) for FLNG and LNGC with operability limits.

Roll Predictions and Orientation Limits Figure F.1 shows excellent agreement between predicted and actual roll values, with R^2 values between 0.986 and 0.992. All roll predictions fall comfortably within the $\pm 0.0524\,\mathrm{rad}$ orientation limit. The scatter points cluster tightly around the 1:1 diagonal, demonstrating that the model does not underpredict roll near the operability boundary. These results indicate that the MLP can reliably support orientation-limit assessment.

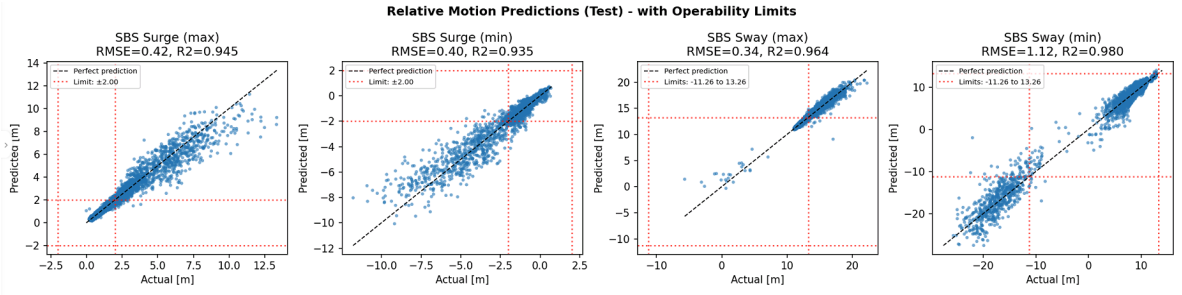


Figure F.2: Relative surge and sway predictions (max/min) with operability limits.

Relative Surge and Sway Predictions The relative motion results in Figure F.2 show strong performance for both surge and sway. Relative surge predictions obtain $R^2 = 0.935$–0.945, with unbiased scatter near the $\pm 2\,\mathrm{m}$ limit, meaning the model correctly identifies when longitudinal clearance constraints are approached. Relative sway predictions achieve $R^2 = 0.964$–0.980. Although sway exhibits slightly larger variance in the extreme regions, the model still separates safe and unsafe conditions well. Importantly, no systematic underprediction is observed near the limits.

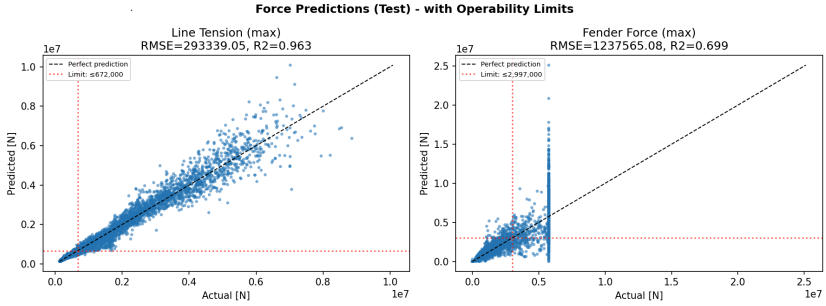


Figure F.3: Predictions for maximum line tension and maximum fender force with operability limits.

Line Tension and Fender-Force Predictions Figure F.3 shows that line-tension predictions reach a high accuracy ($R^2 = 0.963$), with scatter concentrated around the diagonal and no tendency to underestimate values near the 5.672 MN limit. Fender forces are inherently more nonlinear, which is reflected in a lower but still acceptable $R^2 = 0.699$. The model remains conservative near the 2.997 MN limit, avoiding false-safe classifications.

Overall, the model shows strong predictive performance across all variables relevant for operability. Roll and relative motions are predicted with R^2 values above 0.94 and low scaled RMSE, indicating that the network captures both trend and variability well. Line tension prediction is also highly accurate ($R^2 = 0.963$), ensuring correct identification of limiting mooring conditions. Fender-force prediction is less precise—as expected for a highly nonlinear response—but remains conservative near the threshold, preventing unsafe underestimation.

Quantity	R^2	RMSE	sRMSE	Quantity	R^2	RMSE	sRMSE
FLNG Roll (max)	0.992	4.83e–03	0.09	LNGC Roll (max)	0.987	4.05e–03	0.10
FLNG Roll (min)	0.992	4.75e–03	0.09	LNGC Roll (min)	0.986	5.08e–03	0.11
Rel. Surge (max)	0.945	4.22e–01	0.23	Rel. Sway (max)	0.964	3.44e–01	0.19
Rel. Surge (min)	0.935	3.97e–01	0.24	Rel. Sway (min)	0.980	1.12e+00	0.13
Line Tension (max)	0.963	2.93e+05	0.13	Fender Force (max)	0.699	1.24e+06	0.17

Table F.1: Performance metrics (R^2, RMSE, and scaled RMSE) for the summary-based MLP model.

These results show that the model reproduces the main operability-relevant behaviours with good fidelity. LNGC roll—one of the more difficult motions to predict in this study—is captured accurately, and relative surge and sway show strong agreement across the full range of conditions. Combined, these metrics indicate that the model provides sufficiently accurate predictions to support operability screening.

F.2. False Safe Environmental Pattern Analysis

This appendix provides a consolidated overview of the environmental conditions associated with *False Safe* cases, i.e., simulations where the model predicts operability while the MARIN reference classifies the case as non-operable. A total of 67 out of 6816 test cases (0.98%) fall into this critical category. The tables below present the most frequent three-parameter combinations found in the full False Safe set and for each individual limiting operability parameter. All swell-related variables have been removed as requested.

F.2.1. Most Frequent Combinations Across All False Safe Cases

Table F.2: Top 5 environmental combinations in all False Safe cases (swell parameters removed).

Parameter 1	Parameter 2	Parameter 3	Count
windVel = 10.00	seaHs = 1.50	currVel = 0.24	22
windVel = 4.00	seaHs = 1.50	currVel = 0.24	22
windVel = 10.00	seaHs = 2.50	currVel = 0.24	21
windVel = 10.00	seaTp = 11.00	currVel = 0.24	19
windVel = 10.00	seaTp = 10.00	currVel = 0.24	18

Interpretation: False Safe cases in the total set cluster in moderate sea states ($H_s = 1.5$–2.5 m), mid–peak periods ($T_p = 10$–11 s), and always at the maximum current velocity (0.24 m/s). These combinations reflect conditions where wind-driven and current-driven drift forces interact most strongly.

F.2.2. False Safe Patterns for Manifold Surge Limit

Table F.3: Top 5 environmental combinations for False Safe surge-limit violations (swell removed).

Parameter 1	Parameter 2	Parameter 3	Count
windVel = 10.00	seaHs = 2.50	currVel = 0.24	47
seaDir = 105.00	seaHs = 2.50	currVel = 0.24	45
windVel = 10.00	seaTp = 10.00	currVel = 0.24	43
windVel = 10.00	seaHs = 3.00	currVel = 0.24	41
windVel = 4.00	seaHs = 2.50	currVel = 0.24	47

Interpretation: Surge-limit False Safes occur predominantly when moderate sea states ($H_s = 2.5$–3.0 m) combine with the maximum current velocity, amplifying low-frequency drift and mean offsets that the model underestimates.

F.2.3. False Safe Patterns for Manifold Sway Limit

Table F.4: Top 5 environmental combinations for False Safe sway-limit violations (swell removed).

Parameter 1	Parameter 2	Parameter 3	Count
windVel = 10.00	seaTp = 11.00	currVel = 0.24	46
windVel = 4.00	seaTp = 11.00	currVel = 0.24	45
seaDir = 255.00	seaTp = 11.00	currVel = 0.24	31
seaHs = 0.75	windVel = 10.00	currVel = 0.24	30
seaTp = 8.00	windVel = 10.00	currVel = 0.24	36

Interpretation: Sway-limit False Safes cluster around mid-period seas ($T_p = 11$ s) with strong currents. These regimes promote transverse drift amplification and lateral offsets that exceed limits.

F.2.4. False Safe Patterns for Line Tension Limit

Table F.5: Top 5 environmental combinations for False Safe line tension violations (swell removed).

Parameter 1	Parameter 2	Parameter 3	Count
windVel = 10.00	seaHs = 1.50	currVel = 0.24	25
windVel = 4.00	seaHs = 1.50	currVel = 0.24	25
windVel = 10.00	seaTp = 11.00	currVel = 0.24	21
seaHs = 2.50	windVel = 10.00	currVel = 0.24	21
seaTp = 10.00	windVel = 10.00	currVel = 0.24	18

Interpretation: Line-tension False Safes arise in moderate sea states ($H_s = 1.5$–2.5 m) combined with strong currents, which increase slow-drift excursions and tension peaks that the model tends to underpredict.

F.2.5. False Safe Patterns for Fender Force Limit

Table F.6: Top 5 environmental combinations for False Safe fender force violations (swell removed).

Parameter 1	Parameter 2	Parameter 3	Count
windVel = 10.00	seaTp = 11.00	currVel = 0.24	64
windVel = 4.00	seaTp = 11.00	currVel = 0.24	64
seaDir = 120.00	seaTp = 11.00	currVel = 0.24	45
seaTp = 10.00	windVel = 10.00	currVel = 0.24	47
seaHs = 0.75	windVel = 10.00	currVel = 0.24	39

Interpretation: Fender-force False Safes concentrate around mid-period seas ($T_p = 10$–11 s) under strong currents, where relative surge–sway motions intensify contact loads that the model systematically underestimates.

Summary. Across all tables, the same environmental structure dominates: moderate wind speeds (4–10 m/s), mid-range sea states ($H_s = 1.5$–3.0 m and $T_p = 8$–11 s), and the maximum current velocity (0.24 m/s). These conditions produce elevated low-frequency drift responses and mean offsets, even without swell, and appear to be systematically underestimated by the model, leading to the observed False Safe outcomes.

F.3. TSBL Environmental Pattern Analysis

A total of 736 out of 6816 cases (10.8%) resulted in *False Safe* predictions. The tables below list the most frequent environmental combinations associated with these cases, grouped by their relevance to the limiting operability parameters.

F.3.1. Combinations of False Safe Classifications

Table F.7: Top 5 most frequent 3□parameter combinations in False Safe cases.

windVel	seaHs	currVel	Count
10.00	3.00	0.24	156
4.00	3.00	0.24	142
180.00†	0.24	180.00	108
0.00†	0.24	0.00	105
315.00†	0.24	315.00	101

Interpretation: High wind velocity, sea state $H_s = 3$ m, and maximum current speed (0.24 m/s) dominate False Safe regions. Directional alignment of wind and current is a consistent contributor.

F.3.2. Combinations False-Safe Classification

Table F.8: Top 5 most frequent 4-parameter combinations in False Safe cases.

windVel	windDir	currVel	currDir	Count
10.00	180.00	0.24	180.00	55
4.00	180.00	0.24	180.00	53
10.00	0.00	0.24	0.00	53
4.00	0.00	0.24	0.00	52
4.00	315.00	0.24	315.00	52

Interpretation: The most critical 4-parameter combinations involve wind and current moving in the same direction, reinforcing drift loads.

F.3.3. Combinations Causing Surge-Limit False Safe Cases

Table F.9: Top 5 environmental combinations leading to Relative Surge limit violations.

Parameter Set	Value Set	Count
windVel / seaHs / currVel	(10.00, 3.00, 0.24)	156
windVel / seaHs / currVel	(4.00, 3.00, 0.24)	142
windDir / currVel / currDir	(180.00, 0.24, 180.00)	108
windDir / currVel / currDir	(0.00, 0.24, 0.00)	105
windDir / currVel / currDir	(315.00, 0.24, 315.00)	101

Interpretation: $H_s = 3$ m, and strong current-wind alignment dominate surge-limit failures.

F.3.4. Top Combinations Causing Sway-Limit False Safe Cases

Table F.10: Top 5 environmental combinations leading to Relative Sway limit violations.

Parameter Set	Value Set	Count
windVel / seaHs / currVel	(10.00, 3.00, 0.24)	156
windVel / seaHs / currVel	(4.00, 3.00, 0.24)	142
windDir / currVel / currDir	(180.00, 0.24, 180.00)	108
windDir / currVel / currDir	(0.00, 0.24, 0.00)	105
windDir / currVel / currDir	(315.00, 0.24, 315.00)	101

Interpretation: The sway-limit False Safes match the surge pattern—dominated by the $H_s = 3$ m sea state and maximum current velocity.

F.3.5. Combinations Causing Roll-Limit False Safe Cases

Table F.11: Top 5 environmental combinations leading to Roll limit violations.

Parameter Set	Value Set	Count
windVel / seaHs / currVel	(10.00, 3.00, 0.24)	156
windVel / seaHs / currVel	(4.00, 3.00, 0.24)	142
windDir / currVel / currDir	(180.00, 0.24, 180.00)	108
windDir / currVel / currDir	(0.00, 0.24, 0.00)	105
windDir / currVel / currDir	(315.00, 0.24, 315.00)	101

Interpretation: Roll-limit False Safes also cluster around the same dominant forcing regime as surge and sway.

F.3.6. Top Combinations Causing Line-Tension False Safe Cases

Table F.12: Top 5 environmental combinations associated with tension-limit violations.

Parameter Set	Value Set	Count
windVel / seaHs / currVel	(10.00, 3.00, 0.24)	156
windVel / seaHs / currVel	(4.00, 3.00, 0.24)	142
windDir / currVel / currDir	(180.00, 0.24, 180.00)	108
windDir / currVel / currDir	(0.00, 0.24, 0.00)	105
windDir / currVel / currDir	(315.00, 0.24, 315.00)	101