· UNIVERSITY OF OXFORD ·

· DEPARTMENT OF ENGINEERING SCIENCE ·

Structural Identification of Offshore Wind Turbine Foundations: Accounting for Modelling Error

Harry Simpson

Exeter College



A thesis submitted for the degree of

Doctor of Philosophy

Supervised by Professor Manolis Chatzis & Professor Eleni Chatzi

September 2025

Abstract

The offshore wind industry is expected to experience significant growth over the next decade. Structural health monitoring strategies, where measurements from installed sensors are combined with physics-informed models, can help ensure turbines are not at risk of fatigue or damage, thereby maximising availability. Furthermore, structural health monitoring is likely to play a key role in assessing the viability for lifetime extension, offering the potential for substantial cost savings across the industry.

However, the performance of structural health monitoring methods rely on the accuracy of the numerical model representing the real structure, which is often uncertain in practice. A major source of this uncertainty is the soil-structure interaction and previous studies have shown that quantities of interest, such as fatigue, can be highly sensitive to the parameters governing it.

To address this, model updating techniques are explored in which measured natural frequencies and mode shapes identified from ambient vibration data are compared with model predictions to calibrate foundation parameters. A compact macro-element foundation model is used to limit the number of unknown parameters. Deterministic and Bayesian model updating methods are evaluated in numerical case studies. In the absence of modelling error, the foundation parameters are shown to be accurately identified from limited tower accelerometer data. While the deterministic approach is advantageous due to its computational efficiency, Bayesian methods offer more interpretable insights by quantifying both parameter and model uncertainties.

Although the primary focus of this work is to identify the foundation parameters, model uncertainties elsewhere in the structure complicate this task. An additional source of error arises from the flexible rotor nacelle assembly, where blade properties are often unknown and must be scaled from reference turbines. This introduces significant uncertainty that is shown to reduce the robustness of model updating frameworks using modal data.

To overcome this modelling challenge, substructuring techniques for parameter, input, and state estimation are investigated, eliminating the need to model the rotor nacelle assembly explicitly. A new time-domain Bayesian framework is proposed to estimate macro-element foundation parameters and is shown to perform well when rotor nacelle assembly properties are unknown. Additionally, an alternative substructure that excludes both the soil-structure interaction and the rotor nacelle assembly is considered for input-state estimation. Although this unconstrained substructure presents challenges, it is advantageous when the ground-level response must be estimated without making any assumptions about foundation behaviour. With sufficient measurements, numerical case studies demonstrate accurate estimation of system inputs, states, and fatigue.

Acknowledgements

Firstly, I would like to thank my supervisors, Manolis Chatzis and Eleni Chatzi, for their support and encouragement. Their guidance, expertise and constructive feedback have been essential in shaping me into an independent researcher, and this thesis would not have been possible without them. I would also like to extend this gratitude to Ross McAdam for his supervision during the initial months of the project. In addition, I am extremely grateful to RWE for funding my DPhil and for providing data to support this research. In particular, I would like to thank Steven Heald, Russell Mayall and Gianluca Zorzi for their insights, which helped ensure the research remained relevant to industry.

Throughout my DPhil, I was fortunate to benefit from the support of various research groups. Being part of the WAMESS CDT network gave me the opportunity to attend exciting trips and courses that helped develop my career and build valuable connections. My time with the Structural Mechanics and Monitoring group at ETH Zurich played a key role in developing my experience with state-of-the-art methods in structural health monitoring. I also learned a great deal from the Nonlinear Dynamics and Civils research groups at the University of Oxford, where seminars from experienced visiting researchers contributed significantly to broadening my knowledge. I am grateful to the lifelong friends I have made in each of these groups, whose support and welcome distractions kept me motivated. Thank you as well to my old friends from Edinburgh, who reminded me of life outside the DPhil and helped me maintain a healthy work-life balance.

Last but not least, to my family and Rosie, whose unconditional support helped me to persevere during the more demanding times, for which I will always be grateful.

Contents

Abstract ii
Acknowledgements iii

CHAPTER 1: Introduction and Literature Review

1.1 Background and Motivation 1
1.2 Literature Review 4
1.2.1 Offshore wind turbine modelling 4
1.2.2 System identification and operational modal analysis. 10
1.2.3 Model updating 12
1.2.4 Recursive methods for state, input and parameter estimation 18
1.3 Research objectives and thesis structure 23
1.3.1 Chapter Overview 24

CHAPTER 2: Simulink Offshore Wind Turbine Model

2.1 Introduction 27
2.2 Offshore Wind Turbine Modelling Requirements 27
2.3 Hydro-Servo-Aero-Elastic Simulink Model 29
2.3.1 Soil-Structure Interaction Modelling 29
2.3.2 Tower, Transition Piece and Monopile Modelling 31
2.3.3 Rotor Nacelle Assembly Modelling 32
2.3.4 Aerodynamics 34
2.3.5 Hydrodynamics 36
2.3.6 Controller 39
2.3.7 Full-System Integration 39
2.4 Summary 42

CHAPTER 3: Model Updating Using Modal Data

3.1 Introduction 44
3.2 Operational Modal Analysis 46
3.3 Model Updating Frameworks 48
3.3.1 Deterministic Model Updating 48
3.3.2 Bayesian Model Updating 49
3.4 Numerical Case Study 54
3.4.1 Wind Turbine Description 55
3.4.2 Synthetic Data Generation and Sensor Set-up 57
3.4.3 Operational Modal Analysis Results 59
3.4.4 Estimation of the Foundation Parameters (model updating results) 64
3.5 Summary 76

CHAPTER 4: Modelling Challenges for Foundation Parameter Estimation

4.1 Introduction 79

4.2 Baseline Wind Turbine Model . . . 80
4.3 Blade Scaling Methods . . . 81
4.3.1 Reference Turbines . . . 82
4.3.2 Influence on Blade Property Distributions . . . 83
4.3.3 Influence on Modal Parameters . . . 83
4.4 Identification of Foundation Parameters . . . 88
4.4.1 Macro-element Parameterisations and Prior Distributions . . . 88
4.4.2 Modal Data: Mode Set Selection . . . 90
4.4.3 Results and Discussion . . . 93
4.5 Summary . . . 100

CHAPTER 5: A Substructuring Approach for Input-State Estimation

5.1 Introduction . . . 102
5.2 Mathematical Modelling of Substructures . . . 104
5.2.1 Reduced-Order Model . . . 105
5.2.2 State space model . . . 108
5.3 Augmented Kalman Filter for Input-State Estimation . . . 110
5.4 Performance Evaluation using Numerical Case Studies . . . 112
5.4.1 Synthetic data generation . . . 112
5.4.2 Linear finite element model . . . 114
5.4.3 Comparison metrics . . . 117
5.4.4 Case study 1: wind loading only . . . 119
5.4.5 Case Study 2: wind and wave loading . . . 133
5.5 Summary . . . 142

CHAPTER 6: A Substructuring Approach for Parameter Estimation

6.1 Introduction . . . 146
6.2 Time-Domain Bayesian Inference with Substructuring . . . 147
6.3 Numerical Example . . . 149
6.3.1 Wind Turbine Model for Data Generation . . . 149
6.3.2 Tower-Foundation Substructure Model . . . 152
6.3.3 Identification of the foundation parameters . . . 156
6.4 Summary . . . 162

CHAPTER 7: Real-World Case Study

7.1 Introduction . . . 165
7.2 Field Data . . . 166
7.2.1 Operational Modal Analysis . . . 168
7.3 Numerical Models . . . 175
7.3.1 Foundation Model . . . 176
7.3.2 Model for Method 1: Simulink Model . . . 177
7.4 Application of Methods to Field Data . . . 179
7.4.1 Method 1 Results . . . 180
7.4.2 Method 2 Results . . . 186
7.4.3 Discussion . . . 193
7.5 Summary . . . 195

CHAPTER 8: Conclusions

8.1 Key Contributions from Work Packages 197
8.2 Future Work 200
8.2.1 Nonlinear and dynamic soil modelling 200
8.2.2 Experimental validation 201
8.2.3 Optimal sensor placement studies 201
8.2.4 Damage detection and changing structural behaviour 202
Bibliography 202
APPENDIX A: Comparison of Model-Predicted Modal Parameters for Chapter 4 . 211
APPENDIX B: Tabulated Performance Metrics for Case Study 1 from Chapter 5 . 214

List of Figures

1.1 Typical wind and wave loading spectra, showing the rotor rotation frequency range (1P), blade passing frequency range (3P), and first natural frequency (f_n) for three reference turbines: the NREL 5 MW [8], the IEA 10 MW [9], and the IEA 15 MW [10]. 2
1.2 Simplified soil-structure interaction models. (a) Apparent fixity model (b) Macro-element foundation model. 7
1.3 Example of relationship between (a) k_H and h (b) k_M and h (c) k_H^{-1} and h, for a fixed macro-element foundation. The salient features shown on each curve are determined by individual macro-element compliance parameters (c_{11}, c_{22} and c_{12}). 10
1.4 Example highlighting the challenge of identifying closely spaced modes: the identified mode shape lies within the subspace spanned by the two true model mode shapes. 16
1.5 Summary of work packages (WPs) and research questions. Abbreviations: SSI = soil-structure interaction; RNA = rotor nacelle assembly. 24

2.1 Offshore wind turbine model and simulation requirements for this thesis. . . . 28
2.2 Macro-element foundation representation. 30
2.3 Simulink implementation of the macro-element foundation. The x and y degrees of freedom (DOFs) are aligned in the fore-aft and side-side directions, receptively. The z DOF is aligned in the vertical direction. 31
2.4 Rotor-Nacelle Assembly (RNA) 33
2.5 Example of blade model sub-assembly. Each blade is defined by a user-specified number of *Reduced-Order Flexible Solid* blocks. In this example, the blade is modelled using ten blocks, with one block highlighted for illustration. 34
2.6 Schematic showing how the aerodynamics module interacts with the Simscape strucutral model. 35
2.7 Example JONSWAP Spectrum 38
2.8 Full-system integration of the Simulink model for a monopile-supported offshore wind turbine with a macro-element foundation. 40

3.1 (a) Overview of the 15 MW offshore wind turbine and accelerometer locations (b) Macro-element foundation representation. 56
3.2 Example time histories of: (a) Kaimal hub-height wind speed (b) JONSWAP spectrum wave surface elevation. 57
3.3 Example set of acceleration measurements. From left to right: fore-aft time domain response (200 second excerpt), side-side time domain response (200 second excerpt), and frequency domain response. Black lines show the FA response and red lines show the SS response. (a) Accelerometer at 175 m (b) Accelerometer at 123 m (c) Accelerometer at 97 m (d) Accelerometer at 45 m. 58
3.4 Example stabilisation diagrams from OMA on a representative dataset. Left: all stable modes based on criteria of 1 % frequency difference and a MAC value greater than 0.97 for consecutive modes. Right: filtered stabilisation diagram showing only the modes and selected model orders used for model updating. . 60

3.5 True model mode shapes for the six modes selected for model updating. Each row represents a separate mode. Top row shows Fore-Aft (FA) view, middle row shows Side-Side (SS) view, bottom row shows FA-SS view. 60
3.6 Comparison between the identified modes and true model-predicted modes across all 100 datasets. (a) Frequency errors (b) MAC Values. 61
3.7 (a) Modal parameter error vectors (b) Comparison of modal parameter standard deviations from SSI-cov and frequentist estimates. Each point corresponds to one modal parameters (i.e. one modal frequency or mode shape component). The horizontal axis shows the standard deviation estimated using the frequentist approach. The vertical axis shows the average of the 100 standard deviations estimated by SSI-cov, with blue error bars indicating the 95 % confidence interval of those 100 values. The black line indicates the 1:1 reference. (c) Magnified version of Figure 3.7b for small standard deviations. 63
3.8 Updated compliance estimates, normalised by their true values. Comparison between all three model updating frameworks where the error bars indicate a 95 % confidence interval for the Bayesian frameworks. The corresponding % errors from averaging across 100 individual datasets (labelled 'Avg. (100 runs)') are compared to the frequentist approach (labelled 'Freq.') in bar plots on the right. 66
3.9 Updated compliance estimates when using the deterministic framework with different initial guesses. The % errors from averaging across 100 individual datasets (labelled 'Avg. (100 runs)') are compared to the frequentist approach (labelled 'Freq.') in bar plots. 68
3.10 Posterior distributions for the **vector-based likelihood function** when the measurement uncertainty was undefined/neglected and the modal parameters' average from 100 datasets were used for model updating. 69
3.11 Posterior distributions for the **MAC-based likelihood function** when measurement uncertainty was undefined/neglected and the modal parameter average across 100 datasets was used for model updating. 69
3.12 Histograms from propagating uncertainty from mode shape components to MAC values along with the optimised truncated normal for each mode. . . . 73
3.13 Most probable estimates for the compliance parameters for all methods. 'Avg. 100 runs' refers to the case when the most probable estimates are averaged across 100 individual datasets. 74
3.14 Model evidence and model error parameter MAP estimates for the Bayesian methods. 'Avg. 100 runs' refers to the case when the most probable estimates are averaged across 100 individual datasets. (a) Vector-based likelihood funciton (b) MAC-based likelihood function. 75

4.1 Baseline wind turbine model, reproduced from Chapter 3. (a) Baseline IEA 15-MW offshore wind turbine and accelerometer locations (b) Macro-element foundation representation. 81
4.2 Blade property distributions for baseline and up-scaled model. 84
4.3 Comparison of frequencies and mode shapes between the baseline turbine and turbines with the up-scaled blade models. Modal parameters were computed using the baseline macro-element foundation $\mathbf{K}^0_{\text{macro}}$. 86
4.4 Tower mode shapes for turbines with different blade models. The foundation is defined using $\mathbf{K}^0_{\text{macro}}$. 87

4.5 Prior distribution for γ-$\boldsymbol{d}$ macro-element parameterisation in the (a) original γ-$\boldsymbol{d}$ space and (b) mapped to the compliance space. 91
4.6 Averaged fore-aft and side-side accelerometer power spectral densities (PSDs) from a turbulent forward simulation. The modes selected for model updating are indicated by the vertical lines and labelled corresponding to Figures 4.3 and 4.4. 93
4.7 Posterior distributions obtained when using the models defined by the scaling methods and reference turbines listed in Table 4.2 and mode sets defined in Table 4.4. 95
4.8 Trend between model evidence and average MAP estimate error in the compliance parameters. The dashed lines show the line of best fit for each mode set. (a) Compliance parameterisation (b) γ-$\boldsymbol{d}$ parameterisation. 99

5.1 Overview of substructuring method for an offshore wind turbine. The structure is 'cut' at the dashed lines, which introduces three substructures: the Rotor Nacelle Assembly (RNA), the tower, and the foundation. This figure provides a simplified 2D representation, whereas the model used in this study is 3D. . . 103
5.2 Simulink model setup used for generating synthetic data 113
5.3 (a) Linear finite element model with interface forces and moments used for input-state estimation, (b) 'Sparse: θ-a' sensor setup. 115
5.4 Graphical representation of the rigid body modes and retained normal modes for the tower substructure. Each column represents an individual mode, with the top and bottom rows displaying the projections in the fore-aft (FA) and side-side (SS) directions, respectively. 116
5.5 Comparison between actual response and AKF estimated response at selected DOFs for the three sensor setups (Case Study 1). *From left to right:* entire response duration, magnified excerpt, $TRAC/FRAC/MAE$ values. 121
5.6 Estimated modal displacements when using 'Dense: θ-a' sensor setup (red-dashed lines) vs actual response (solid-black line). Sub-figures (a)-(d): *rigid body modes*, sub-figures (e)-(h): *normal modes*, sub-figures (i)-(p): *attachment modes*. 123
5.7 Accumulated damage at each node in the finite-element model (Left figure: fore-aft damage distribution. Right figure: side-side damage distribution) . . . 125
5.8 Shear forces and bending moments at the **ground level interface** for the three sensor setups (Case Study 1). *From left to right:* entire response duration, magnified excerpt, $TRAC/FRAC/MAE$ values. 126
5.9 Shear forces and bending moments at the **tower top interface** for the three sensor setups (Case Study 1). *From left to right:* entire response duration, magnified excerpt, $TRAC/FRAC/MAE$ values. 127
5.10 Results for input estimation with synthetic data that excludes gravity forces (Case Study 1) (a) Fore-aft shear force at the tower top interface: $\mathtt{F175}_{\mathtt{FA}}$. Comparison between actual response and AKF-estimated response when using the 'Dense: θ-a' sensor setup. (b) MAEs in estimated inputs when the measured data do not include gravity forces, compared with the original results where gravity forces are included. Both sets of results are obtained using the 'Dense: θ-a' sensor setup. 129

5.11 Estimated macro-element parameters from the 'Dense: x-a' sensor setup. The annotated percentage errors represent the average over the final 300 seconds. (a) Lateral compliance c_{11}; (b) Rotational compliance c_{22}; (c) Cross-coupling compliance c_{12}. Compliance values are normalised by their true values: c_{11}^0, c_{22}^0, and c_{12}^0. 131
5.12 Estimated macro-element parameters from the sensor setups without displacements (i.e., only inclinometers and accelerometers). (a) Lateral compliance c_{11}; (b) Rotational compliance c_{22}; (c) Cross-coupling compliance c_{12}. Compliance values are normalised by their true values: c_{11}^0, c_{22}^0, and c_{12}^0. 132
5.13 Estimated macro-element parameters from the 'Sparse: θ-a' sensor setup when estimated displacements are ignored. The annotated percentage errors represent the average over the final 300 seconds. (a) Rotational compliance c_{22} (b) Cross-coupling compliance c_{12}. Compliance values are normalised by their true values: c_{22}^0, and c_{12}^0. 133
5.14 Triangular wave loading approximation 134
5.15 Case Study 2 damage distributions. (solid black line: actual damage, blue dash-dotted line: damage when no wave inputs were considered, red dashed line: damage when triangular approximation was used). 136
5.16 Ground level interface input estimation for Case Study 2 (wind and wave loading). For each sub-figure, left to right shows: entire response duration, 10 s magnified excerpt, power spectral density. Top row shows results for wave yaw = 0° and the bottom row shows wave yaw = 30°. 137
5.17 Time series excerpts of estimated triangular wave inputs ($\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$) from simulations with a range of wind-wave misalignment angles. 139
5.18 Illustration showing the wind and wave directions, and how the resultant of the estimated wave inputs $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$ relates to the wave misalignment angle. 139
5.19 Estimated wave misalignment angles from simulations with varying wind-wave misalignment, using (a) the Kalman filtering approach and (b) the Arias Intensity tensor approach. In (a), the solid lines indicate mean values for α_w, and the shaded area represents the 95 % confidence interval. In (b), the direction of the major axis aligns with the wave misalignment angle and the subplots from top to bottom correspond to simulations with true misalignment angles of 90°, 60°, 30°, and 0°, respectively. 142

6.1 (a) Finite element (FE) model used for data generation. (b) FE model for the tower-foundation substructure and locations of sensors. The FE model is 3D, constructed with 1D beam elements, but the figure provides a 2D schematic representation in the fore–aft plane. The side–side direction is perpendicular to the figure (into the page). 150
6.2 (a) Tower top wind speed defined by a Kaimal spectrum. (b) Wave surface elevation defined by a JONSWAP spectrum. 152
6.3 Distribution of strain gauges (SGs) at the transition piece. Abbreviations: FA = fore-aft, SS = side-side. 152
6.4 Example of retained normal modes computed using $\mathbf{K}_{\text{macro}}^0$. $\tilde{z}$ denotes the normalised height, where $\tilde{z} = 0$ corresponds to the mudline and $\tilde{z} = 1$ to the tower top. x and y are aligned with the fore-aft and side-side directions, respectively. 155
6.5 Posterior distribution for compliance parameterisation. Parameters are normalised by their true values. Probability density functions are normalised by the maximum value. 158

6.6 Posterior distribution for γ. The probability density function is normalised by its maximum value. 159
6.7 Influence of the underlying foundation stiffness on the identifiability of γ. . . 160
6.8 Posterior estimates for γ obtained from data with different parameterisations of the internal water level ζ_i, external water level ζ_e and hydrodynamic added mass coefficient C_a. The solid line indicates the *Maximum a Posteriori* estimates and the shaded region represents the 95 % confidence interval. (a) Results when ζ_e is varied and ζ_i is fixed equal to the mean water level. (b) Results when $\zeta_i = \zeta_e$. 163

7.1 (a) Wind turbine geometry and tower sensor positions. (b) Transition piece strain gauges: SG-TP (c) Transition piece accelerometer: ACC-TP (d) Tower top accelerometer: ACC-TT (e) Tower-middle accelerometer: ACC-TM 167
7.2 Ten-minute averages and corresponding ± 2 standard deviations of nacelle yaw angles from the SCADA system. 169
7.3 (a) Example set of cleaned clusters from a representative dataset—grey dots show rejected modes and colours indicate common clusters. (b) Example set of selected clusters and representative modes used for model updating. 173
7.4 Natural frequencies selected stable modes over idling period. (a) All selected modes. (b) Modes SS2 and FA2. (c) Modes FA1 and SS1. 174
7.5 Averaged identified modal parameters over the idling period $\pm$ one coefficient of variation. Height is the normalised distance from mudline to tower top (mudline = 0, tower top = 1). (a) Frequencies normalised by average value for each mode. (b)-(g) Mode shapes projected in the fore-aft (ϕ_{FA}) and side-side (ϕ_{SS}) directions. 174
7.6 Foundation models 176
7.7 Posterior distributions for macro-element compliances and model error parameters from updating with the **averaged idling dataset**. All units are dimensionless: compliances are normalised by their initial design values and probability densities are normalised by the *maximum a posteriori* estimate. 181
7.8 MAP estimates and 95 % confidence intervals for macro-element compliances and model error parameters when updating with **individual ten-minute datasets** across the idling period. All units are dimensionless: compliances are normalised by their initial design values. 182
7.9 Posterior distributions for macro-element γ scaling factor and model error parameters when updating with the **averaged idling dataset**. All units are dimensionless: γ is relative to the initial design and probability densities are normalised by the *maximum a posteriori* estimate. 183
7.10 MAP estimates and 95 % confidence intervals for macro-element γ scaling factor and model error parameters when updating with **individual ten-minute datasets** across the idling period. All units are dimensionless: γ is relative to the initial design. 184
7.11 Comparison of model predictions for Method 1. Error bars/shaded areas indicate $\pm$ one coefficient of variation in the measured values. (a) Frequencies normalised by averaged idling measured frequency for each mode. (b)-(g) Mode shapes. $\tilde{z} = 0$ = mudline; $\tilde{z} = 1$ = tower top. *Solid black: Measured; Solid red: Initial model; Dashed blue: Compliance MAP; Dash-dot green: γ MAP.* 185

7.12 Example posterior distribution for the compliance parameters obtained using a representative ten-minute dataset. All units are dimensionless: compliances are normalised by their initial design values and probability densities are normalised by the *maximum a posteriori* estimate 187
7.13 Posterior distributions for soil parameter γ, obtained from Method 2. γ is relative to the initial design. (a) Example distribution from a representative ten-minute dataset. (b) MAP estimates and corresponding 95% confidence interval across idling period. 188
7.14 Estimated tower-top input forces from three parameterisations of the macro-element foundations: initial design (solid black), compliance MAP estimated (dashed blue), and γ MAP estimated (dotted red). (a) Normalised East–West force; (b) Normalised South–North force. 189
7.15 Excerpts of estimated ground-level states, internal shear forces and bending moments. Units have been removed for confidentiality reasons. *Solid black: Initial design; Dashed red: Compliance MAP; Dash-dot blue: γ MAP*. 190
7.16 Excerpts of recreated measurements compared with sensor data. Units are hidden for confidentiality reasons. *Legend: Solid black = Measured; Dashed red = AKF Reconstructed; Dash-dot blue = Forward Simulation*. 192
7.17 Comparison of mean absolute errors between measured sensor data and corresponding responses from Forward Simulations driven by estimated inputs: (a) accelerometer measurements; (b) strain gauge measurements. Acceleration and strain units are anonymised. 193
7.18 Comparison of methods for foundation parameter estimation. Blue: marginal distribution obtained using Method 1 with the averaged idling dataset. Red: frequentist distribution of MAP estimates obtained from Method 2. 194

List of Tables

2.1 Simscape Multibody blocks used to model the offshore wind turbine structural mechanics. 30
2.2 Application of the Simulink model throughout this thesis. 43

3.1 Lower and upper bounds for each updated parameter in the Bayesian frameworks. 64

4.1 Blade property distributions used for flexible blade modelling, with corresponding units and scale factors. 81
4.2 Definition of up-scaled models used for comparison with the baseline model to investigate the effects of blade modelling error. 83
4.3 Prior distributions assigned to each parameter for both foundation parameterisations. $\mathcal{U}(a,b)$ denotes a uniform distribution with lower bound a and upper bound b. $\mathcal{N}(\mu,\sigma)$ denotes a normal distribution with mean μ and standard deviation σ. 90
4.4 Mode sets used for model updating to identify the foundation parameters. . . 92

5.1 List of user-defined inputs required to model the tower substructure. 114
5.2 Summary of sensor configurations considered in the numerical case studies. Note that at each sensor location, the response in both the fore-aft and side-side directions are used in the measurement set. 116

6.1 Lower and upper bounds used to define the uniform prior distributions for all parameter. The initial design compliance values (c_{11}^0, c_{22}^0, and c_{12}^0) are given in Equation 6.11. 157
6.2 *Maxima a Posteriori* (MAP) estimates for γ for different assumed errors in the baseline compliance matrix $\mathbf{C}_{\text{macro}}^0$. 159
6.3 Influence of the underlying foundation stiffness on the *Maximum a Posteriori* (MAP) estimates for γ. 161

7.1 Lower and upper bounds of the uniform prior distributions applied to macro-element foundation parameters and the model error parameters within likelihood function (Equation 3.11). The initial design compliance values, c_{11}^0, c_{22}^0 and c_{12}^0, are provided in Equation 6.11. 179
7.2 Comparison of model-predicted and measured modal parameters for *maximum a posteriori* (MAP) foundation estimates. Results correspond to those obtained from the averaged idling dataset and consider two parameterisations: individual compliances or the γ scaling factor. 185
7.3 MAP estimates for each compliance parameter obtained with Method 2 from a representative ten-minute dataset. These estimates were used to the produce the input-state estimation results presented in Section 7.4.2 (iii). Normalised compliance values are relative to the initial design. 188

A.1 Comparison of MAP model-predicted modal parameters from each blade-scaled model (see Table 4.2) against those from the baseline model, based on updating with the **full mode-set**. 212

A.2 Comparison of MAP model-predicted modal parameters from each blade-scaled model (see Table 4.2) against those from the baseline model, based on updating with the **reduced mode-set**. 213

B.1 $TRAC$, $FRAC$, and MAE for selected degrees of freedom in the fore-aft and side-side directions from Case Study 1 in Chapter 5. 215

CHAPTER 1

Introduction and Literature Review

1.1 Background and Motivation

There is an urgent need to decarbonise our global energy mix. Between 2011 and 2020, global surface temperature increased by 1.1°C above pre-industrial levels, primarily as a result of human activities [1]. The impacts of climate change are already evident, with increasing frequency and severity of extreme weather events, such as heatwaves, wildfires, heavy precipitation, droughts and tropical cyclones. The continued accumulation of greenhouse gases threatens human health, food security and natural ecosystems [2].

To reduce carbon dioxide emissions and meet net-zero targets, many nations are relying on wind energy, particularly offshore wind. For example, the UK has set a target of achieving an installed capacity of 50 GW of offshore wind by 2030 [3]. Compared to their onshore counterparts, offshore wind turbines benefit from stronger and more consistent winds, lower visual and noise impacts, and the potential to accommodate larger turbines sizes. At the end of 2024, global offshore wind capacity reached 83 GW, with projections indicating an increase to 441 GW by 2034 [4].

Alongside this rapid growth, the offshore wind industry has experienced a significant increase in the average turbine size over the past decade. In Europe, the average rated power of newly installed offshore wind turbines increased from 4.2 MW in 2015 to 10.1 MW in 2024 [5], with 15 MW models now entering commercial deployment. While this increase in size has been a significant driver in reducing the levelised cost of energy (LCOE) [6], it also introduces heightened structural challenges. Fixed-bottom offshore wind turbines are usually designed as

soft-stiff structures [7], where their first natural frequency is positioned between the rotor rotation frequency (1P) and the blade passing frequency (3P), as shown in Figure 1.1. This design criterion helps avoid resonance from loads caused by rotor imbalance (1P) and aerodynamic forces from each blade passing the tower (3P). Figure 1.1 provides typical frequency ranges for three reference turbines of increasing size—the NREL 5 MW [8], the IEA 10 MW [9] and the IEA 15 MW [10] reference turbines—showing how larger turbines are associated with lower natural frequencies, which shift further into the higher-energy regions of the wind and wave loading spectra. This effect, coupled with the increased regularity of extreme weather events, raises concerns about fatigue accumulation and highlights the need for a deeper understanding of the dynamic behaviour of offshore wind turbines. Furthermore, deployment in deeper waters and the use of larger support structures introduce additional structural demands and associated risks.

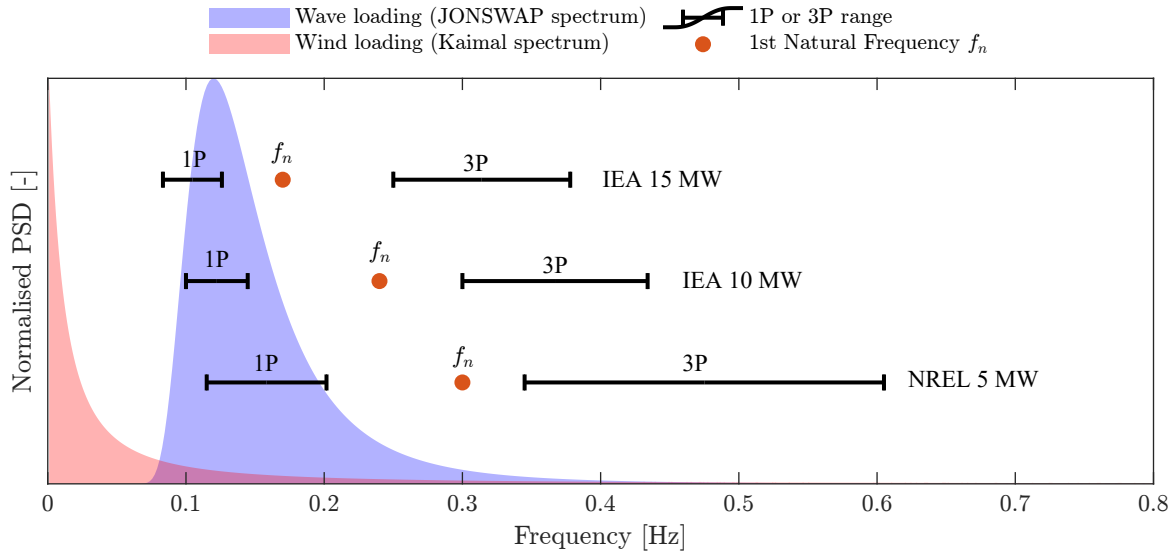


Figure 1.1: Typical wind and wave loading spectra, showing the rotor rotation frequency range (1P), blade passing frequency range (3P), and first natural frequency (f_n) for three reference turbines: the NREL 5 MW [8], the IEA 10 MW [9], and the IEA 15 MW [10].

To help mitigate these risks, structural health monitoring techniques can be utilised, whereby sensors are installed on the structure and the data is continuously processed and analysed to monitor the condition of the *in-situ* structure [11]. In some structural health monitoring methods, the condition of a structure is assessed solely on the measured data. For example, modal parameters identified from vibration data can be tracked over time, where changes may indicate shifts in stiffness or mass distribution due to damage or degradation [12].

While such data-driven methods have demonstrated significant value in civil engineering applications, their effectiveness for offshore wind turbines can be limited by the spare deployment of sensors. Mainly owing to cost and access constraints, only a small number of sensors (typically accelerometers, strain gauges or inclinometers) are installed along the tower and foundation. As a result, it can be difficult to confidently detect or localise structural changes at unmeasured yet critical locations, such as within the subsoil or submerged portions of a monopile foundation. To address these limitation, structural health monitoring approaches that integrate physics-based models with measured data offer a promising alternative. By combining the predictive capabilities of numerical models with *in-situ* monitoring data, this hybrid approach can provide a more complete picture of the structure's condition. This concept has recently evolved into a broader framework of digital twins, whereby a virtual model is used to represent an offshore wind turbine structure and can provide valuable insight into the structure's current and future condition, as well as potential risks [13]. These insights can inform design improvements and help reassess remaining useful life [14, 15]. However, digital twins must be updated to match monitored data to reflect the as-built structure, rather than the as-designed model.

A key source of discrepancy between the two is the challenge of accurately modelling the soil-structure interaction [16, 17], which is further complicated by the variability and limited availability of *in-situ* soil data from site investigations [18]. For example, Kallehave *et al.* [7] found that the first natural frequency predicted by the as-designed model was consistently lower than values estimated from monitored data. Therefore, a greater understanding of the soil-structure interaction is essential to improve the accuracy of numerical models and the reliability of digital twins [19, 20].

This uncertainty in the soil-structure interaction formed the basis of the doctoral research preceding this thesis by Dr Emily Anderson [21], which demonstrated the feasibility of identifying foundation parameters using structural health monitoring techniques within a deterministic framework. This work highlighted and addressed challenges specific to monopile-supported offshore wind turbines. However, a key limitation acknowledged by Anderson was the lack of uncertainty quantification in the estimated parameters within the deterministic setting, limiting confidence in the results and their use in decision-making. Therefore, there is motivation

to extend this work by investigating and developing structural health monitoring methods to identify foundation parameters in a probabilistic framework.

An additional challenge when using offshore wind turbine numerical models for structural health monitoring is the lack of detailed design information for the rotor nacelle assembly (RNA). In practice, Original Equipment Manufacturers (OEMs) often withhold certain design parameters—such as blade properties—which can significantly influence the dynamic behaviour of an offshore wind turbine. **Therefore, the central motivation for this thesis lies in advancing structural health monitoring methods that combine physics-informed models with vibration-based measurements to identify foundation parameters and quantify their uncertainty, while accounting for potential modelling errors introduced by limited knowledge of the RNA.** It should be noted that this research only considers offshore wind turbines supported by monopile foundations, as this foundation type forms the majority of installations—in 2020, 80.5 % of all turbines installed in Europe were supported with monopile foundations [22].

The remainder of this chapter presents a review of the relevant literature, followed by an outline of the main research objectives and the overall structure of the thesis.

1.2 Literature Review

1.2.1 Offshore wind turbine modelling

Numerical models are often employed to predict the response of offshore wind turbines. However to do so, it is important to consider all aspects of the structure that contribute to its dynamics. For example, Anderson [21] showed that the flexibility of the blades results in a strong interaction with the tower, leading to different dynamics compared with a rigid rotor model. This highlights the importance of using an accurate rotor model and suggests that simplifying the rotor nacelle assembly will result in lost detail. The blade-tower interaction has also been investigated by Shkara *et al.* [23], where it is shown how loading is induced on the structure due to this coupling. Furthermore, the controller and the hydrodynamic loading can also affect the dynamic response. Therefore, accurately modelling offshore wind turbines requires a fully coupled servo-hydro-aero-elastic model.

Many models have been developed to achieve this. Three of the most popular models include OpenFast [24], HAWC2 [25] and Bladed [26]. In the Offshore Code Comparison Collaboration (OC3) project [27], these three simulation tools, as well as many others, were used to provide a comparison of offshore wind turbines with various foundation types. One conclusion from OC3 was that all codes performed similarly for the case where the structure is modelled with a flexible monopile foundation.

More recently (and not included in OC3), the Non-Linear Dynamics Group at the University of Oxford developed a wind turbine model in Simulink for floating wind applications [28, 29]. After reviewing the OpenFast and HAWC2 models, for this thesis it was decided to adopt this Simulink wind turbine model and extend its application to a fixed-bottom monopile foundation. The model's key features and implementation details are discussed in Chapter 2.

(i) Soil-structure interaction modelling

The soil-structure interaction plays a critical role in governing the dynamic behaviour of offshore wind turbines. For monopile-supported turbines, the axial loading is often small relative to the pile's axial capacity, and as a result, the lateral response is usually of more interest. Accurately modelling this interaction is extremely challenging [18], and a range of modelling approaches have been proposed. Each method involves trade-offs in terms of accuracy, computational cost and applicability to specific engineering tasks.

From a geotechnical perspective, high-fidelity 3D constitutive models are considered the most accurate approach, where the soil continuum is composed of finite elements that embed complex constitutive laws to capture non-linear, anisotropic, and rate-dependent soil behaviour [30, 31]. While they are capable of simulating realistic soil behaviour with high accuracy, they require a significant number of parameters to define and are computationally expensive. Therefore, they are generally impractical for inverse methods, such as the task of parameter identification considered in this work.

To reduce computation cost in dynamic simulation, a commonly adopted alternative is the use of distributed spring models. In these approaches, the soil is represented by non-linear springs distributed along the embedded length of the pile. Traditional design standards for

laterally loaded piles, such as those from the American Petroleum Institute (API) and DNV, use the p-y method to define the spring stiffnesses [32, 33]. In this method, only the distributed lateral soil reaction is considered to determine the uncoupled lateral springs that simulate forces.

The concept of the p-y method was first introduced by Reese & Matlock [34] and later supported by experimental work from Matlock [35], who validated the assumption of uncoupled springs. Cox *et al.* [36] then carried out lateral load tests on steel tube piles embedded in saturated sand, with a diameter of $D = 0.61\,\mathrm{m}$ and an embedment length of $L = 21\,\mathrm{m}$, corresponding to a length-to-diameter ratio $L/D = 34$. These results allowed Reese *et al.* [37] to empirically derive the p-y curves.

However, the p-y method was developed primarily for the offshore oil and gas industry, where piles are relatively slender compared with the offshore wind industry, which often have L/D ratios below 6. Applying the p-y method to these shorter, stiffer piles can lead to inaccuracies [38]. In particular, it tends to underestimate the soil stiffness, leading to overly conservative design estimates and higher construction costs [7].

To address these limitations, the Pile Soil Analysis (PISA) project was initiated to develop improved design methods tailored to offshore wind applications [39–41]. The PISA model retains the Winkler-style framework of the p-y method but extends it by including additional soil reaction components that become important for monopiles with low L/D ratios. Specifically this incorporates: a distributed lateral load $p(z, v)$, a distributed moment $m(z, \psi)$, a horizontal base force $H_B(\psi_B)$, and a base moment $M_B(\psi_B)$. Here, z is the depth below ground, v is the lateral displacement, ψ is the cross-sectional rotation, and the subscript $_B$ dontes quantities at the monopile base. Each of these components is represented by a normalised soil reaction curve using a four parameter conic function:

$$-n\left(\frac{\overline{y}}{\overline{y}_u}-\frac{\overline{x}}{\overline{x}_u}\right)^2+(1-n)\left(\frac{\overline{y}}{\overline{y}_u}-\frac{\overline{x}k}{\overline{y}_u}\right)\left(\frac{\overline{y}}{\overline{y}_u}-1\right)=0, \tag{1.1}$$

where $\overline{x}$ is the normalised displacement or rotation, and $\overline{y}$ is the corresponding normalised load or moment. The other four parameters of this function—ultimate displacement $\overline{x}_u$, ultimate load $\overline{y}_u$, initial stiffness k, and curvature parameter n—are depth dependent and were calibrated

against high-fidelity 3D finite element analyses in the PISA project [42, 43]. These analyses were in turn validated against field tests [44–47].

Compared to the traditional p-y method, the results from the PISA project showed a significant reduction in the discrepancy between design and measurements for monotonic lateral and moment loading. However, additional research is needed to address the PISA design method for cyclic loading [48] and to accommodate layered soil profiles [49]. Building on the PISA project, the ongoing PICASO project was launched in 2018 to further develop methods for modelling the cyclic behaviour of monopiles for offshore wind turbines [50].

As an alternative to these distributed spring models, several simplified models have been proposed that reduce the number of modelling parameters and further improve computational efficiency. Among the most common are the apparent fixity method and the macro-element foundation, which are illustrated in Figure 1.2.

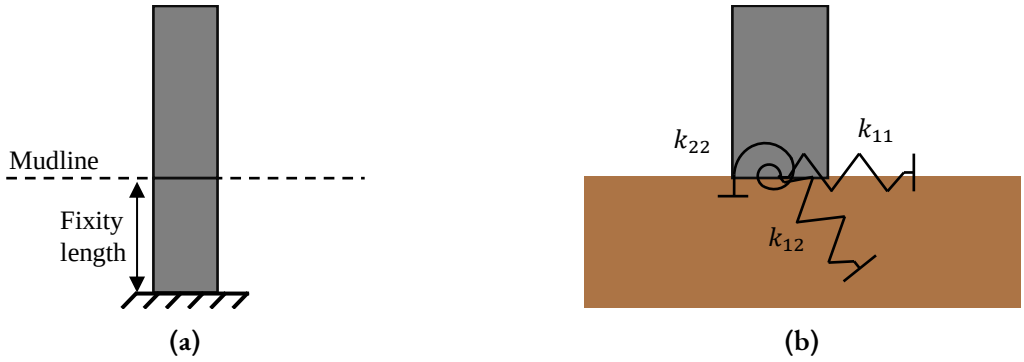


Figure 1.2: Simplified soil-structure interaction models. (a) Apparent fixity model (b) Macro-element foundation model.

In the apparent fixity method, the lateral flexibility of the monopile–soil system is idealised as an equivalent cantilever beam embedded in rigid soil. The fixity depth and fictitious beam properties are calibrated so that the lateral and rotational response at the mudline matches that of the full monopile-soil system. Løken & Kaynia [51] applied this foundation model to simulate the dynamic response of an offshore wind turbine and estimate fatigue, concluding that this simplified model can overpredict loads and lead to conservative designs. Nonetheless, the authors highlight the important role that soil-structure interaction plays in dynamic analysis and fatigue estimation.

The macro-element replaces the embedded foundation with a set linear-elastic springs introduced at the mudline: a translation spring k_{11}, a rotational spring k_{22}, and a cross-coupling spring k_{12}. These springs relate the ground-level displacement to the applied loads through Equation 1.2:

$$\begin{bmatrix} H_G \\ M_G \end{bmatrix} = \begin{bmatrix} k_{11} & k_{12} \\ k_{12} & k_{22} \end{bmatrix} \begin{bmatrix} v_G \\ \psi_G \end{bmatrix}. \tag{1.2}$$

Here, H_G and M_G are the ground-level horizontal loading and overturning moment, respectively, and v_G and ψ_G are the ground level displacement and rotation. The axial and torsional degrees of freedom can also be included; however, for monopile-supported offshore wind turbines, they are typically much stiffer than the lateral and bending directions, so can often be neglected. Similar to the apparent fixity method, the macro-element spring stiffnesses (k_{11}, k_{22} and k_{12}) can be calibrated to reproduce the response of a more detailed soil–structure interaction model, such as one based on distributed springs or finite elements.

Building on the computational efficiency of the macro-element, the REDWIN project [52] developed a more advanced formulation that incorporates key aspects of realistic soil behaviour, including plasticity, kinematic hardening, and hysteresis. The REDWIN macro-element requires, as primary inputs, an elastic stiffness matrix at the seabed and a nonlinear load–displacement curve obtained obtained from a pushover analysis. Its performance in dynamic simulations was demonstrated by Page *et al.* [53] and compared against the traditional p-y approach, showing improved accuracy in capturing different stiffness during loading and unloading cycles, as well as foundation damping dependent on the load history. Furthermore, Phase II of the Offshore Code Comparison, Collaboration, Continued with Correlation and unCertainty (OC6) project [54] provides a comprehensive benchmark of various soil–structure interaction models—including the apparent fixity method, coupled-springs macro-element, distributed springs, and the REDWIN model—by comparing results from multiple participants using different offshore wind turbine simulation tools.

One aim of this thesis is to identify the parameters of the macro-element foundation model (i.e., k_{11}, k_{22} and k_{12} in Equation 1.2). However, Anderson [21] highlighted the difficultly in comparing sets of stiffness parameters, as changes in k_{11}, k_{22}, and k_{12} can often compensate for

one another, resulting in minimal change in the overall foundation response. To address this, McAdam *et al.* [19] derived an equivalent uncoupled macro-element formulation, defined by a lateral stiffness k_H and a rotational stiffness k_M, where:

$$\begin{bmatrix} H_G \\ M_G \end{bmatrix} = \begin{bmatrix} k_H & 0 \\ 0 & k_M \end{bmatrix} \begin{bmatrix} v_G \\ \psi_G \end{bmatrix} \tag{1.3}$$

This transformation is achieved by introducing the loading eccentricity h, defined by the ratio between the ground-level moment M_G and the horizontal force H_G:

$$h = \frac{M_G}{H_G}. \tag{1.4}$$

It can then be shown that the uncoupled stiffnesses are related to the original coupled parameters as:

$$\begin{aligned} k_H &= \frac{k_{11}k_{22} - k_{12}^2}{k_{22} - hk_{12}} \\ k_M &= \frac{h\left(k_{11}k_{22} - k_{12}^2\right)}{hk_{11} - k_{12}}. \end{aligned} \tag{1.5}$$

For a given set of k_{11}, k_{22} and k_{12}, the variation of k_H, k_M and k_H^{-1} with respect to h can be plotted (see Figure 1.3). This graphical representation provides a clear and more interpretable comparison between different macro-element formulations.

Additionally, the salient features shown in Figure 1.3—the intercept of k_H, the asymptote of k_M, and the gradient of k_H^{-1}—are governed by an individual component of the compliance matrix, defined as the inverse of the stiffness matrix:

$$\begin{bmatrix} c_{11} & c_{12} \\ c_{12} & c_{22} \end{bmatrix} = \begin{bmatrix} k_{11} & k_{12} \\ k_{12} & k_{22} \end{bmatrix}^{-1} = \frac{1}{k_{11}k_{22} - k_{12}^2} \begin{bmatrix} k_{22} & -k_{12} \\ -k_{12} & k_{11} \end{bmatrix}. \tag{1.6}$$

Therefore, the compliance parameters (c_{11}, c_{22} and c_{12}) are considered to be more robust metrics for comparing foundation behaviour and have been adopted in recent studies as optimisation variables for parameter identification [19, 21, 55].

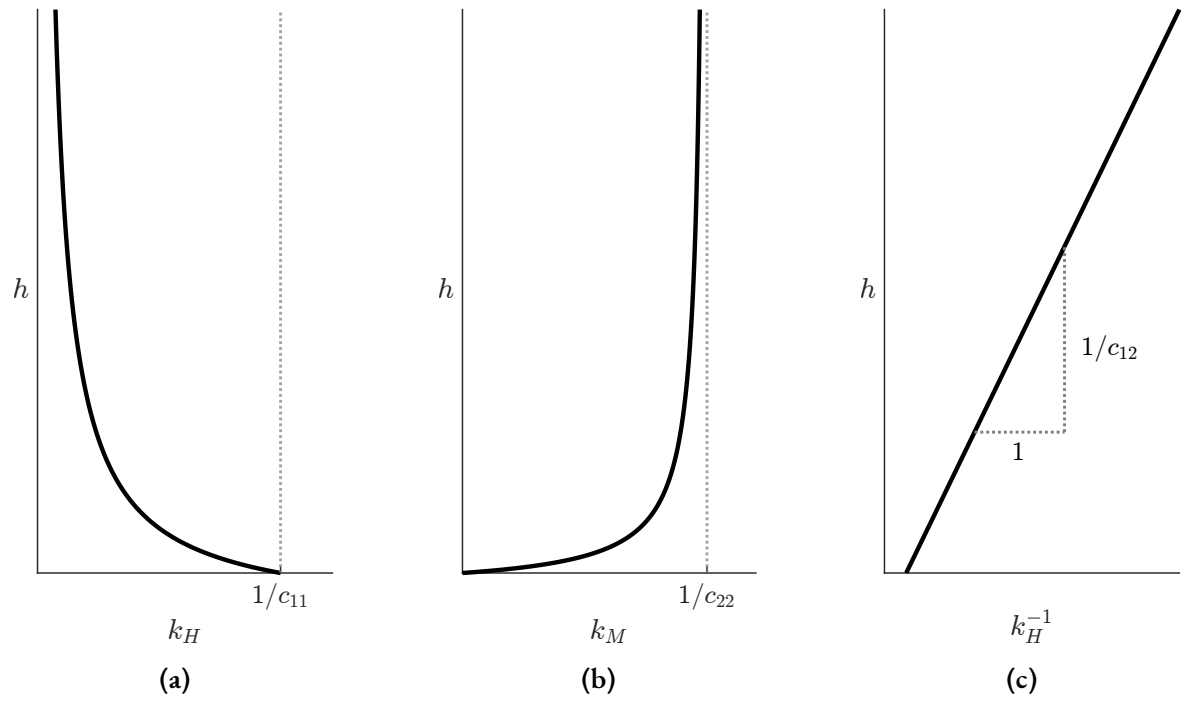


Figure 1.3: Example of relationship between (a) k_H and h (b) k_M and h (c) k_H^{-1} and h, for a fixed macro-element foundation. The salient features shown on each curve are determined by individual macro-element compliance parameters (c_{11}, c_{22} and c_{12}).

1.2.2 System identification and operational modal analysis

System identification refers to the estimation of a dynamic system's model parameters or structure based on input and output data [56]. For offshore wind turbines, the inputs are typically the wind or wave loading, which are unmeasured in practice, and the outputs are a set of measured vibration data (e.g., from tower accelerometers or strain gauges). System identification can be broken down into three stages [57]: data collection, subspace identification and estimation of modal properties.

One of the earliest methods used for estimating modal parameters is the so-called experimental modal analysis (EMA) [58], which typically involves exciting the structure with known forces and measuring the response. For example, Abdulraheem & Al-Kindi [59] applied EMA to detect faults in wind turbine blades by identifying the modal parameters after excitation with an impact hammer. In the case of large scale structures where laboratory testing is challenging, it becomes necessary to measure the response in environmental conditions. However, this poses a challenge for EMA based methods, since the inputs cannot be controlled or measured.

To overcome this limitation, output-only methods have been developed, where the modal

parameters are identified from the measured response of the wind turbine during operational conditions. This type of system identification is known as operational modal analysis (OMA) [57]. Various OMA methods exist, for which an extensive review is provided in the paper by Reynders [57]. However, the stochastic subspace identification (SSI) methods [56] have so far dominated output only modal identification. Two primary instances of this scheme are the covariance-driven approach (SSI-cov) and the data-driven approach (SSI-data). These methods estimate the system model based on projections of the measured response data and the use of linear algebra techniques. In addition to providing estimates of the modal parameters, subspace identification methods have also been developed to estimate the corresponding uncertainty associated with the modal parameters [60–62]. This uncertainty can be very valuable for interpreting the results. SSI-cov and SSI-data are based on the assumption that the system is linear and time invariant (LTI), and the input is broadband [56]. When offshore wind turbines are operating at high wind speeds, this assumption can easily be violated due to the non-proportional aeroelastic damping and harmonic excitation caused by the rotating blades [21]. Recent research has extended OMA methods to account for harmonic loading [63], as well as the development of time series models that are able to track the non-stationary modulation of modal properties [64].

OMA methods have been applied in an experimental setting in [65, 66], for the continuous monitoring of an onshore wind turbine in [67, 68], for an offshore wind turbine on a monopile foundation in [69, 70], as well as for floating turbines in [71]. Continuously tracking the modal parameters offers an approach for detecting damage due to a change in the systems parameters. For example, Chen *et al.* [72] used modal analysis to detect scour of an offshore wind turbine monopile due to a change in the mode shapes. Shirzadeh *et al.* [73] have also applied operational modal analysis techniques for the estimation of damping in offshore wind turbines. However, it is important to note that while output-only modal analysis methods can accurately identify frequencies and mode shapes, they tend to struggle when it comes to estimating damping ratios, as shown in [74, 75].

1.2.3 Model updating

For numerical models to be provide reliable predictions, it is crucial that their parameters are well calibrated. This can be done using model updating methods, in which uncertain model parameters are optimised to improve agreement between model predictions and measured data [76].

The measured data typically consist of vibration-based time histories from accelerometers, inclinometers, and strain gauges. For large-scale structures, where direct measurement of system inputs is often not possible, operational modal analysis is commonly applied as a first step to extract a set of identified modal parameters from the output-only data. The identified modal parameters can then be compared with the model-estimated modal parameters to calibrate the uncertain parameters. Therefore, these methods rely on the modal parameters being sufficiently sensitive to a change in the structural parameters of the model. Although model updating can also be performed directly using time series data (as is the focus of the approach developed in Chapter 6), this section only reviews the methods based on modal-data.

Model updating methods have been widely applied to civil engineering structures and can be loosely categorised as deterministic [21, 55, 77–79] or Bayesian [80–85]. In the deterministic approach, an objective function that quantifies the discrepancy between model predictions and measurements is minimised to determine the optimal set of uncertain model parameters. For example, Augustyn *et al.* [78] applied a deterministic approach to update the parameters of an offshore wind turbine supported with a jacket foundation, where the parameters to be updated were selected based on a sensitivity analysis. Using *in-situ* data, the authors in [78] were able to reduce the initial maximum eigenfrequency discrepancy from 30 % down to 1 %. In addition, they attributed the initial discrepancy to the high level of uncertainty in the initial estimate of the soil stiffness. This study not only showed the impact that the uncertainty in the foundation can have, but also showed how the modal parameters are sensitive to the foundation parameters. Deterministic model updating has also been applied to identify foundation properties for monopile-supported offshore wind turbines [19, 21, 55, 79, 86].

Model updating involves an inverse optimisation process. Its success depends on several

factors, including the type and placement of sensors, the number and nature of estimated modes, and the observability properties of the system to be estimated [87]. However, the optimisation process can be equally sensitive to uncertainties introduced through modelling or measurements as it is to the structural parameters of interest [85]. Even in cases where the model is sufficiently sensitive to the parameters being estimated, uncertainty remains inevitable and should always be accounted for during model updating [88].

The two primary sources of uncertainty in model updating are modelling uncertainty and measurement uncertainty. Measurement uncertainty commonly arises from sensor noise or signal processing errors. In modal-based updating, an additional source of measurement uncertainty is introduced during the operational modal analysis process. As discussed in Section 1.2.2, one advantage of subspace-based identification methods, such as SSI-cov or SSI-data, is their ability to quantify this uncertainty [60–62]. However, it is interesting to note that uncertainty estimates from subspace-identification methods have not yet been widely incorporated into model updating frameworks.

Common types of modelling uncertainty include parameter, structure and code uncertainty [77]. The model code uncertainty refers to the numerical errors introduced by the software or hardware required for the computational methods. This type of uncertainty can typically be discarded since its contribution is negligible relative to the other forms of uncertainty. Model parameter uncertainty relates to assuming structural parameters are time-invariant and ignoring the inherent variability of structural parameters that can result from changing operational or environmental conditions. For example, assuming a structural stiffness is constant despite variability with temperature or humidity. Finally, model structure uncertainty relates to the assumptions made about the underlying physics of the model. For example, using a linear model to represent a system that exhibits nonlinear behaviour.

While deterministic methods are advantageous due to the relatively low computational cost required to update the uncertain parameters, they struggle to account for model and measurement uncertainties that may be present. Moreover, they provide no insight into the associated uncertainty of the updated estimates. These limitations motivate the use of probabilistic frameworks for model updating.

In probabilistic model updating methods, uncertain parameters are treated as random variables described by probability distributions, as opposed to fixed values to be identified. This approach allows for the incorporation of prior knowledge and provides a natural way to quantify uncertainty in both the parameters and the model predictions. One of the most widely adopted probabilistic approaches is Bayesian model updating, first introduced to the field of structural dynamics by Beck and Katafygiotis [89, 90]. This framework is based on the well known Bayes' Theorem, and defines a probabilistic model of the uncertainties and uses measured data to formulate a likelihood function. The likelihood updates the prior probability density function (PDF) to obtain the posterior PDF. This framework allows for not only the estimation of the *Maximum a Posteriori* (MAP) parameter values but also for the quantification of their uncertainty. Such uncertainty quantification is particularly valuable for assessing the robustness of the model predictions [77], or for propagating uncertainty to obtain probabilistic estimates of fatigue life or reliability [91].

The Bayesian probabilistic framework was successfully applied by Simoen *et al.* [85] for the identification of damage and the quantification of uncertainty of a seven-storey reinforced concrete building. This work was motivated by the deterministic approach undertaken by Moaveni *et al.* [92], where it was shown that significant uncertainties were present during the progressive damage assessment. Simoen *et al.* [82] showed how a correlated prediction error can impact both the MAP estimates and the parameter uncertainty. The authors then go on to recommend and apply a strategy to select the optimal probabilistic model for the prediction error correlation. Prajapat & Ray-Chaudhuri [93] applied the Bayesian framework in a numerical investigation of a 12-storey shear building to incorporate the modal sensitivities into the prediction error. They showed that the information extraction from less sensitive data points can be restricted when using the more commonly applied equal variance approach. Argyris *et al.* [83] introduced a novel likelihood function which evaluates the mode shape discrepancy using the Modal Assurance Criterion (MAC) [94] and compared this to the more typically used vector-based approach. Furthermore, the use of Bayesian methods for identifying foundation parameters of offshore wind turbines has been demonstrated in numerical case studies [95], and in real-world applications to a 6 MW offshore wind turbine [86].

The previous works consider the classical Bayesian model updating approach. One limitation of classical Bayesian model updating is that, as the number of datasets increases, the posterior PDF approaches a Dirac Delta function, as shown by Behmanesh *et al.* [84]. This highlights how the classical Bayesian framework does not capture the inherent parameter uncertainty that large scale structures exhibit. To overcome this limitation, hierarchical approaches were introduced [84, 96–100]. In the hierarchical approach, a probability distribution is assumed for the structural parameters, and the hyperparameters of the chosen distribution are updated instead of the structural parameters. In doing so, as more datasets become available, the hyperparameters' posterior PDF approaches a Dirac Delta function whilst the parameter uncertainty is accounted for in the MAP estimates.

Behmanesh *et al.* [84] applied the hierarchical framework for the probabilistic identification of damage of a three-storey shear building in a numerical case study. Here it was shown that the hierarchical framework is better suited than the classical framework for damage assessment when the modal parameters are variable due to changing environmental conditions. The hierarchical framework can be taken one step further to update parameters at higher information levels. For example, Song *et al.* [97] successfully captured the stiffness-temperature dependence of a footbridge by updating the higher-level hyperparameters. Furthermore, Song *et al.* [101] applied a hierarchical framework to identify aerodynamic stiffness and damping of a 6 MW offshore wind turbine supported by a jacket foundation. Two information levels were considered: one using only the identified modal parameters, and another which also included measurements from the Supervisory Control And Data Acquisition (SCADA) system. Their study concluded that the higher information level offers improved capability to account for aerodynamic variability arising from changing environmental and operational conditions.

(i) Offshore wind turbine system complexities for model updating

The doctoral thesis by Anderson [21] focused on identifying the foundation parameters of a monopile-supported offshore wind turbine using a deterministic framework, where the challenges that are unique to this type of structure were addressed. In this section, two of those most relevant to this thesis are discussed.

Firstly, the quasi-axisymmetric nature of offshore wind turbine structures leads to closely spaced modes [19]. While this results in unique (but nearly repeated) eigenvalues and orthogonal eigenvectors, operational modal analysis methods often struggle to separate these modes [102]. As a result, the identified mode shapes tend to fall within the subspace spanned by the two underlying closely-spaced modes. Figure 1.4 illustrates this concept using the first pair of tower modes from an offshore wind turbine. In this example, the true modes align with the fore-aft and side-side directions; however, the identified mode shows mixed contributions from both, reflecting the difficulty in isolating them using operational modal analysis.

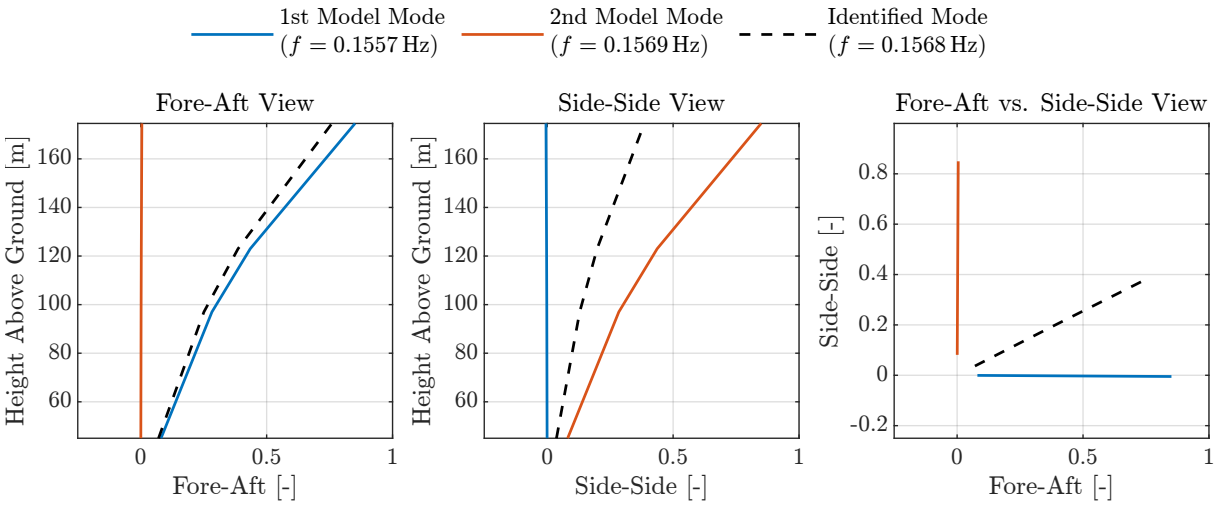


Figure 1.4: Example highlighting the challenge of identifying closely spaced modes: the identified mode shape lies within the subspace spanned by the two true model mode shapes.

Anderson [21] showed that this artefact can be problematic when comparing individual model-predicted and identified mode shapes using similarity metrics such as the Modal Assurance Criterion (MAC) [94]. This is because the poorly identified modes can lead to misleadingly low MAC values, which may result in incorrect mode matching during model updating. To address this, a higher-order similarity metric, known as the S2MAC, has been proposed [103]. The S2MAC can evaluate how well an identified mode shape aligns with the subspace spanned by the two model-predicted modes, making it a more robust metric for handling closely-spaced modes.

The second significant challenge identified by Anderson [21] is the non-proportional nature of damping in offshore wind turbines, which is primarily introduced due to the velocity-dependent aerodynamic damping matrix. Non-proportional damping increases the complexity

when model updating, due to the effect it has on the system's mode shapes. Due to the arbitrary scaling of mode shapes, it is important to use the same normalisation scheme to compare the model-estimated and identified mode shapes. For an undamped system, where the mode shapes are real valued, a widely-used technique is to normalise so that the maximum value of each mode shape is equal to one. When damping is present, which is usually the case, then the mode shapes depend on the type of damping. A common approach for modelling damping is to use proportional damping, where the damping matrix is proportional to the mass and stiffness matrices [104]. For this case, it can be shown that the proportionally damped and undamped mode shapes are equivalent. However, for the non-proportionally damped system, the damped mode shapes are no longer equivalent to their corresponding undamped mode shapes.

Therefore, either the damped or undamped mode shapes should be used in the model updating framework. However, it is generally preferred to use undamped mode shapes, since damped mode shapes can be difficult to estimate for two reasons: (1) challenges with the identification of damping using output-only operational modal analysis [105] and (2) modelling non-proportional damping effects [106]. Anderson [21] showed that when the normalisation scheme presented by Fillod [107], termed *complex normalisation*, is applied to a damped mode shape (whether proportionally or non-proportionally damped), the result is a good approximation of the corresponding undamped mode shape. In contrast, using the more common approach of discarding the imaginary part and normalising by its maximum real value can lead to significant errors in identified foundation parameters [21, 55].

Therefore, when carrying out model updating in this thesis, the undamped model and identified mode shapes are compared by applying the complex normalisation approach, shown in Equation 1.7:

$$\bar{\boldsymbol{\phi}}^{C} = \mathrm{Re}\left(\frac{\boldsymbol{\phi}_d}{\phi_{d,j}}\right), \tag{1.7}$$

where $\bar{\boldsymbol{\phi}}^{C}$ is the complex normalised mode shape, which is the undamped approximation of the damped mode shape $\boldsymbol{\phi}_d$ and j corresponds to the mode shape component that maximises $|\boldsymbol{\phi}_d|$.

1.2.4 Recursive methods for state, input and parameter estimation

Recursive estimation methods aim to infer unknown quantities in a dynamical system, such as states, inputs, or parameters, by combining a physics-based model with measurements from a limited number of sensors. The system model typically involves a set of dynamic equations that relate the state of the system to its inputs and parameters. The dynamic state of the system refers to time-varying quantities such as displacement or velocity. The inputs to the system represent external excitation acting on the structure (e.g., wind and wave loading). The system parameters define the properties of the model, such as a structure's stiffness, mass, or damping.

One important application of recursive estimation methods is virtual sensing [108]. Here, the idea is to make use of observed data and a system model to estimate the response at unmeasured locations [109, 110]. This is particularly useful for offshore structures, since installing sensors at submerged or foundation-embedded locations can be both costly and lead to poor data quality. Therefore, virtual sensing provides a means to estimate the response at these challenging locations based on a limited set of measurements.

Various approaches have been explored within the domain of virtual sensing and, similar to Section 1.2.3, can be broadly categorised as deterministic and probabilistic. Among the deterministic approaches, modal decomposition and expansion (MD&E) is one of the most widely-used techniques [111]. This method is based on the decomposition of the measured data into the contribution of each mode, which can be used to reconstruct the response using the system model at unmeasured locations. The MD&E method relies on the accuracy of the finite element model and typically requires a deterministic model updating approach to calibrate the model parameters. Therefore, MD&E faces the same limitations relating to ill-conditioning and uncertainty previously discussed. Examples of MD&E being applied to offshore structures include the work by Song *et al.* [112], where it was used to predict the strain response of an offshore platform and the works by Henkel [79] for estimating the subsoil and submerged response of offshore wind turbines. Additionally, Nabiyan *et al.* [113] successfully applied an MD&E approach to accurately estimate the moment time history at the mudline of a 2 MW offshore wind turbine.

Since MD&E is a deterministic approach, it does not account for uncertainties in the measurements or model. To overcome this limitation, focus has shifted towards probabilistic approaches. One such option is found in Bayesian schemes, where the model and measurement uncertainties are incorporated into the framework. If the aim is to achieve online estimation (real-time or near real-time), such schemes are often implemented in variants of the Kalman filter [114]. In this framework, the prior knowledge of a state is predicted in a time update step, before computing the innovation (discrepancy) between measurements and estimated quantities, to estimate the posterior of the entire state vector. The Kalman filter has been successfully applied by Papadimitriou *et al.* [115], who used acceleration measurements to estimate the strain response at unmeasured locations of a metallic structure.

Although the classical Kalman filter has been widely used for state estimation across various disciplines, for large scale engineering structures, the inputs which excite the system are often unknown. These unknown inputs can be accounted for within the measurement and modelling uncertainty terms, which are often (although not necessarily) assumed to be Gaussian noise sources—but they are not directly estimated. Since such knowledge can be valuable, several extensions to the classical Kalman filter have been proposed to address the joint input-state estimation problem. For example, Gillijns and de Moor [116, 117] developed a filter which estimates both the state and the input at the same time. This approach, often referred to as the GDF in the literature, requires no prior assumptions about the input characteristics. Instead, it determines the inputs by enforcing conditions of minimum variance and unbiasedness.

Another common filter for joint input-state estimation is the Augmented Kalman Filter (AKF) [118], which augments the state-space model from the classical Kalman filter to also include the unknown inputs to be estimated. While effective in many settings, the AKF can suffer from observability issues when only acceleration measuremetns are available. However, to overcome these limitations, Naets *et al.* [119] proposed the AKF with "dummy" displacements (AKF-dm) and showed how the use of artificial displacements stabilises the state estimation. An alternative filter that has been proposed to tackle the unobservability problem is the Dual Kalman Filter (DKF) [120], where similar to the GDF but unlike the AKF, the input and state estimation is performed in two successive steps. Additionally, Tchemodanova *et al.* [121]

proposed a hybrid approach that combines modal expansion with an augmented Kalman filter (ME-AKF).

Maes *et al.* [122, 123] compared a Kalman filter, GDF, and MD&E methods to estimate the dynamic strain of a real world 3 MW offshore wind turbine with a monopile foundation. They conclude that in general, the three methods perform competitively for this task. Tatsis & Lourens [124] compared the DKF and GDF for estimating the modal states of an offshore wind turbine jacket structure, where they showed that the GDF slightly outperformed the DKF due to the DKF's need for an optimal regularisation parameter. However, they also discussed the requirement of either strain or displacement measurements in order for the GDF estimates to be stable. The AKF has also been applied for estimating the response of an experimental wind turbine blade [125], for providing probabilistic estimates of accumulated fatigue damage in an onshore wind turbine [91], and for estimating tower loads on an onshore wind turbine from simulated data [126]. Mehrjoo *et al.* [127] introduced a Bayesian method to recursively estimate the wind load on an offshore wind turbine, validating it both numerically and experimentally. Furthermore, the OWI-lab has made significant contributions for fatigue and input estimation on real-world offshore wind turbines [111, 128, 129].

One common challenge across many of these filtering methods is the need to tune the process and measurement noise covariance matrices. This tuning process can significantly influence the filter's performance, particularly when estimating a large number of inputs. A few methodologies have been developed to address this issue. For example the L-curve [130] has been used in many works for virtual sensing [124, 131, 132], which aids the user by providing information about the regularisation properties. However, the L-curve becomes more complicated when multiple parameters require tuning. Moreover, in practice it has been shown to produce curves which do not exhibit a clear L shape, rendering the selection of optimal parameter values difficult [118, 133]. This shortcoming led Vettori *et al.* [125] to develop the adaptive-AKF, in which the process noise is adaptively tuned based on a set of optimal parameters for a given batch of data. This approach also has the advantage of tuning in an online manner, in contrast to the L-curve, which has to be carried out offline in a preliminary step. While these methods can help to alleviate the challenges associated with tuning, they still

have their own limitations. As a result, a trial-and-error approach to tuning can often achieve satisfactory results in less time.

While many of the previous works assume a random walk to model the evolution of the inputs, more sophisticated methods can also be used. Alvarez *et al.* [134] introduced the idea of a latent force model (LFM), which models the unknown inputs as a Gaussian process (GP) defined by a covariance function [135]. This concept of combining the statistical model with a mechanics-based model is known as the Gaussian process latent force model (GPLFM). In the GPLFM, the tuning of the input noise is replaced by the optimisation of the covariance function hyperparameters. Furthermore, the choice of covariance function allows for greater flexibility into the evolution of the force and incorporation of prior knowledge. However, careful consideration should be given to the choice of covariance function for the expected dynamics of the input [136]. A significant contribution was made by Hartikainen and Sarkka [137, 138] in constructing a state-space formulation for the GPLFM, which allowed the joint input-state problem to be solved using a Kalman filter approach with RTS smoother [139]. This has been applied by Zou *et al.* [140] for joint input-state estimation, in the first real world validation of subsoil strain from an offshore wind turbine monopile structure. In addition, their work showed that the estimated input force can be used to compensate for modelling errors that were introduced due to underestimating the damping. Bilbao *et al.* [141] also applied a GPLFM to predict the strain response and fatigue for a full scale onshore wind turbine. In this study, they present a simplified procedure to find the covariance function hyperparameters for improved computational efficiency. Additionally, Vettori *et al.* [142] compared different types of GPLFM covariance functions for input-state estimation in an experimental study involving a 3D-printed, scaled wind turbine blade.

Recently, the extension towards solving the joint parameter-input-state problem has started to receive more attention in the literature. Naets *et al.* [143] first considered this problem for structural dynamics using an augmented version of the extended Kalman filter (EKF), where the EKF was introduced for nonlinear systems in [144]. Nabiyan *et al.* [113] applied an iterative EKF-based method for the joint parameter-input-state estimation of an offshore wind turbine, where they identify the rotational stiffness at the mudline. However, since the EKF is based

on a first order linearisation using a Taylor series expansion, it is limited for the case of highly non-linear systems. This led to the development of the unscented Kalman filter (UKF) and particle filter (PF) based methods for highly non-linear problems [145]. The work presented in [132, 146, 147] propose an augmented UKF for the joint input-parameter estimation problem, where the joint estimation is performed in two stages: first the classical Kalman filter is used for the estimation of the input, which is required for the augmented UKF to estimate the states and parameters (a.k.a. the DKF-UKF algorithm). A similar windowed UKF algorithm was applied by Song *et al.* [112] for the joint estimation of foundation stiffness, inputs and states for an offshore platform. Impraimakis & Smyth [148] considered a sequential Kalman filter (SKF) which uses two Kalman filters. Ebrahimian *et al.* [149] introduced an alternative Bayesian framework for model parameter and input estimation, based on a sequential *maximum a posteriori* estimation approach, which is shown to be equivalent to a sequential iterative EKF method. Additionally, to allow greater flexibility over the inputs, Rogers *et al.* [150] extended the GPLFM framework for the joint parameter-input-state problem.

For successful system identification, it is essential that the dynamic states, inputs, and parameters are observable given the system model and available measurements. If the system is theoretically unobservable, any identification method is likely to produce incorrect results. Recent works have tackled the theoretical problem of estimating the observability of systems under the presence of unmeasured inputs [87, 151], with the aim of determining sensor setups that allow estimation not only of the system states and parameters, but also of the unmeasured inputs or the corresponding Lie symmetries, prior to carrying out a measurement campaign [152]. In [87], an algorithm is presented to test for the observability of non-linear systems with unmeasured inputs, where its application to a wind turbine structure is demonstrated using a numerical example.

(i) Dynamic substructuring

Dynamic substructuring is a powerful engineering tool for analysing complex dynamical systems by decomposing them into smaller, more manageable substructures. This decomposition allows for more efficient modelling and analysis of the system's dynamic behaviour [153]. While

various substructuring methods exist, the most common approach is component mode synthesis [154, 155], where a reduced-order model is created by retaining the most important vibration modes for each substructure. These reduced-order models can subsequently be assembled to accurately represent the dynamic response of the full structure.

While substructuring techniques were originally developed to solve large complex structural dynamics problems, they have also been shown to offer significant advantages in virtual sensing applications [156]. When parts of a structure are difficult to model accurately—such as the rotor nacelle assembly in wind turbines—substructuring allows for a framework where only the component of interest has to be modelled and combined with measurements from sensors within its subdomain.

Tatsis *et al.* [131] presented a general framework for input-state estimation at the component level, using a limited number of output measurements. Their study compared both fixed- and free-interface model reduction methods and demonstrated the performance of the augmented Kalman filter across three numerical case studies involving wind turbine components. Similarly, Vettori [157] developed a reduced-order model of a large-scale wind turbine blade, which they applied for input-state estimation using strain measurements distributed along its length.

In addition, Caglio [158] proposed a substructure-based Kalman filter for joint input–state estimation and damage identification of an offshore jacket platform. Using limited measurements from the topside of the platform, they showed that the nonlinear response of the entire structure could be accurately estimated by modelling only a reduced substructure and treating the interface forces as unknown inputs. This demonstrates the potential of substructuring-based estimation methods to avoid modelling the entire system while still achieving accurate results.

1.3 Research objectives and thesis structure

Building on the literature, this thesis aims to explore and develop structural health monitoring techniques to address the two key industry challenges discussed in Section 1.1: (1) the high uncertainty in the parameters that govern the soil-structure interaction [16], and (2) the lack of detailed design inputs required to model the rotor nacelle assembly [19].

To achieve this, the work packages and associated research questions shown in Figure 1.5 were defined. Work Packages 1–3 focus on numerical case studies, allowing a comparison with the true solutions in a controlled setting. These case studies provide valuable insight into the performance of the structural health monitoring methods considered in this work, which could potentially contribute to more reliable and improved monitoring strategies within the industry. To complement this, Work Package 4 applies the developed methods to real-world data, contributing evidence of their practical applicability.

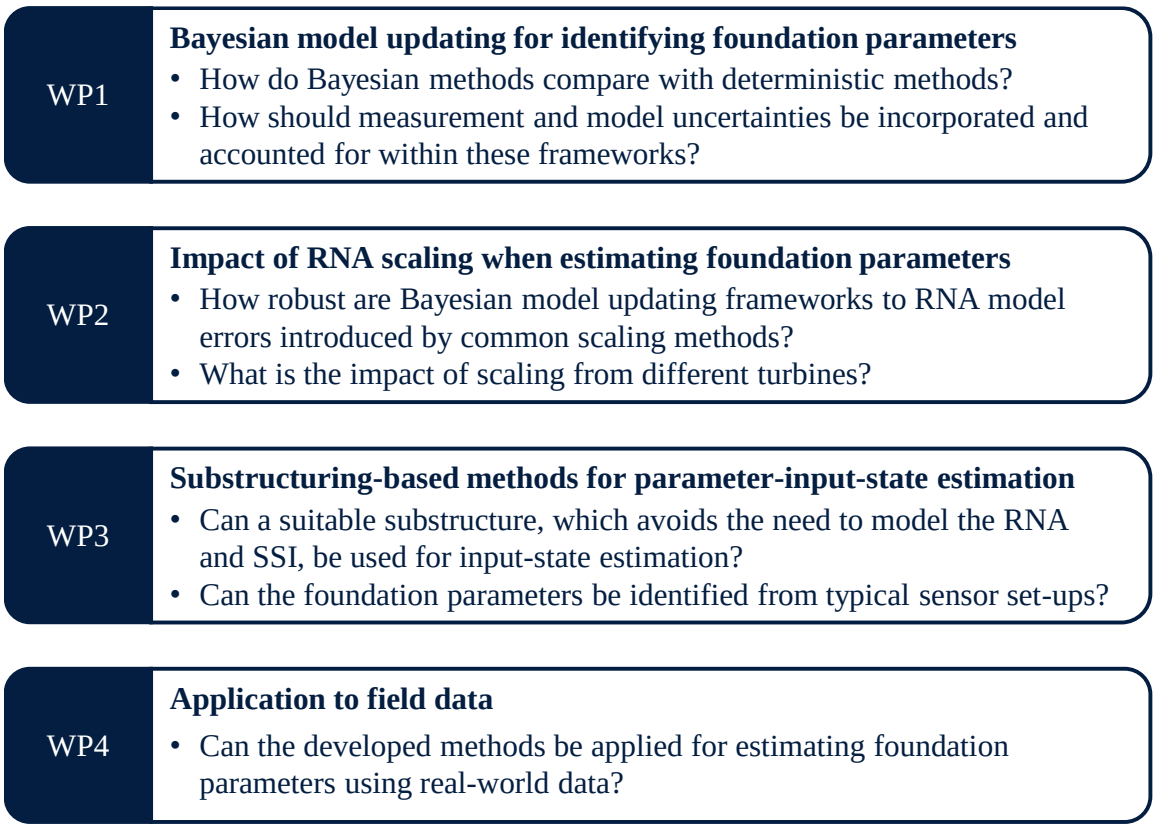


Figure 1.5: Summary of work packages (WPs) and research questions. Abbreviations: SSI = soil-structure interaction; RNA = rotor nacelle assembly.

1.3.1 Chapter Overview

This chapter (**Chapter 1**) has introduced the thesis, providing a background and motivation for the work. The relevant literature has been summarised and the main research objectives were defined.

Chapter 2 presents the aero-elastic offshore wind turbine simulator used for data generation and model updating in various chapters of this thesis. The model is adapted from the Simulink

floating offshore wind turbine model developed in [28, 29] and modified for a fixed-bottom, monopile-supported turbine. The key modelling features that influence the dynamics of this type of structure are discussed, along with details of how they were implemented.

Chapter 3 addresses **Work Package 1** by comparing deterministic and Bayesian model updating frameworks for identifying the parameters of a macro-element foundation. A numerical case study is presented using modal data obtained from operational modal analysis of a 15 MW turbine. Simulated data is generated under parked conditions with wind and wave loading, assuming no modelling error. The performance of each framework is evaluated, with a focus on how measurement uncertainties should be incorporated.

Building on Chapter 3, **Chapter 4** investigates the robustness of Bayesian model updating frameworks to modelling errors, addressing **Work Package 2**. Model error is introduced by assuming the blade properties are unknown and applying common scaling methods. The impact of different scaling methods from different sized reference turbines are analysed in terms of resulting blade properties, changes in modal parameters, and accuracy of the identified macro-element foundation parameters.

Chapters 5 and 6 shift away from model updating using modal data and focus on substructuring-based methods for parameter-input-state estimation, addressing **Work Package 3**. **Chapter 5** considers an unconstrained substructure that does not model the soil-structure interaction or the rotor nacelle assembly. A reduced-order model of the substructure is combined with vibration-based measurements for input-state estimation using the augmented Kalman filter. Numerical case studies examine performance under wind-only and combined wind–wave loading, with different sensor types assessed for their effectiveness in input, state, and fatigue estimation. **Chapter 6** extends this by introducing a substructuring approach for parameter estimation. A Bayesian inference framework is developed to identify the macro-element foundation parameters based on time-domain sensor data and corresponding quantities obtained from input-state estimation results.

Chapter 7 applies both the Bayesian model updating framework (from Chapters 3 and 4) and the substructuring-based inference framework (from Chapter 6) to identify the macro-

element foundation parameters in a real-world case study. The two independent identification methods are compared for cross-validation, with a discussion of their benefits and limitations. This chapter addresses **Work Package 4**.

Finally, **Chapter 8** concludes the thesis by summarising the main findings, discussing limitations, and offering recommendations for future research.

CHAPTER 2

Simulink Offshore Wind Turbine Model

2.1 Introduction

The structural health monitoring methods developed in this thesis rely on a physics-informed model to accurately represent the dynamic behaviour of an offshore wind turbine. This chapter discusses the key modelling requirements specific to this work and outlines the reasons for selecting the Simulink model adopted throughout the thesis. The main features of the model are then described, including details about its implementation using Matlab and Simulink.

2.2 Offshore Wind Turbine Modelling Requirements

The dynamic behaviour of an offshore wind turbine is governed by complex interactions between structural, aerodynamic, hydrodynamic, and controller-induced effects. To ensure robust analyses in the subsequent chapters, a model was required that could account for these interacting effects. In particular, due to the strong interaction between the dynamics of the tower and blades, the flexibility of the blades—i.e., the distribution of stiffness and mass along the blade length—can influence the structure's dynamic behaviour. Since this can in turn impact the identifiability of the foundation parameters, it was essential to use a model capable of capturing this non-linear, coupled behaviour.

For application to the methods and numerical case studies presented in this thesis, two types of simulation were required and are summarised in Figure 2.1. First, dynamic time-domain simulations were required to generate realistic synthetic response data under turbulent wind and irregular wave loading, for the numerical examples in Chapters 3–5. Second, modal analyses were necessary for identification of the foundation parameters using the model updating

frameworks applied in Chapters 3, 4 and 7. This required the ability to compute modal parameters from the non-linear model, which in turn necessitated linearisation about a steady-state operating point. Therefore, steady-state simulations to obtain these operating points were also essential for this work.

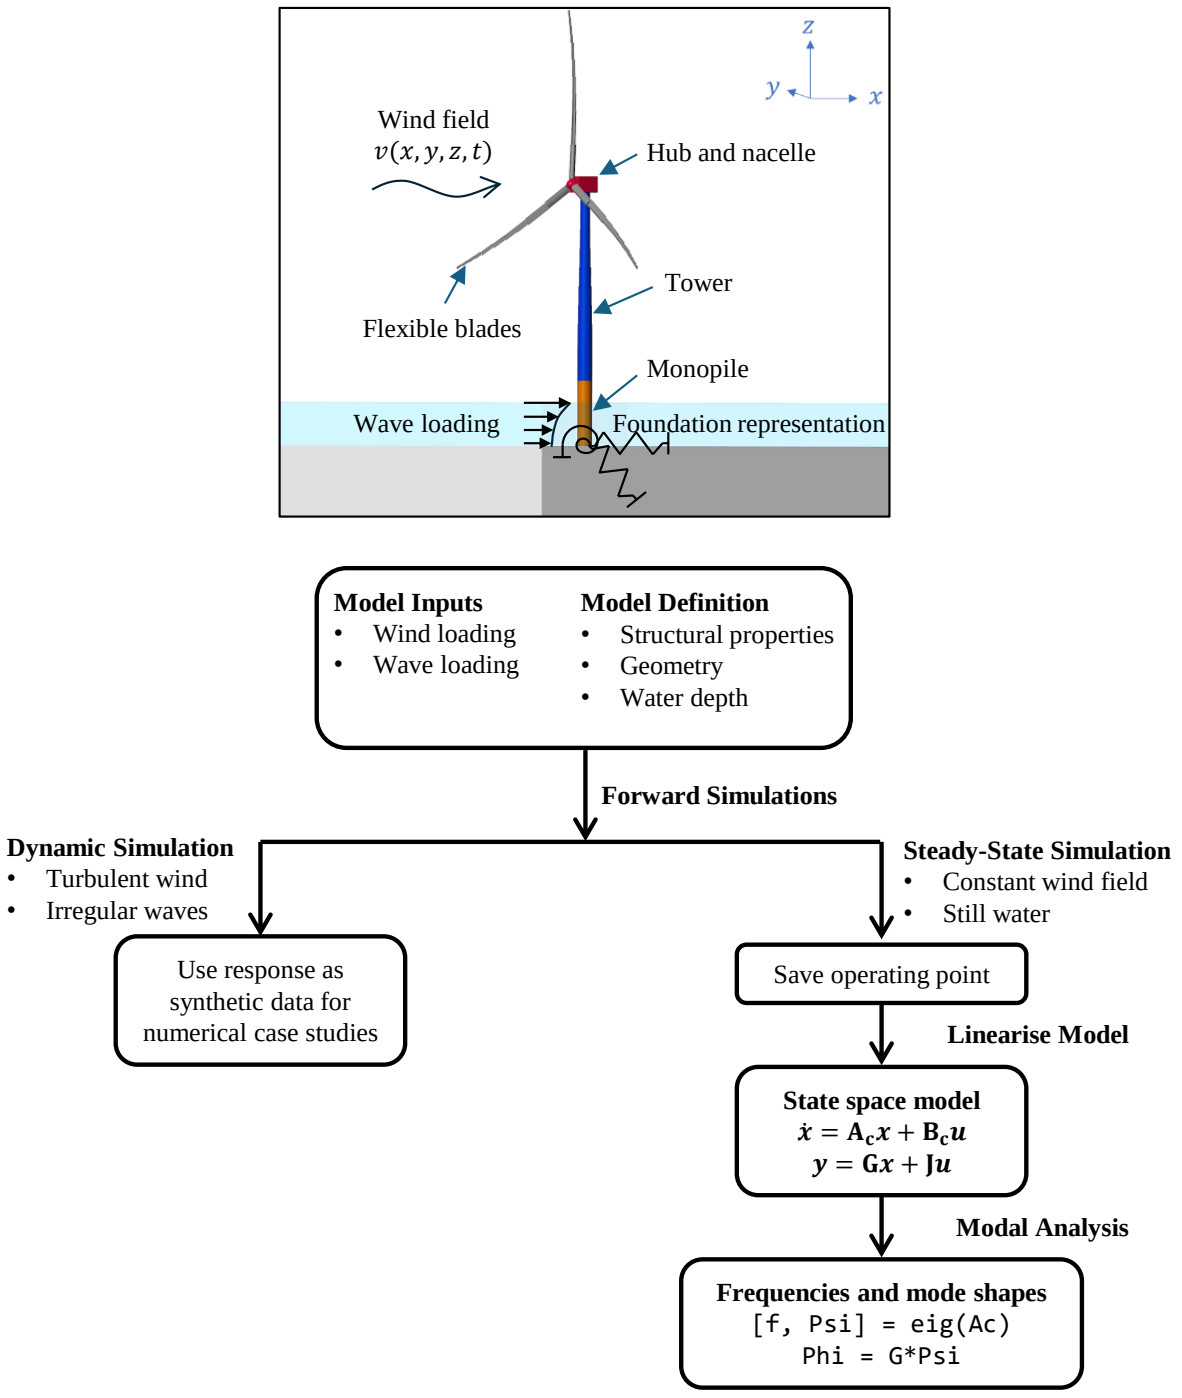


Figure 2.1: Offshore wind turbine model and simulation requirements for this thesis.

To meet these requirements, a fully coupled hydro-servo-aero-elastic model was implemented in Simulink. The model was initially developed by Dr Jiayao Meng for applications to floating

offshore wind turbines [28, 29]. For this work, the model was adapted to include a monopile foundation and a representation of the soil-structure interaction. While alternative and widely used simulation tools such as OpenFAST [24] or HAWC2 [25] were considered, the use of an in-house model offered greater flexibility for adapting to the specific needs and modifications required throughout this project. The following sections of this chapter describe the main features of this model.

2.3 Hydro-Servo-Aero-Elastic Simulink Model

The Simulink Model utilises the Simscape Multibody Simulation Environment [159] to represent the structural mechanics of an offshore wind turbine with a monopile foundation. This approach leverages Simscape's built-in blocks, allowing each component to be modelled individually and then assembled as needed. Table 2.1 lists the Simscape Multibody blocks used in this thesis. Moreover, implementing the structural model within Simscape allowed for seamless integration with Simulink's wider capabilities, so that the controller dynamics, aerodynamics and hydrodynamics could be easily coupled with the structural model.

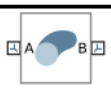
2.3.1 Soil-Structure Interaction Modelling

The soil-structure interaction is modelled using the macro-element foundation introduced in Chapter 1. As a reminder, this foundation is shown again in Figure 2.2 and consists of three linear elastic springs (k_{11}, k_{22}, and k_{12}) that define the soil reaction force and moment through Equation 2.1:

$$\begin{bmatrix} H_G \\ M_G \end{bmatrix} = \begin{bmatrix} k_{11} & k_{12} \\ k_{12} & k_{22} \end{bmatrix} \begin{bmatrix} v_G \\ \psi_G \end{bmatrix} \tag{2.1}$$

This foundation is based on the assumption that the monopile undergoes relatively small deformations, so the soil response behaves linearly. Additionally, it is assumed that the soil mass, inertia, and damping have a negligible influence on the response. Note that Equation 2.1 describes the soil-structure interaction in a single lateral direction for simplicity; however, identical springs are also applied in the perpendicular direction. Throughout this thesis, the axial and torsional DOFs are rigidly constrained. These simplifications are reasonable given the symmetric geometry of the monopile and the dominance of lateral wind and wave loading in

Table 2.1: Simscape Multibody blocks used to model the offshore wind turbine structural mechanics.

Simscape Multibody Block		Description/Usage
	Flexible Cylindrical Beam	Uniform Euler-Bernoulli beam used for tower and monopile elements
	Reduced-Order Flexible Solid	Used to model flexible blade elements. Mass and stiffness matrices are derived from a separate reduced-order model.
	Spherical Solid	Represent a component as a point mass with inertia. Used to model the hub and nacelle.
	Revolute Joint	Allows one rotational degree of freedom between two frames. Used to model the blade pitch, hub rotation, and ground-level rotation.
	Rectangular Joint	Allows two translation degrees of freedom between two frames. Used to model the ground-level lateral displacements.
	Rigid Transform	Defines a 3D transformation between two frames.
	Transform Sensor	Measures the response (position, velocity, acceleration, internal forces) between two frames.
	External Force and Torque	Applies an external force and torque to the attached frame. Used to implement hydro-dynamics, aerodynamics and soil forces.

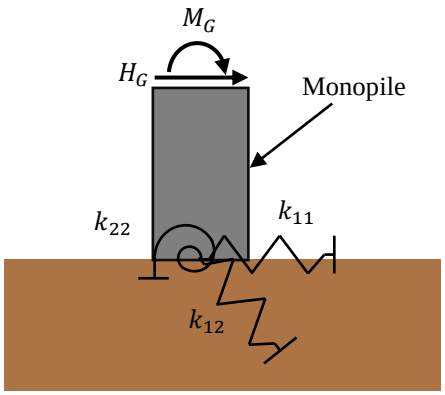


Figure 2.2: Macro-element foundation representation.

offshore environments.

To implement the macro-element foundation model in the Simscape environment, a series combination of one *Rectangular Joint* block and two *Revolute Joint* blocks are used to release the required lateral and rotational degrees of freedom. The displacement and rotation response at ground level are then fed into a gain block that contains the macro-element stiffness matrix. The resulting soil forces are applied to the frame corresponding to the monopile at ground level. The Simulink user interface for this soil-structure interaction model is shown in Figure 2.3.

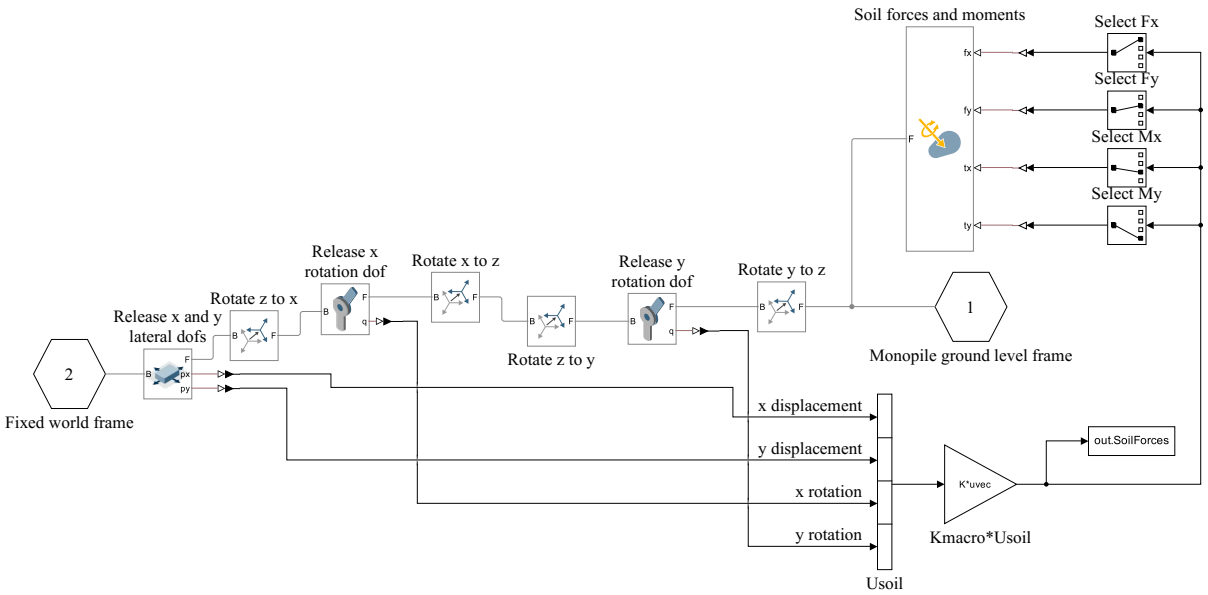


Figure 2.3: Simulink implementation of the macro-element foundation. The x and y degrees of freedom (DOFs) are aligned in the fore-aft and side-side directions, receptively. The z DOF is aligned in the vertical direction.

2.3.2 Tower, Transition Piece and Monopile Modelling

The monopile, transition piece and tower are modelled as a combined subcomponent, using beam elements. This was implemented in Simscape using either the built-in *Flexible Cylindrical Beam* block or the *Reduced-Order Flexible Solid* block (see Table 2.1). The Flexible Cylindrical Beam block models an Euler-Bernoulli beam, defined by geometric and material properties (typically steel for these components). As this block can only model uniform elements, tapered sections must be approximated by connecting multiple beam blocks. While this increases computational cost, it also improves modelling accuracy and allows the model to capture geo-

metric non-linearities due to large deflections. However, for monopile-supported offshore wind turbines, tower deflections generally remain within the linear regime.

If Timoshenko beams are required then the Flexible Cylindrical Beam block is not suitable. Instead, the Reduced-Order Flexible Solid block can be used. This block requires the user to provide the component's mass and stiffness matrices, derived from an external reduced-order model. In this thesis, a finite-element model defined by Timoshenko beam theory was used to assemble mass and stiffness matrices for the tower, transition piece and monopile. These matrices were subsequently reduced using a Craig-Bampton reduction [160] retaining a user-specified number of modes, to generate the required inputs for the Simscape block. The number of modes retained will be specified in Chapter 7 where the Reduced-Order Flexible Solid block is employed.

2.3.3 Rotor Nacelle Assembly Modelling

The rotor nacelle assembly (RNA) refers to the components mounted on top of the tower, consisting of the nacelle, hub and blades. The nacelle houses the key mechanical and electrical components, including the generator. All case studies in this thesis use wind turbines with direct-drive generators, meaning no gearbox is present; however the Simulink model has the capability to include one, as shown in [28, 29]. Therefore, the hub connects directly to the generator's rotor, and each blade is connected to the hub via a pitch bearing, which allows for blade rotation about its longitudinal axis. This sub-assembly is illustrated in Figure 2.4.

(i) Hub and Nacelle Simscape Implementation

The contribution of the hub and nacelle to the structural dynamics of the full offshore wind turbine system is modelled using lumped masses with associated moments of inertia about each axis. This was implemented in Simscape using two *Spherical Solid* blocks: one to represent the (possibly rotating) hub, and another to represent the nacelle—including all housed components—which moves rigidly with the tower top and does not rotate relative to it.

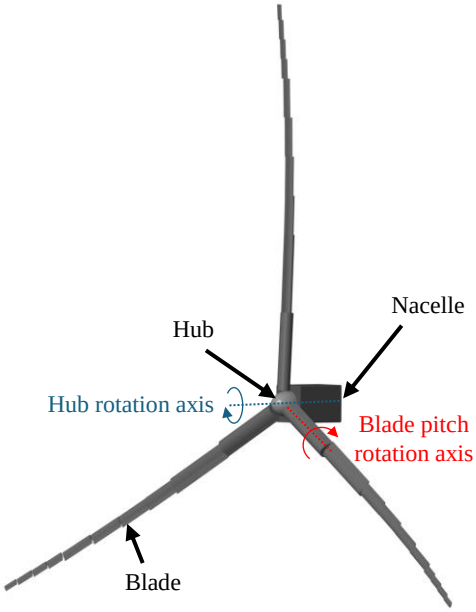


Figure 2.4: Rotor-Nacelle Assembly (RNA)

(ii) Flexible Blade Simscape Implementation

The blades are modelled as Euler-Bernoulli beam elements. Due to their complex geometry, they are implemented in Simscape using *Reduced-Order Flexible Solid* blocks. The user can specify how many of these blocks to include per blade, depending on the trade-off between model accuracy and computation cost. However, due to the high flexibility of the blades, their deflections can become large under operational loading conditions, increasing the influence of geometric non-linearities on the dynamic response. While each of these Reduced-Order Flexible Solid blocks follow the linear theory of elastodynamics, the connecting nodes between such elements in Simscape allow for large rotations and better capture geometric non-linearities. Therefore, to capture these effects accurately, it is recommended to use multiple connected blocks. Figure 2.5 highlights a single Reduced-Order Flexible Solid block within the blade sub-assembly.

For each block, a finite-element model comprising 1D Euler-Bernoulli beam elements was used to compute mass and stiffness matrices, based on material properties and geometry along the block's longitudinal axis. A Craig-Bampton reduction was then applied to the assembled matrices, retaining a user-specified number of modes per block, to produce the reduced-order mass and stiffness matrices required for the Simscape model.

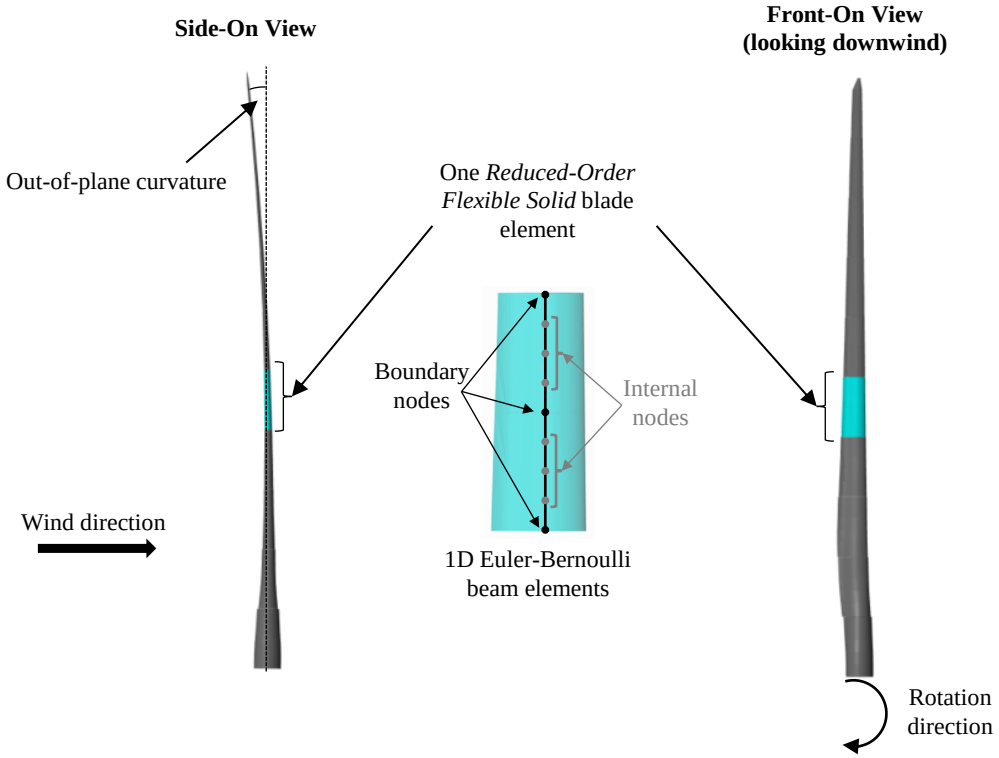


Figure 2.5: Example of blade model sub-assembly. Each blade is defined by a user-specified number of *Reduced-Order Flexible Solid* blocks. In this example, the blade is modelled using ten blocks, with one block highlighted for illustration.

For large-scale wind turbines, the blades are often curved out-of-plane or swept in-plane to reduce internal stresses under loading and to prevent tower strike during large deflections. The side-on view in Figure 2.5 shows an example of out-of-plane curvature. Although the finite-element mode can capture blade twist, each *Reduced-Order Flexible Solid* block in Simscape assumes a straight element. To account for geometric features such as curvature and sweep, *Rigid Transform* blocks are used between each reduced-order element to apply the necessary rotations. This provides an additional motivation for using multiple reduced-order elements: increased discretisation improves the model's ability to accurately represent complex blade geometry.

2.3.4 Aerodynamics

The aerodynamics module was implemented by Dr Jiayao Meng and has been validated against OpenFAST [28, 29]. This module computes the aerodynamic forces on each blade using unsteady blade element momentum theory [161] and applies these to the structural model within Simscape. Figure 2.6 summarises the key inputs for the aerodynamics module and how

it is coupled with the Simscape model.

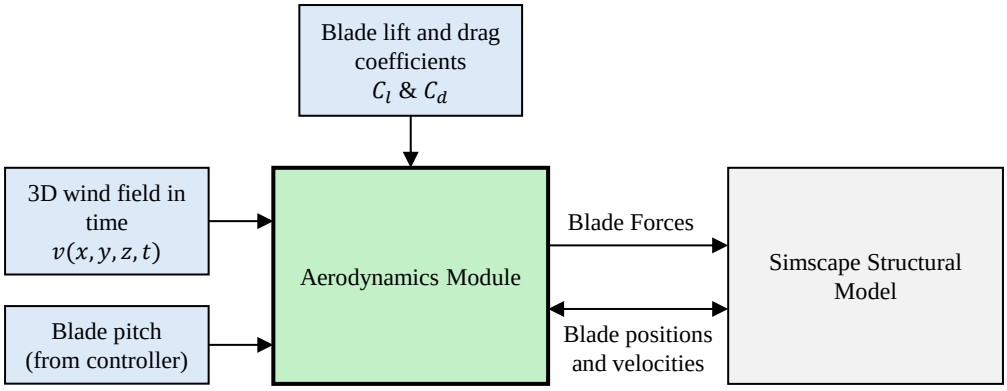


Figure 2.6: Schematic showing how the aerodynamics module interacts with the Simscape strucutral model.

The blade forces are calculated based on the relative wind speed at each blade position. To compute this, the module requires information on the blade positions and velocities, aerofoil lift and drag coefficients, blade pitch angles, and the incoming wind speed. The blade pitch is provided by the controller (see Section 2.3.6), and the blade kinematics are obtained from the Simscape structural model using *Transform Sensor* blocks. The user must define the aerofoil lift and drag coefficient data and provide a pre-defined wind field through which the blades pass during a simulation. In this thesis, both steady and turbulent wind fields are used. For the steady wind field, a constant wind speed is specified across the rotor area and remains constant through the simulation duration. For the turbulent case, a wind field is generated using Turbsim [162] prior to running the Simulink model, based on a user-defined mean wind speed at hub height and a turbulence intensity.

Once the aerodynamic forces have been computed, they are applied to the Simscape model using *External Force and Torque* blocks. The forces are applied at the aerodynamic centres of each *Reduced-Order Flexible Solid* block. Since the aerodynamic centres do not coincide with the boundary nodes of the flexible blocks, *Rigid Transform* blocks are used to offset the application points, ensuring the forces are applied at the correct locations.

2.3.5 Hydrodynamics

(i) Hydrodynamic added mass

For submerged elements, the inertia of the surrounding water that moves with the structure must be accounted for. This hydrodynamic added mass is incorporated by increasing the density of submerged monopile elements. The total effective mass per unit length, M_{eff} includes contributions from the monopile steel, the water inside the monopile, and the added mass of displaced external fluid, as shown in Equation 2.2:

$$M_{\text{eff}} = \underbrace{\rho_s A_s}_{\text{monopile steel}} + \underbrace{\rho_w A_i}_{\text{internal water}} + \underbrace{C_a \rho_w A_o}_{\text{added mass (external water)}}, \tag{2.2}$$

where ρ_s and ρ_w are the densities of the steel and water, respectively. A_s is the monopile's cross-sectional area, A_i is the internal water column's cross-sectional area, and A_o is the outer displaced area (i.e., $A_o = A_i + A_s$). The hydrodynamic added mass coefficient C_a accounts for the inertial effect of the surrounding fluid.

The effective density used for the submerged structural elements is then computed as:

$$\rho_{\text{eff}} = \frac{M_{\text{eff}}}{\rho_s} = A_s \left(1 + \frac{\rho_w}{\rho_s} \cdot \frac{A_i + C_a A_o}{A_s}\right). \tag{2.3}$$

(ii) Hydrodynamic loading

When sea states are such that the wave loading can influence the dynamic response of an offshore wind turbine, wave effects must be included in the model. In this thesis, irregular wave loading is modelled using a combination of linear wave theory [163] and the Morison Equation [164], as described below.

For an individual regular wave in finite water depth d, the horizontal water particle velocity, u_n, and accelerations, $\frac{du_n}{dt}$ with respect to the depth z below the mean water level, are derived

from linear wave theory as:

$$\begin{aligned} u_n &= \frac{H_n}{2}\frac{gk_n}{2\pi f_n}\frac{\cosh(k_n[d+z])}{\cosh(k_n d)}\sin(2\pi f_n t - \phi_n), \\ \frac{du_n}{dt} &= \frac{H_n}{2} gk_n \frac{\cosh(k_n[d+z])}{\cosh(k_n d)}\cos(2\pi f_n t - \phi_n), \end{aligned} \tag{2.4}$$

where H_n is the wave height, g is the acceleration due to gravity, ϕ_n is the wave phase, f_n is the wave frequency in Hz, and the wave number k_n is obtained from the linear wave dispersion relation (Equation 2.5):

$$(2\pi f_n)^2 = gk_n \tanh(k_n d). \tag{2.5}$$

The subscript n used above indicates that those variables are associated with the harmonic of an individual regular wave. In irregular seas, the surface elevation and associated kinematics result from the linear superposition of many such harmonics with different wave heights and phases. Therefore, the irregular sea water kinematics (u and $\frac{du}{dt}$) are defined as the linear combination of each nth harmonic:

$$\begin{aligned} u &= \sum_n u_n, \\ \frac{du}{dt} &= \sum_n \frac{du_n}{dt}. \end{aligned} \tag{2.6}$$

Consequently, to compute the irregular sea wave kinematics, the wave height H_n and phase ϕ_n of each harmonic must be specified for a given frequency f_n. These parameters can be derived by assuming a power spectral density $S(f)$ representative of the sea state. A common choice is to use the JONSWAP spectrum [165], which characterises realistic sea conditions and is defined by a significant wave height H_s and peak wave period T_p. An example spectrum is shown in Figure 2.7.

For a Gaussian random sea surface, Tucker *et al.* [166] showed that the individual wave heights H_n are not deterministic (except in the limit of infinitely many harmonics). Instead, the wave amplitudes $H_n/2$ should be drawn from a Rayleigh distribution with a root-mean-square (RMS) value of $\sqrt{2S(f_n)\Delta f_n}$, where Δf_n is the frequency separation of the harmonics.

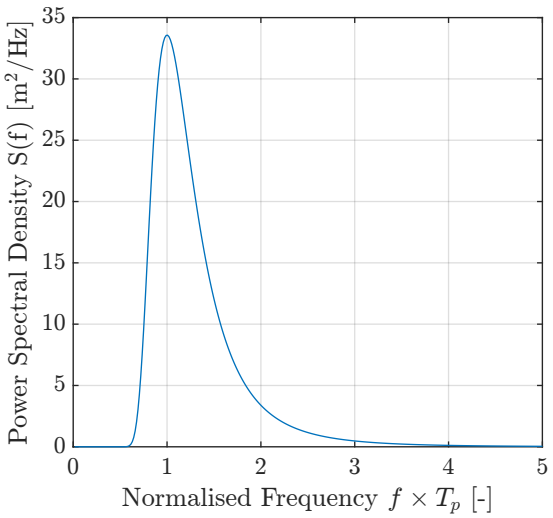


Figure 2.7: Example JONSWAP Spectrum

The corresponding phases ϕ_n can be randomly obtained from a uniform distribution bounded between 0 and 2π, and the wave number k_n is calculated from Equation 2.5. With f_n, k_n, H_n and ϕ_n defined, the wave kinematics follow from Equations 2.4 and 2.6.

One limitation of linear wave theory is that it is only valid up to the mean water level (i.e., $z = 0$). Therefore, to account for waves rising above and falling below this level, the Wheeler stretching method [167] is applied. This method substitutes the vertical coordinate z with z':

$$z' = d\frac{d+z}{d+\eta} - d, \tag{2.7}$$

where $\eta = \sum_n \frac{H_n}{2}\cos(2\pi f_n t + \phi_n)$ is the surface elevation. This transformation stretches or compresses the water column such that it remains within the range $-d$ to 0.

With the water particle velocities u and accelerations $\frac{du}{dt}$ known, the wave-induced force per unit length f_I acting on the cylindrical monopile is computed using the Morison Equation:

$$f_I = \rho_w R_o C_D u|u| + \rho_w \pi R_o^2 C_M \frac{du}{dt}, \tag{2.8}$$

where R_o is the outer radius of the monopile and ρ_w is the fluid density. The inertia coefficient, C_M (which typically ranges from 1.5-2), and the drag coefficient, C_D (which typically ranges from 0.5-1) are dependent on the Reynolds number, surface roughness, Keulegan-Carpenter number, and interaction parameters [168].

To apply the distributed wave forces described by Equation 2.8 to the Simulink model, it must be integrated over the length of each monopile segment to obtain the equivalent external nodal forces and moments. This integration is carried out in MATLAB using 4-point Gauss quadrature and executed prior to running a forward simulation. The computed wave loading time histories are then applied within the Simulink model using Simscape's *External Force and Torque* blocks.

2.3.6 Controller

The Simulink implementation of the NREL Reference OpenSource Controller (ROSCO) [169, 170] was used to represent the dynamics introduced by the wind turbine controller. At below-rated wind speeds, the ROSCO controller uses a proportional-integral (PI) controller to regulate generator torque, maximising power output by maintaining an optimal tip-speed ratio. At above-rated wind speeds, ROSCO uses a separate PI controller to maintain constant power output by regulating blade pitch. This also helps prevent excessive aerodynamic loading on the structure as wind speed increases.

2.3.7 Full-System Integration

Having described each of the model's main subsystems individually, they can be easily combined within Simulink. The overall model architecture is shown in Figure 2.8, illustrating how the full system is assembled from its constituent modules and components.

One key advantage of using Simulink and Simscape is the ease in which uniform global forces can be applied across each body. This feature was particularly useful for incorporating gravitational effects, and for investigating the influence of neglecting gravity in simulations.

Forward simulations can be executed using Simulink's built-in solvers. To linearise the model, Simulink's Model Linearizer App can be used, allowing for computation of the model's modal parameters, as described next.

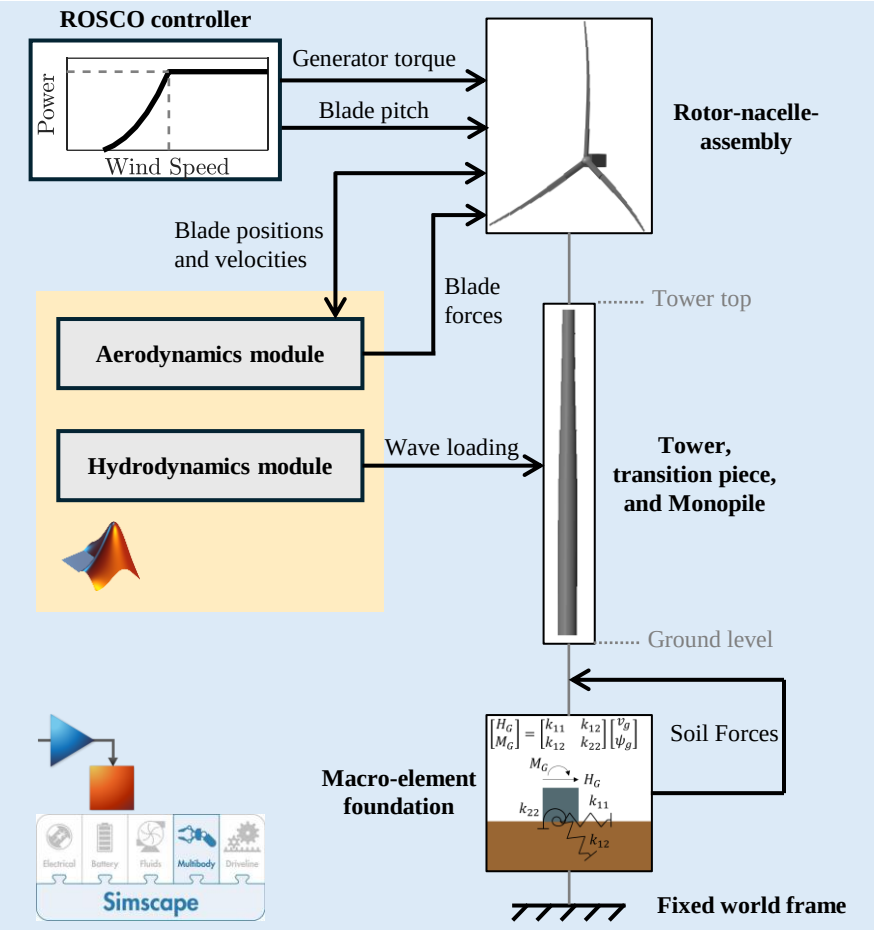


Figure 2.8: Full-system integration of the Simulink model for a monopile-supported offshore wind turbine with a macro-element foundation.

(i) Modal Analysis

As mentioned earlier in this chapter, computing the model's modal parameters (natural frequencies, mode shapes and damping ratios) is essential for the model updating methods introduced in later chapters. Modal parameters are properties of a linear system and therefore depend on the operating point about which the non-linear Simulink model is linearised.

In this thesis, the model was always linearised about a steady-state operating point, in which still water conditions were assumed and a constant steady wind field was assigned, where the wind speed was dependent on the specific scenario considered for each case.

The Model Linearization App was then used to generate a continuous-time state-space

representation of the system:

$$\begin{aligned} \dot{\boldsymbol{x}} &= \mathbf{A}|_{\text{op}}\, \boldsymbol{x} + \mathbf{B}|_{\text{op}}\, \boldsymbol{u}, \\ \boldsymbol{y} &= \mathbf{G}|_{\text{op}}\, \boldsymbol{x} + \mathbf{J}|_{\text{op}}\, \boldsymbol{u}, \end{aligned} \tag{2.9}$$

where $\boldsymbol{x}$ is the state vector, $\boldsymbol{u}$ is the input vector, $\boldsymbol{y}$ is the output vector, and $\dot{}$ indicates a time derivative. The matrices $\mathbf{A}|_{\text{op}}$, $\mathbf{B}|_{\text{op}}$, $\mathbf{G}|_{\text{op}}$, $\mathbf{J}|_{\text{op}}$, represent the state, input-state, output and feed-through matrices, respectively. The subscript 'op' indicates that these matrices are evaluated with respect to a specific operating point.

Wave-induced nodal forces and moments along the monopile, as well as wind speed, were defined as system inputs. By using wind speed rather than explicit blade forces as inputs, the effects of aerodynamic damping are captured in the linearised system matrices and therefore the resulting modal parameters. The choice of model outputs depended on the case study, but typically included acceleration responses at key locations along the tower, transition piece, and monopile to correspond with available measurements.

An eigenvalue analysis on the state matrix $\mathbf{A}|_{\text{op}}$ leads to:

$$\mathbf{A}|_{\text{op}}\, \boldsymbol{\Psi} = \boldsymbol{\Psi} \boldsymbol{\Lambda}, \tag{2.10}$$

where $\boldsymbol{\Lambda}$ is the diagonal matrix of eigenvalues, and $\boldsymbol{\Psi}$ contains the corresponding mode shape vectors in the state space.

Each diagonal entry λ_i of $\boldsymbol{\Lambda}$ represents an eigenvalue from which the corresponding ith natural frequency f_i (in Hz) and damping ratio ζ_i can be computed as:

$$\lambda_i = -2\pi f_i \zeta_i \pm 2\pi f_i \sqrt{1 - \zeta_i^2}, \tag{2.11}$$

$$f_i = \frac{|\lambda_i|}{2\pi}. \tag{2.12}$$

The mode shapes in the state space, $\boldsymbol{\Psi}$, can be easily converted to the output (or sensor) space for comparison with measured data by premultiplying by the linearised output matrix

$\mathbf{G}|_{\text{op}}$:

$$\mathbf{\Phi} = \mathbf{G}|_{\text{op}} \mathbf{\Psi}, \tag{2.13}$$

where $\mathbf{\Phi}$ is the mode shape matrix in the output basis.

2.4 Summary

This chapter has presented the implementation of a non-linear offshore wind turbine model with a monopile foundation, developed in Simulink. The model captures the key physical phenomena that influence structural dynamics, including aerodynamics, hydrodynamics and controller-induced effects. Structural components are modelled using the Simscape Multibody environment, allowing a modular and intuitive framework. Simulink's built-in solvers are used to perform time-domain simulations and system linearisation.

The developed model forms a core part of this thesis. It is used in subsequent chapters both as a forward model for generating synthetic response data and to identify the foundation parameters as part of the reverse problem. Table 2.2 provides an overview of how the model is used throughout the thesis. The specific user-defined inputs for each use case are detailed in the corresponding chapters.

Table 2.2: Application of the Simulink model throughout this thesis.

Chapter	Application of Simulink Model
1. Introduction	Not used
2. Simulink Offshore Wind Turbine Model	Not used
3. Model Updating Using Modal Data	Synthetic data generation Steady-state simulation Model linearisation Identification of foundation parameters Validation of results
4. Modelling Challenges for Foundation Parameter Estimation	Steady-state simulation Model linearisation Identification of foundation parameters Validation of results
5. A Substructuring Approach for Input-State Estimation	Synthetic data generation Validation of results
6. A Substructuring Approach for Parameter Estimation	Not used
7. Real-World Case Study	Steady-state simulation Model linearisation Identification of foundation parameters
8. Conclusions	Not used

CHAPTER 3

Model Updating Using Modal Data

3.1 Introduction

The literature review in Chapter 1 highlighted how previous studies have successfully applied both deterministic and Bayesian methods for parameter identification after first obtaining an estimate of the structure's modal properties using operational modal analysis. However, there still exists significant variability in the functional form of the objective/likelihood function. For example, Argyris *et al.* [83] introduced a novel likelihood function where the probabilistic model for the mode shapes is based on the modal assurance criterion (MAC) value rather than using the full mode shape vector. This MAC-based likelihood function was compared to a vector-based likelihood function and applied to update model parameters of the Metsovo Bridge.

A comparison of different methods for identifying foundation parameters of offshore wind turbines has not yet been conducted in a controlled setting, despite the salient influence of these parameters in fatigue assessment. To fill this gap, this chapter presents numerical case studies that compare one deterministic framework with two Bayesian frameworks for estimating parameters that define soil-structure interaction of a 15 MW offshore wind turbine supported by a monopile foundation. The deterministic framework follows a similar approach to [21, 55], while the Bayesian frameworks are based on the vector-based likelihood function from [77] and the MAC-based likelihood function from [83].

To identify modal parameters (natural frequencies and mode shapes) from simulated data, operational modal analysis was performed using the covariance-driven Stochastic Subspace Iden-

tification algorithm (SSI-Cov) [60]. As mentioned in Section 1.2.2, subspace identification methods have also been developed to provide an estimate of the uncertainty associated with the identified modal parameters [60–62]. This uncertainty can be incorporated into the model updating frameworks to account for measurement errors. However, the accuracy of the estimated uncertainty has not been investigated for offshore wind turbine structures, where the loading conditions can differ from the underlying white noise assumption used in these methods. Furthermore, the influence of including this uncertainty in the model updating frameworks on the estimated foundation parameters has not been investigated.

Therefore, this chapter aims to address these shortcomings by answering the following research questions:

1. Can the foundation parameters of an offshore wind turbine be accurately estimated using modal data?
2. What are the key differences between the various model updating frameworks in terms of most probable foundation estimates and the associated uncertainty?
3. How does the uncertainty in the modal parameters, estimated using SSI-cov for an individual dataset, compare with the uncertainty estimated from a frequentist approach across multiple datasets?
4. Does incorporating the modal parameters' measurement uncertainty into the model updating frameworks improve the identifiability of the foundation parameters?

The remainder of this chapter is structured as follows. A review of the operational modal analysis algorithm, SSI-cov, is provided in Section 3.2 for completeness, before describing the deterministic and Bayesian model updating frameworks used in this chapter in Section 3.3. The numerical case study is then presented in Section 3.4, which includes discussions on: the wind turbine description and model set-up in Section 3.4.1, the synthetic data generation in Section 3.4.2, the operational modal analysis results in Section 3.4.3 and the model updating results in Section 3.4.4.

3.2 Operational Modal Analysis

This section provides a brief review of the operational modal analysis method, SSI-cov, used as part of the numerical case study in Section 3.4.

The SSI-cov algorithm starts by constructing a covariance data matrix for a given model order n_{mo}, using the time history responses $\boldsymbol{y}(t) \in \mathbb{R}^{n_y}$ from the n_y measurements, sampled at time step dt. Linear algebra techniques are then used on this data matrix to estimate the state $\hat{\mathbf{A}}_d \in \mathbb{R}^{n_{\text{mo}} \times n_{\text{mo}}}$ and output $\hat{\mathbf{G}} \in \mathbb{R}^{n_y \times n_{\text{mo}}}$ matrices, for a discrete linear time-invariant system:

$$\boldsymbol{x}_{k+1} = \hat{\mathbf{A}}_d \boldsymbol{x}_k + \boldsymbol{w}_k \quad \& \quad \boldsymbol{y}_k = \hat{\mathbf{G}} \boldsymbol{x}_k + \boldsymbol{v}_k, \tag{3.1}$$

where $\boldsymbol{y}_k \in \mathbb{R}^{n_y}$ and $\boldsymbol{x}_k \in \mathbb{R}^{n_{\text{mo}}}$ are the measurement and state vectors at time step k, respectively. $\boldsymbol{w}_k \in \mathbb{R}^{n_{\text{mo}}}$ and $\boldsymbol{v}_k \in \mathbb{R}^{n_y}$ are the process and measurement noise, which are unknown but assumed to correspond to a zero-mean discrete white noise structure. Throughout this thesis, the notation $\hat{}$ is used to represent an identified or measured parameter and distinguish it from model-estimated parameters. Once $\hat{\mathbf{A}}_d$ has been identified, it is then converted to its continuous form through Equation 3.2:

$$\hat{\mathbf{A}}_c = \log\left(\hat{\mathbf{A}}_d\right)/dt, \tag{3.2}$$

where dt is the sampling time step and log denotes the matrix logarithm.

The modal parameters can then be extracted from an eigenvalue analysis on the state space matrix $\hat{\mathbf{A}}_c \in \mathbb{R}^{n_{\text{mo}} \times n_{\text{mo}}}$:

$$\hat{\mathbf{A}}_c \hat{\boldsymbol{\Psi}} = \hat{\boldsymbol{\Psi}} \hat{\boldsymbol{\Lambda}}, \tag{3.3}$$

where $\hat{\boldsymbol{\Lambda}} \in \mathbb{R}^{n_{\text{mo}} \times n_{\text{mo}}}$ is the diagonal eigenvalue matrix and $\hat{\boldsymbol{\Psi}} \in \mathbb{R}^{n_{\text{mo}} \times n_{\text{mo}}}$ is the mode shape matrix.

The diagonal terms in $\hat{\boldsymbol{\Lambda}}$ contain the eigenvalues $\hat{\lambda}_i$, from which the ith natural frequency

can be calculated:

$$\hat{\lambda}_i = -2\pi \hat{f}_i \hat{\zeta}_i \pm 2\pi \hat{f}_i \sqrt{1 - \hat{\zeta}_i^2},$$
$$\hat{f}_i = \frac{|\hat{\lambda}_i|}{2\pi}, \tag{3.4}$$

where $\hat{f}_i$ and $\hat{\zeta}_i$ are the ith natural frequency in Hz and damping ratio respectively.

$\hat{\mathbf{\Psi}}$ contains the mode shape vectors, which correspond to unknown states of the system. Therefore, to convert $\hat{\mathbf{\Psi}}$ to the known sensor basis so that it can be compared with model predictions, the output matrix $\hat{\mathbf{G}}$ can be used:

$$\hat{\mathbf{\Phi}} = \hat{\mathbf{G}}\hat{\mathbf{\Psi}}, \tag{3.5}$$

where $\hat{\mathbf{\Phi}} \in \mathbb{R}^{n_y \times n_{\text{mo}}}$ is the mode shape matrix corresponding to the measured vibration data in the sensor basis. Finally, $\hat{\mathbf{\Phi}}$ are complex normalised using Equation 1.7 to provide an approximation of the corresponding undamped mode shape.

As discussed in Section 1.2.3, a common challenge in identifying the mode shapes for offshore wind turbines is the close spacing of some modal frequencies [19]. This can make it difficult for subspace identification algorithms to distinguish between modes, resulting in identified mode shapes that are a linear combination of the true underlying modes. However, in this chapter, modes were chosen that were sufficiently spaced apart, so this issue was not observed (see Section 3.4.3).

Additionally, OMA methods tend to struggle with identifying damping ratios [74]. Therefore, only frequencies and mode shapes were considered for further analysis. The SSI-cov algorithm has been extended in [60] to estimate the uncertainty in the modal parameters, which includes the covariance of the ith frequency $\text{Cov}\left(\hat{f}_i\right)$ and mode shape $\text{Cov}\left(\hat{\boldsymbol{\phi}}_i\right)$. These estimated uncertainties provide valuable insight into the accuracy of the identified modes. In this chapter, they are used for selecting stable modes in Section 3.4.3. Additionally, since these uncertainties offer an estimate of the measurement error, they can be incorporated into model updating frameworks. However, the influence of these uncertainty estimates on model updat-

ing applications has not received much attention in the literature. Section 3.4.4 investigates the impact of incorporating this information into the presented model updating frameworks by comparing it to the case where it is discarded.

Throughout this thesis, the SSI-cov algorithm was implemented using the MACEC toolbox [171] for modal parameter identification and uncertainty quantification.

3.3 Model Updating Frameworks

3.3.1 Deterministic Model Updating

The deterministic approach to formulating a model updating framework consists in defining an objective function $J(\boldsymbol{\theta})$, which evaluates discrepancies between the model predictions and the measured data. Once OMA has been performed on the vibration data, the objective function can be defined based on the modal parameters.

A commonly used metric to compare two mode shapes, $\boldsymbol{\phi}_1$ and $\boldsymbol{\phi}_2$, is the modal assurance criterion (MAC) [94]:

$$\text{MAC}(\boldsymbol{\phi}_1, \boldsymbol{\phi}_2) = \frac{|\boldsymbol{\phi}_1 \boldsymbol{\phi}_2^*|^2}{\boldsymbol{\phi}_1 \boldsymbol{\phi}_1^* \boldsymbol{\phi}_2 \boldsymbol{\phi}_2^*}, \tag{3.6}$$

where $(\cdot)^*$ indicates the (complex conjugate) transpose of a vector. The MAC takes a value between zero and one, where a MAC of one indicates that two vectors are colinear and a MAC of zero indicates that two vectors are orthogonal.

In this chapter, the objective function is defined as:

$$J(\boldsymbol{\theta}) = \sum_{i=1}^{n_m} W_{f,i} \left(\frac{f_i(\boldsymbol{\theta}) - \hat{f}_i}{\hat{f}_i} \right)^2 + \sum_{i=1}^{n_m} W_{\text{MAC},i} \left(1 - \text{MAC}(\boldsymbol{\phi}_i(\boldsymbol{\theta}), \hat{\boldsymbol{\phi}}_i) \right), \tag{3.7}$$

where $f_i(\boldsymbol{\theta})$ and $\boldsymbol{\phi}_i(\boldsymbol{\theta})$ are the ith model-predicted frequency and mode shape respectively, for a chosen realisation of the parameters $\boldsymbol{\theta}$, and $\hat{f}_i$ and $\hat{\boldsymbol{\phi}}_i$ are the ith identified frequency and mode shape. n_m is the total number of modes identified from the data. $W_{f,i}$ and $W_{\text{MAC},i}$ are weighting factors that can be assigned to change the influence of the frequencies and/or mode shapes on the objective function. This work implements both an unweighted case, where $W_{f,i} = W_{\text{MAC},i} = 1$, and a weighted case where $W_{f,i}$ and $W_{\text{MAC},i}$ are defined based on estimates

of measurement uncertainty in the modal parameters. The specific formulation of these weights is described in later sections where they are applied.

The updated parameter estimates $\boldsymbol{\theta}^{\text{opt}}$ are those which minimise the objective function $J(\boldsymbol{\theta})$:

$$\boldsymbol{\theta}^{\text{opt}} = \arg\min_{\boldsymbol{\theta}} J(\boldsymbol{\theta}), \tag{3.8}$$

where $\boldsymbol{\theta}$ is a set of model parameters to be updated.

In this chapter, MATLAB's `fminsearch` [172] algorithm was used to solve Equations 3.7 and 3.8. This algorithm requires an initial guess for the updating parameters $\boldsymbol{\theta}$. It should be noted that this is a local optimisation method and therefore, the solution may not correspond to the global minimum. The significance of this is discussed in more detail in the numerical case study (Section 3.4).

3.3.2 Bayesian Model Updating

Bayesian model updating frameworks incorporate prior information and progressively update it as new data become available, resulting in a posterior distribution of the uncertain parameters. Equation 3.9 presents this mathematically, where $\boldsymbol{\theta}$ is a set of uncertain model parameters from the joint model class $\mathcal{M}$ and $\hat{\boldsymbol{d}}$ is the measured data.

$$p(\boldsymbol{\theta}|\hat{\boldsymbol{d}}, \mathcal{M}) = \frac{p(\hat{\boldsymbol{d}}|\boldsymbol{\theta}, \mathcal{M})p(\boldsymbol{\theta}|\mathcal{M})}{p(\hat{\boldsymbol{d}}|\mathcal{M})}. \tag{3.9}$$

In Equation 3.9, $p(\boldsymbol{\theta}|\hat{\boldsymbol{d}}, \mathcal{M})$ is the posterior distribution, $p(\hat{\boldsymbol{d}}|\boldsymbol{\theta}, \mathcal{M})$ is the likelihood, $p(\boldsymbol{\theta}|\mathcal{M})$ is the prior distribution and $p(\hat{\boldsymbol{d}}|\mathcal{M})$ is the evidence.

$p(\hat{\boldsymbol{d}}|\mathcal{M})$ provides the probability of observing the data given the model and can be used for model class selection. Assuming that each model class has equal prior probability, the most probable model class for a given set of data is the one with the highest $p(\hat{\boldsymbol{d}}|\mathcal{M})$ [82]. However, for a fixed model class and given data set, $p(\hat{\boldsymbol{d}}|\mathcal{M})$ is independent of $\boldsymbol{\theta}$ and can be considered as a constant which ensures that the posterior probability density function (PDF) integrates to one.

The prior PDF takes into account any previously acquired knowledge about the uncertain parameters $\boldsymbol{\theta}$, before measurements have been observed. The prior is often chosen based on expertise or engineering judgement. Common approaches for selecting the prior include the use of the principle of Maximum Entropy [85] or conjugate priors [97]. However, in this chapter, an uninformative uniform distribution is used to define $\boldsymbol{\theta}$. This simplifies Equation 3.9, reducing the posterior to being proportional to the likelihood function, $p(\boldsymbol{\theta}|\hat{\boldsymbol{d}}, \mathcal{M}) \propto p(\hat{\boldsymbol{d}}|\boldsymbol{\theta}, \mathcal{M})$, and making it dependent only on $\boldsymbol{\theta}$ for a fixed $\mathcal{M}$ and $\hat{\boldsymbol{d}}$. This choice was made to focus solely on the influence of the likelihood functions on the updated parameter estimates, providing a better comparison with the deterministic results.

The likelihood function is used to determine how well the measured data agree with the model predictions and can be considered as the Bayesian counterpart to the deterministic objective function. To formulate the likelihood function, a probabilistic model for the error between the model predictions and measurements must be assumed. Similarly to the deterministic objective function, the prediction error can be defined in terms of the modal data. However, the results can depend on how the probabilistic model for the prediction error is formulated. Therefore, in this chapter, two different formulations for the likelihood function are considered and compared. Both likelihood functions assume a normal distribution for the prediction error in the frequencies. The main difference between the two likelihood functions is in the treatment of the mode shape information. In the first approach, termed *vector-based* likelihood function, the error in each model-predicted and measured mode shape component is modelled as a normal distribution. Whereas, in the second approach, termed *MAC-based* likelihood function, a truncated normal distribution is used, defined based on the MAC between the model-predicted and measured mode shapes. Both likelihood functions are explained in more detail in the following sections.

(i) Vector-based likelihood function

To formulate the vector-based likelihood function, the prediction error is first evaluated. For a fixed model and dataset, the prediction error vector $\boldsymbol{\eta}$ is a function of the uncertain foundation parameters only and is constructed by stacking the frequency error on top of the mode shape

error: $\boldsymbol{\eta} = \left[\boldsymbol{\eta_f}; \boldsymbol{\eta_\phi}\right] \in \mathbb{R}^{n_m(n_y+1)}$, and,

$$\begin{aligned} \boldsymbol{\eta_f} &= \boldsymbol{f}(\boldsymbol{\theta}) - \hat{\boldsymbol{f}} \ \in \quad \mathbb{R}^{n_m} \\ \boldsymbol{\eta_\phi} &= \boldsymbol{\phi}(\boldsymbol{\theta}) - \hat{\boldsymbol{\phi}} \ \in \ \mathbb{R}^{n_m n_y}. \end{aligned} \tag{3.10}$$

n_m is the number of modes and n_y is the number of mode shape components.

The probabilistic model for the prediction error is commonly assumed to be a zero-mean multivariate Gaussian distribution, so that the likelihood function can be defined as:

$$p(\hat{\boldsymbol{d}}|\boldsymbol{\theta}, \mathcal{M}) \propto \det(\boldsymbol{\Sigma})^{-1/2} \exp(-\frac{1}{2}\boldsymbol{\eta}^T \boldsymbol{\Sigma}^{-1} \boldsymbol{\eta}). \tag{3.11}$$

The covariance matrix, $\boldsymbol{\Sigma}$, provides a useful way to take into account any measurement or modelling uncertainties that may be present. In this thesis, it is assumed that $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_G + \boldsymbol{\Sigma}_D$, where $\boldsymbol{\Sigma}_G$ and $\boldsymbol{\Sigma}_D$ are the model and measurement covariance matrices, respectively. Furthermore, all frequencies and mode shape components are assumed to be uncorrelated so that $\boldsymbol{\Sigma}$ is a diagonal matrix. Although this formulation accounts for both measurement and modelling uncertainties, it can be viewed as a limitation of the framework to group them together. This is because only the total error can be identified, making it difficult to differentiate between the two types of uncertainty [85].

If estimates of the uncertainty associated with the identified frequencies and mode shapes are available, they can be used to define the diagonal terms of the fixed measurement error matrix $\boldsymbol{\Sigma}_D$. Section 3.4 compares results when this uncertainty is included (i.e., $\boldsymbol{\Sigma}_D \neq 0$) to the case where it is ignored (i.e., $\boldsymbol{\Sigma}_D = 0$). The specific formulation for the case when $\boldsymbol{\Sigma}_D \neq 0$ is explained in Section 3.4.4 (ii).

Defining the model covariance matrix, $\boldsymbol{\Sigma}_G$, is more challenging due to the often limited information on modelling error prior to updating. Therefore, $\boldsymbol{\Sigma}_G$ is defined by a set of parameters $\boldsymbol{\theta}_G$, which are updated alongside the structural parameters $\boldsymbol{\theta}_S$. Thus, the combined set of model parameters to be updated is $\boldsymbol{\theta} = \{\boldsymbol{\theta}_S; \boldsymbol{\theta}_G\}$. In this thesis, $\boldsymbol{\theta}_G = \{\sigma_f; \sigma_\phi\}$, where σ_f and σ_ϕ are unknown dimensionless parameters representing the uncertainty in the frequencies

and mode shapes, respectively. σ_f and $\sigma_{\boldsymbol{\phi}}$ are defined based on the standard deviations for the ith frequency and mode shape, which are assumed to be proportional to the ith identified frequency $\hat{f}_i$ and the norm of the ith identified mode shape $\hat{\boldsymbol{\phi}}_i$:

$$\sigma_{f,i} = \sigma_f \hat{f}_i \quad \& \quad \sigma_{\boldsymbol{\phi},i} = \sigma_{\boldsymbol{\phi}} \|\hat{\boldsymbol{\phi}}_i\|. \tag{3.12}$$

The model frequency and mode shape variances, $\sigma^2_{f,i}$ and $\sigma^2_{\boldsymbol{\phi},i}$, are then entered into the corresponding diagonal entries in $\boldsymbol{\Sigma}_G$. In order to keep the number of updating parameters to a minimum, σ_f and σ_ϕ take the same value for all modes. Therefore, the number of parameters to be updated for the vector-based likelihood function is equal to $n_S + 2$, where n_S is the dimension of $\boldsymbol{\theta}_S$.

(ii) MAC-based likelihood function

One limitation of the vector-based likelihood function is that, due to the increasing number of mode shape components, it can result in unrealistically narrow posterior distributions as the number of sensors increases [83]. Therefore, to overcome this limitation, Argyris *et al.* [83] introduced a Bayesian framework that uses the MAC value rather than mode shape components. The MAC-based likelihood function implemented in this work is very similar to that introduced in [83]. All frequencies and MAC values are assumed to be independent and univariate distributions, such that $p(\hat{\boldsymbol{d}}, \mathcal{M}|\theta)$ assumes the following form:

$$p(\hat{\boldsymbol{d}}, \mathcal{M}|\theta) = \prod_{i=1}^{n_m} p(\hat{f}_i^2|\boldsymbol{\theta})\, p(\text{MAC}_i = 1|\boldsymbol{\theta}), \tag{3.13}$$

where $p(\hat{f}_i^2|\boldsymbol{\theta})$ is the probability density function for the ith frequency and $p(\text{MAC}_i = 1|\boldsymbol{\theta})$ is the probability density function for the ith MAC taking a value equal to 1.

Each modal frequency is assumed to follow a zero-mean Gaussian distribution so that:

$$p(\hat{f}_i^2|\boldsymbol{\theta}) = \mathcal{N}(\hat{f}_i^2; f^2(\boldsymbol{\theta}), \sigma_f^2 \hat{f}_i^4), \tag{3.14}$$

where $\mathcal{N}(X; \mu, \sigma^2)$ denotes a normal distribution for the random variable X, which has mean,

μ, and standard deviation σ.

To compare mode shape information, the MAC value is used. In this case, the MAC value is comparing two estimates of the same physical mode. In the absence of modelling and measurement errors, the MAC value should be equal to 1 at the true structural parameter values. Since the MAC value is bounded between 0 and 1, a normal distribution could result in non-zero probability for values outside these bounds. Therefore, Argyris *et al.* [83] suggested to use a truncated normal distribution to model the probability density function for each MAC value:

$$p(\text{MAC}_i = 1|\boldsymbol{\theta}) = \mathcal{TN}(1; \text{MAC}_i(\boldsymbol{\theta}), \sigma^2_{\text{MAC}_i}, 0, 1). \tag{3.15}$$

The notation $\mathcal{TN}(X; \mu, \sigma^2, a, b)$ defines a truncated normal distribution for a random variable X, which is centered on μ with variance σ^2. a and b define the lower and upper bounds of the distribution, respectively.

This probabilistic model assumes that there is no systemic error or bias in both the measurements and the model predictions, and is equivalent to the zero mean assumption in the vector-based likelihood function. However, in practice these errors cannot be avoided. Nonetheless, they can be accounted for by increasing the standard deviation of the distribution, which is controlled by the uncertainty parameters, σ_{f_i} and σ_{MAC_i}. The standard deviation for each frequency is defined as:

$$\sigma_{f_i} = \sigma^G_{f_i} + \sigma^D_{f_i}, \tag{3.16}$$

where $\sigma^G_{f_i}$ and the $\sigma^D_{f_i}$ reflect the uncertainties in the model and measurement, respectively.

The standard deviation for each MAC value is defined as:

$$\sigma_{\text{MAC}_i} = \sigma^G_{\text{MAC}_i}\left(1 + \frac{1}{\sqrt{n_y}}\right) + \sigma^D_{\text{MAC}_i}, \tag{3.17}$$

where $\sigma^G_{\text{MAC}_i}$ and $\sigma^D_{\text{MAC}_i}$ account for modelling and measurement errors. Following the advice in [83], Equation 3.17 also includes a term which is dependent on the number of mode shape components, n_y. The benefit of this term is that it penalises mode shapes with too few components. This is included because the accuracy of the mode shape is expected to increase

with the number of mode shape components, therefore more weight (lower standard deviation) should be assigned to these sensor setups.

The measurement uncertainty parameters, $\sigma_{f_i}^D$ and $\sigma_{\text{MAC}_i}^D$, can be defined based on the estimated uncertainty in the modal parameters. If no information is available on these parameters, then they can be set to zero. The specific formulations used in this work are defined in Section 3.4.

Similarly to the vector-based approach, the model uncertainty parameters ($\sigma_{f_i}^G$ and $\sigma_{\text{MAC}_i}^G$) are considered as unknown parameters to be updated alongside the structural parameters $\boldsymbol{\theta}_S$. To keep the number of updating parameters to a minimum, $\sigma_f^G = \sigma_{f_i}^G$ and $\sigma_{\text{MAC}}^G = \sigma_{\text{MAC}_i}^G$; i.e., each mode is assigned the same value. Therefore, the number of parameters to be updated for the MAC-based likelihood function is equal to $n_S + 2$, where n_S is the dimension of $\boldsymbol{\theta}_S$.

(iii) Evaluating the posterior

Evaluating $p(\boldsymbol{\theta}|\hat{\boldsymbol{d}}, \mathcal{M})$ typically involves solving complex high-dimensional integrals, particularly when the number of unknown parameters is large. Therefore, a common approach to compute the posterior is to use sampling methods, such as Markov Chain Monte Carlo (MCMC) methods, nested sampling, or Transitional MCMC (TMCMC) [173, 174]. Developing on some of the shortcomings of the TMCMC method, Betz *et al.* [175] developed the improved-TMCMC (iTMCMC) approach. In Section 3.4, the iTMCMC sampling method is used and was found to produce acceptable results. Compared to other methods, iTMCMC has several advantages, including the requirement for a relatively low number of user-defined parameters and the ability to estimate the model evidence $p(\hat{\boldsymbol{d}}|\mathcal{M})$ as a by-product.

3.4 Numerical Case Study

In this section, the model updating frameworks outlined in Section 3.3 are applied to a numerical case study focused on identifying the compliance parameters of an offshore wind turbine. The Simulink model described in Chapter 2 is used for forward simulations to generate realistic synthetic vibration data and is then further used to identify the compliance parameters that characterise the soil-structure interaction. It is important to note that the same model is

used for both data generation and model updating, thus this chapter does not consider any modelling error. This decision was made to allow for a more direct comparison between the different model updating frameworks and approaches for identifying the foundation parameters. Nevertheless, due to the artificially added sensor noise, along with identification errors arising from the OMA process, these numerical investigations do consider measurement error in the modal data. The results sections will compare the performance of the various model updating frameworks considered. In addition, the influence of incorporating known measurement uncertainty into each framework will be discussed.

3.4.1 Wind Turbine Description

For the numerical case study considered in this chapter, the Simulink offshore wind turbine model introduced in Chapter 2 was defined based on the IEA Wind 15-MW reference turbine [10] with a monopile foundation. An overview of the main dimensions of the structure are provided in Figure 3.1a.

Ten uniform cylindrical Euler-Bernoulli beam elements were used to model the steel monopile and tower, defined by an elastic modulus and Poisson's ratio equal to 210 GPa and 0.3 respectively. The monopile extends from ground level to 15 m above the mean water level with density 7850 kg/m^3. The tower then sits on top of the monopile and extends to 145 m above the mean water level with density 8500 kg/m^3. The diameter and wall thickness distributions can be found in the reference document [10].

Each blade was defined by a single reduced-order flexible beam element, where six external nodes were used to apply the aerodynamic forces and four normal modes were retained. This reduced blade model was found to significantly improve the computation time, which allowed multiple simulations to be run in a feasible time frame. The stiffness and mass matrices required for each blade were assembled using Euler-Bernoulli beam theory, based on the mass, stiffness and inertia distributions defined in [10]. The components that make up the hub and nacelle, as defined in [10], were modelled as individual lumped masses with associated inertias.

The soil-structure interaction was modelled using the macro-element foundation, introduced in Section 2.3.1 and illustrated in Figure 3.1b, which governs the ground level horizontal

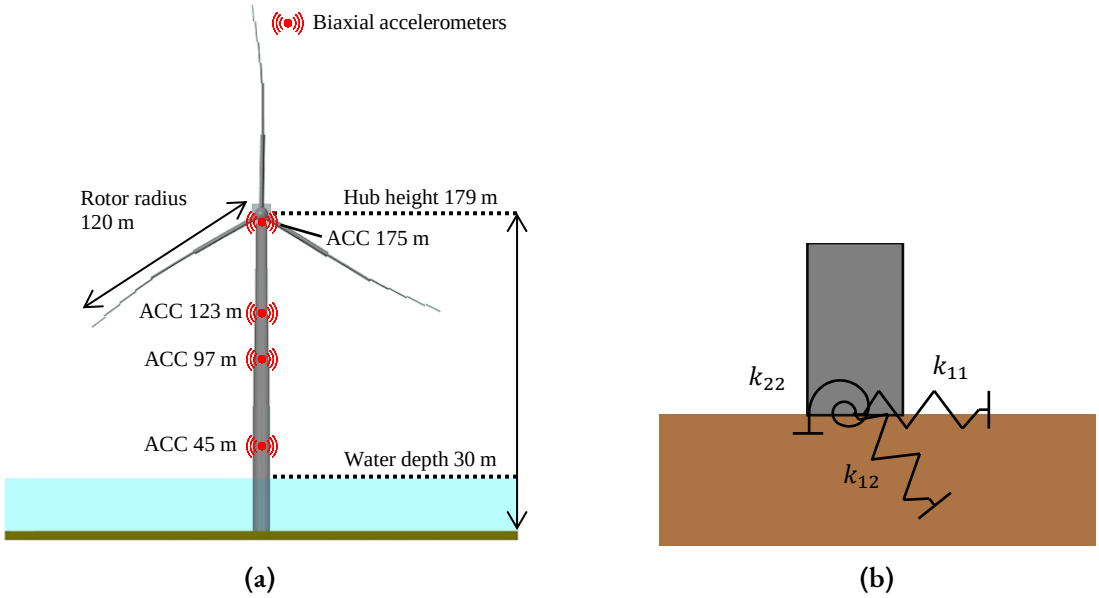


Figure 3.1: (a) Overview of the 15 MW offshore wind turbine and accelerometer locations (b) Macro-element foundation representation.

soil force and moment, H_G and M_G:

$$\begin{bmatrix} H_G \\ M_G \end{bmatrix} = \begin{bmatrix} k_{11} & k_{12} \\ k_{12} & k_{22} \end{bmatrix} \begin{bmatrix} v_G \\ \psi_G \end{bmatrix}, \tag{3.18}$$

where v_G and ψ_G are the ground level displacement and rotation.

To establish a baseline macro-element stiffness matrix for data generation and validation purposes in this chapter, the PISA design method [40] was initially applied to a finite element model where the monopile was modelled using Timoshenko beam elements and embedded in Cowden clay, for which the PISA approach has been calibrated . The embedded length of the monopile was 22.5 m with a diameter of 10 m, resulting in a L/D ratio of 2.25. The baseline stiffness matrix for the macro-element, $\mathbf{K}^0_{\text{macro}}$, was then calculated as the set of ground-level springs that are statically equivalent to the stiffness matrix obtained using the depth-wise PISA design method:

$$\mathbf{K}^0_{\text{macro}} = \begin{bmatrix} k^0_{11} & k^0_{12} \\ k^0_{12} & k^0_{22} \end{bmatrix} = \begin{bmatrix} 0.0532 & -0.3224 \\ -0.3224 & 5.1368 \end{bmatrix} \times 10^{11}, \tag{3.19}$$

where the units of k^0_{11}, k^0_{22} and k^0_{12} are N/m, Nm/rad and Nm/m, respectively.

3.4.2 Synthetic Data Generation and Sensor Set-up

To obtain realistic vibration data for modal analysis using SSI-cov, a ten-minute forward simulation was first carried out. Parked conditions were considered, where the rotor was locked and the blades pitched to 90 degrees. A turbulent wind field was generated using Turbsim [162], defined by a 2 m/s mean wind speed at hub height and 10 % turbulence intensity, which was then applied to the Simulink model. Since the wind turbine was operating below rated wind speed, the controller remained inactive and did not influence the dynamic response. In addition to wind loading, unidirectional wave loading was applied, calculated using the Morison Equation based on a JONSWAP wave spectrum with $H_s = 3.5\,\text{m}$ and $T_p = 8\,\text{s}$. Figure 3.2 shows example time histories for the hub height wind speed and wave elevation when using these hyperparameters. The wind and waves were assumed to be aligned and applied in the fore-aft direction (i.e. yaw angle equal to zero). These low loading conditions were chosen to better align with the linear spring assumption of the macro-element foundation. It should be noted that to obtain a reasonable comparison between the identified and model-estimated modal parameters, the model should be linearised about an operating point that is representative of the conditions for which the data are recorded. Therefore, the model-estimated modal parameters were calculated using Equations 2.9–2.13 about the steady state operating point corresponding to a constant 2 m/s wind field in still water.

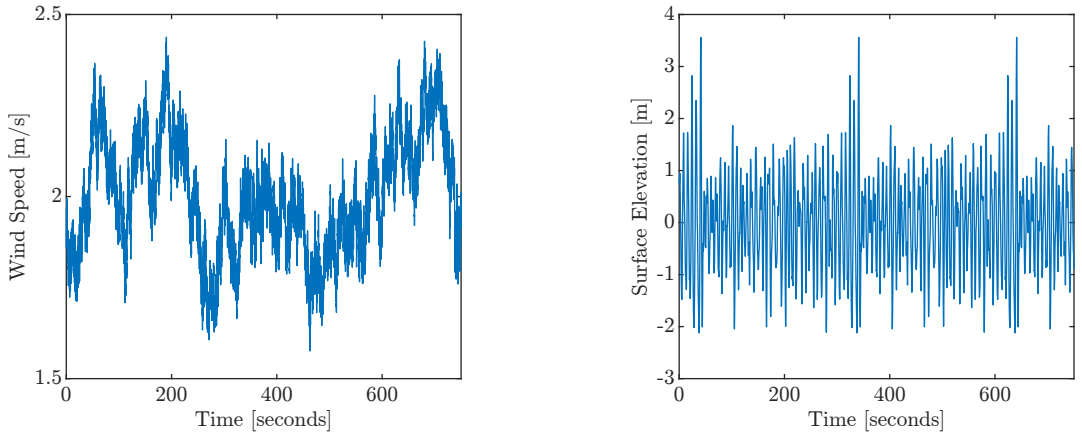


Figure 3.2: Example time histories of: (a) Kaimal hub-height wind speed (b) JONSWAP spectrum wave surface elevation.

Simulations were run for 750 seconds with the first 150 seconds discarded to remove any

transient effects related to the initial conditions, resulting in a ten-minute dataset. Four biaxial accelerometers in the fore-aft and side-side directions were used to measure the response along the wind turbine tower, as shown in Figure 3.1a. Then, to replicate the noise that exists in practical sensors, the acceleration time histories were corrupted by random Gaussian noise with a standard deviation equal to 3% of the RMS value of each signal.

This process was repeated 100 times using different random realisations of the turbulent wind and irregular wave loading, where the same mean wind speed, turbulence intensity, significant wave height and peak wave period were used for each realisation; only the random seeds used within Turbsim and for irregular wave loading were changed between different realisations. Figure 3.3 provides an example of the selected acceleration signals for one of the ten minute forward simulations. All forward simulations used the same initial conditions, corresponding to the steady state response from a constant 2 m/s wind field and still water conditions.

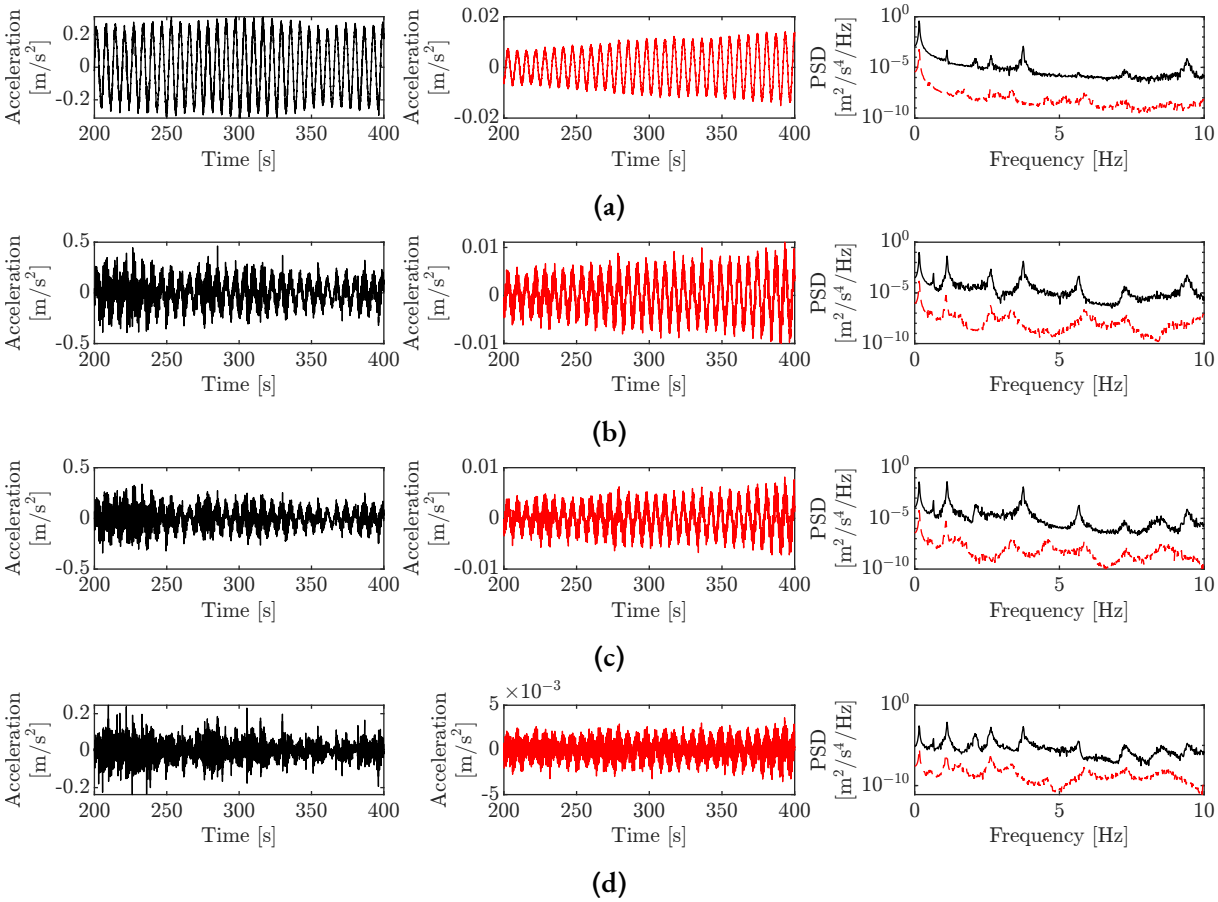


Figure 3.3: Example set of acceleration measurements. From left to right: fore-aft time domain response (200 second excerpt), side-side time domain response (200 second excerpt), and frequency domain response. Black lines show the FA response and red lines show the SS response. (a) Accelerometer at 175 m (b) Accelerometer at 123 m (c) Accelerometer at 97 m (d) Accelerometer at 45 m.

3.4.3 Operational Modal Analysis Results

After each forward simulation, operational modal analysis was carried out on the selected noisy acceleration signals—corresponding to the sensors shown in Figure 3.1a and the example dataset in Figure 3.3—to identify the modal parameters. For a given model order, SSI-cov provides estimates of the modes of the system, which can either correspond to physical modes or mathematical modes. To filter out mathematical modes, a stabilisation diagram can be used [57]. In this approach, the model order is gradually increased in steps of 2 and a mode for a given model order is considered stable if it shows less than 1% difference in frequency and a MAC value greater than 0.97, when compared with a mode from the previous model order.

(i) Mode set and model order selection

Figure 3.4 (left) shows a sample stabilisation diagram for one of the ten-minute datasets. The figure shows a large number of modes that are considered stable. However, not all of these modes should be used for model updating. A sufficient number should be used to ensure that the foundation parameters are identifiable, but using too many modes risks introducing significant measurement error, which may lead to bias in the foundation parameters' estimates. As the response was dominated in the fore-aft direction (see Figure 3.3), the fore-aft modes were expected to be more accurately estimated. As a result, only fore-aft modes were selected for model updating. In addition, only the six modes associated with the largest peaks in the PSD were used, as these were expected to correspond to tower-dominated modes. These six modes are illustrated in Figure 3.5. Figure 3.4 (right) shows an example of the cleared stabilisation diagram.

Finally, a model order should be chosen for each mode. This can be a crucial step since each model order has a different level of measurement error, which can subsequently impact the identified foundation parameters. In this work, the estimated uncertainty from SSI-cov was used to select a suitable mode; for each mode, the model order with the lowest frequency standard deviation was chosen. Figure 3.4 (right) highlights the selected model order for each mode for the example dataset.

It should be noted that the set of six reference modes was manually selected only for the

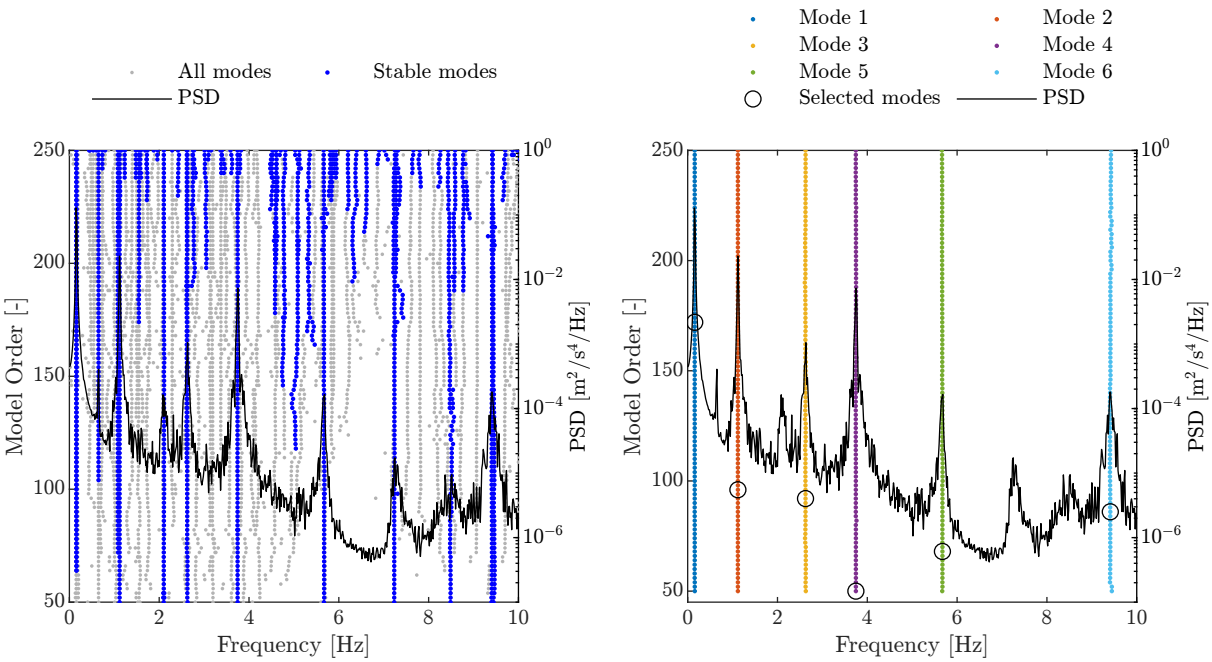


Figure 3.4: Example stabilisation diagrams from OMA on a representative dataset. Left: all stable modes based on criteria of 1 % frequency difference and a MAC value greater than 0.97 for consecutive modes. Right: filtered stabilisation diagram showing only the modes and selected model orders used for model updating.

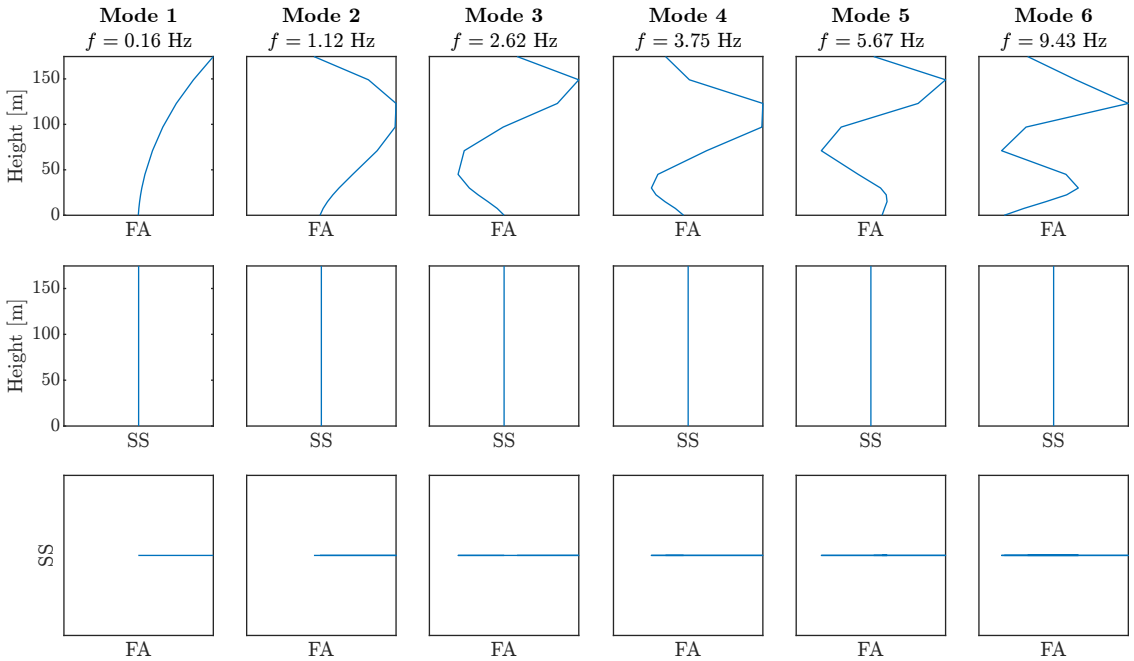


Figure 3.5: True model mode shapes for the six modes selected for model updating. Each row represents a separate mode. Top row shows Fore-Aft (FA) view, middle row shows Side-Side (SS) view, bottom row shows FA-SS view.

case of a single dataset. For the remaining 99 datasets, the process was automated using the same set of reference modes. Figure 3.6 compares the identified frequencies and mode shapes to the true model frequencies and mode shapes for each dataset. The figure shows that the frequencies were identified with errors of less than 3 %, and the mode shapes had MAC values greater than 0.988 across all datasets, indicating accurate identification of the modal parameters.

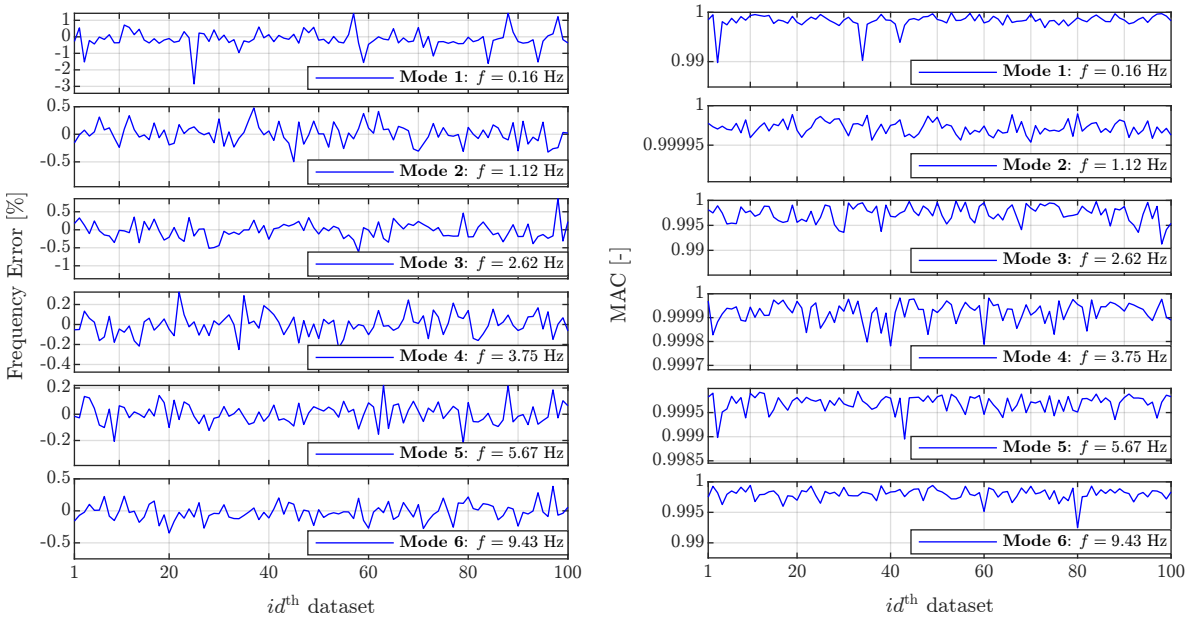


Figure 3.6: Comparison between the identified modes and true model-predicted modes across all 100 datasets. (a) Frequency errors (b) MAC Values.

(ii) Modal parameter bias and uncertainty: individual datasets vs frequentist approach

To compare the error in the identified modal parameters from individual datasets to that from a frequentist approach, the following error vectors were defined. The error vector between the modal parameters from the idth dataset and the true modal parameters was defined as:

$$\boldsymbol{\eta}_0^{id} = \begin{bmatrix} \hat{\boldsymbol{f}}^{id} - \boldsymbol{f}(\mathbf{K}_{\text{macro}}^0) \\ \hat{\boldsymbol{\phi}}^{id} - \boldsymbol{\phi}(\mathbf{K}_{\text{macro}}^0) \end{bmatrix}, \tag{3.20}$$

whereas the error vector corresponding to a frequentist approach was defined as the difference between the average of each modal parameter across all 100 datasets and the true modal

parameters:

$$\bar{\boldsymbol{\eta}}_0 = \begin{bmatrix} \left(\frac{1}{100}\sum_{id=1}^{100} \hat{\boldsymbol{f}}^{id}\right) - \boldsymbol{f}(\mathbf{K}^0_{\text{macro}}) \\ \left(\frac{1}{100}\sum_{id=1}^{100} \hat{\boldsymbol{\phi}}^{id}\right) - \boldsymbol{\phi}(\mathbf{K}^0_{\text{macro}}) \end{bmatrix}. \tag{3.21}$$

Note that the bar notation $\bar{}$ is used throughout this chapter to denote an ensemble average corresponding to the frequentist approach.

Figure 3.7a plots the sum of the error vector for each dataset $\sum |\boldsymbol{\eta}_0^{id}|$ and compares it with the sum corresponding to the frequentist approach $\sum |\bar{\boldsymbol{\eta}}_0|$. Here it can be seen that the overall error from a frequentist perspective is smaller than the error from any given dataset. This implies that as the total number of datasets increases, the average value for a modal parameter shifts closer to its true value. Therefore, the bias is closer to zero, which better aligns with the zero mean assumption used in the model updating frameworks. This is important to keep in mind when analysing the updated estimates for the foundation parameters in Section 3.4.4.

The estimated uncertainty from SSI-cov for a given modal parameter and dataset can also be validated against the uncertainty computed via a frequentist approach. For the im^{th} modal parameter, the standard deviation estimated from SSI-cov for the id^{th} dataset is defined as $\hat{\sigma}_{im}^{id}$, whereas the standard deviation of the im^{th} modal parameter from all 100 datasets using a frequentist approach is defined as $\bar{\sigma}_{im}$ and calculated using Equation 3.22:

$$\bar{\sigma}_{im}^2 = \frac{1}{99}\sum_{id=1}^{100} \left| \hat{\xi}_{im}^{id} - \frac{1}{100}\sum_{id=1}^{100} \hat{\xi}_{im}^{id} \right|^2, \tag{3.22}$$

where $\hat{\xi}_{im}^{id}$ is the im^{th} modal parameter from the id^{th} dataset: $\hat{\boldsymbol{\xi}}^{id} = \left[\hat{\boldsymbol{f}}^{id};\ \hat{\boldsymbol{\phi}}^{id}\right] \in \mathbb{R}^{n_m(n_y+1)}$.

Figures 3.7b and 3.7c provide a comparison between $\hat{\sigma}_{im}^{id}$ and $\bar{\sigma}_{im}$. The first observation from these figures is the presence of large error bars, indicating significant variation in the estimated uncertainty provided by SSI-cov across different datasets. One possible explanation for this is that the SSI-cov method assumes a white noise input, which is not true for offshore wind turbines due to the turbulent wind and irregular wave loading that follow Kaimal and JONSWAP spectra, respectively. Although this did not result in significant modal parameter bias (as evidenced in Figure 3.6), it could potentially explain the variability in the estimated

uncertainty observed in Figure 3.7b. Nonetheless, the figure also shows that the average standard deviations, estimated by SSI-cov across all 100 datasets (indicated by the marker on each error bar), approximately follow the 1:1 line, indicating a reasonable agreement with the standard deviations from a frequentist perspective.

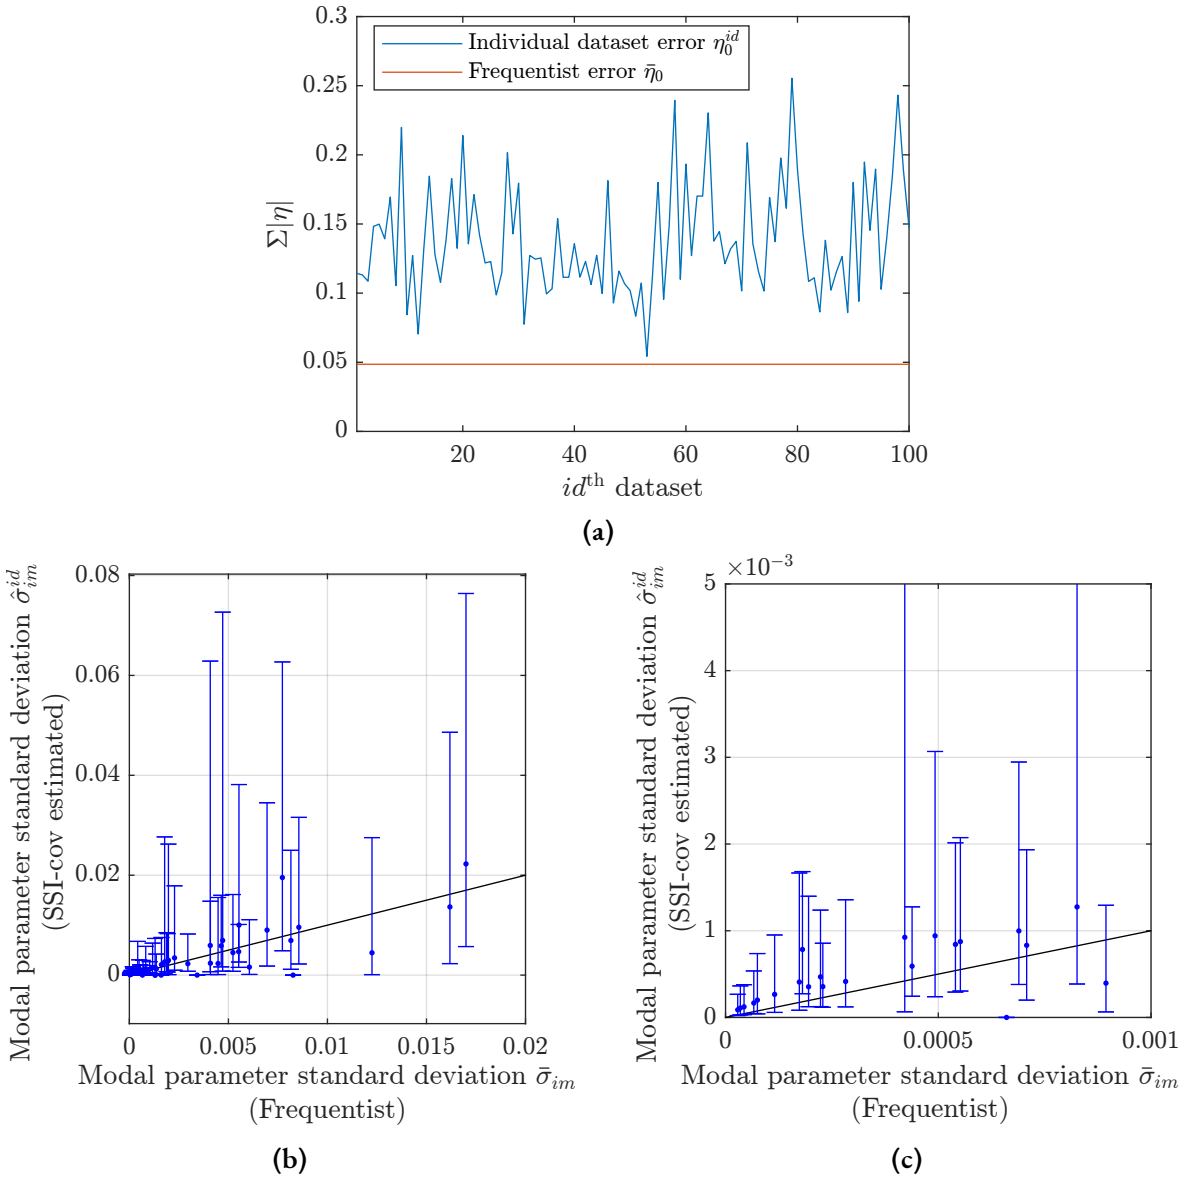


Figure 3.7: (a) Modal parameter error vectors (b) Comparison of modal parameter standard deviations from SSI-cov and frequentist estimates. Each point corresponds to one modal parameters (i.e. one modal frequency or mode shape component). The horizontal axis shows the standard deviation estimated using the frequentist approach. The vertical axis shows the average of the 100 standard deviations estimated by SSI-cov, with blue error bars indicating the 95 % confidence interval of those 100 values. The black line indicates the 1:1 reference. (c) Magnified version of Figure 3.7b for small standard deviations.

3.4.4 Estimation of the Foundation Parameters (model updating results)

Having obtained 100 datasets, the model updating frameworks described in Section 3.3 can be applied to estimate the uncertain model parameters. In this chapter, the structural parameters to be updated consist of the three compliance parameters required to define the foundation's macro-element: $\boldsymbol{\theta}_S = \{c_{11}, c_{22}, c_{12}\}$, where:

$$\mathbf{C}_{\text{macro}} = \mathbf{K}_{\text{macro}}^{-1} = \frac{1}{k_{11}k_{22} - k_{12}^2} \begin{bmatrix} k_{22} & -k_{12} \\ -k_{12} & k_{11} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} \\ c_{12} & c_{22} \end{bmatrix}. \tag{3.23}$$

For the deterministic case, this constitutes the set of uncertain parameters to be updated: $\boldsymbol{\theta}^{\text{det}} = \boldsymbol{\theta}_S$. However, for the Bayesian frameworks, the model error parameters are also considered as uncertain and included in the set of parameters to be updated. Therefore, for the vector and MAC based likelihood functions, the set of parameters to be updated are $\boldsymbol{\theta}^{\text{vec}} = \{\boldsymbol{\theta}_S; \sigma_f; \sigma_{\boldsymbol{\phi}}\}$ and $\boldsymbol{\theta}^{\text{MAC}} = \{\boldsymbol{\theta}_S; \sigma_f^G; \sigma_{\text{MAC}}^G\}$, respectively.

For the Bayesian frameworks, uniform priors with lower and upper bounds were assumed for each updated parameter. The bounds for each parameter are summarised in Table 3.1, with the bounds for the compliance parameters defined as multiples of the baseline compliance matrix $\mathbf{C}_{\text{macro}}^0$:

$$\mathbf{C}_{\text{macro}}^0 = \left(\mathbf{K}_{\text{macro}}^0\right)^{-1} = \begin{bmatrix} c_{11}^0 & c_{12}^0 \\ c_{12}^0 & c_{22}^0 \end{bmatrix}. \tag{3.24}$$

Table 3.1: Lower and upper bounds for each updated parameter in the Bayesian frameworks.

Parameter	Lower bound	Upper bound	Likelihood function type
c_{11}	$0.5c_{11,0}$	$2c_{11,0}$	Vector & MAC
c_{22}	$0.5c_{22,0}$	$2c_{22,0}$	Vector & MAC
c_{12}	$0.5c_{12,0}$	$2c_{12,0}$	Vector & MAC
σ_f	0	0.05	Vector
$\sigma_{\boldsymbol{\phi}}$	0	0.05	Vector
σ_f^G	0	0.05	MAC
σ_{MAC}^G	0	0.05	MAC

Results are presented for two cases: (i) when the modal parameters are obtained from

individual datasets, $\hat{\boldsymbol{\xi}}^{id}$, corresponding to the 'Individual datasets' label in Figure 3.7a; and (ii) when modal parameters are averaged across 100 datasets using a frequentist approach, $\bar{\boldsymbol{\xi}} = E\left[\hat{\boldsymbol{\xi}}^{id}\right] = \frac{1}{100}\sum_{id=1}^{100}\hat{\boldsymbol{\xi}}^{id}$, corresponding to the 'Frequentist' label in Figure 3.7a.

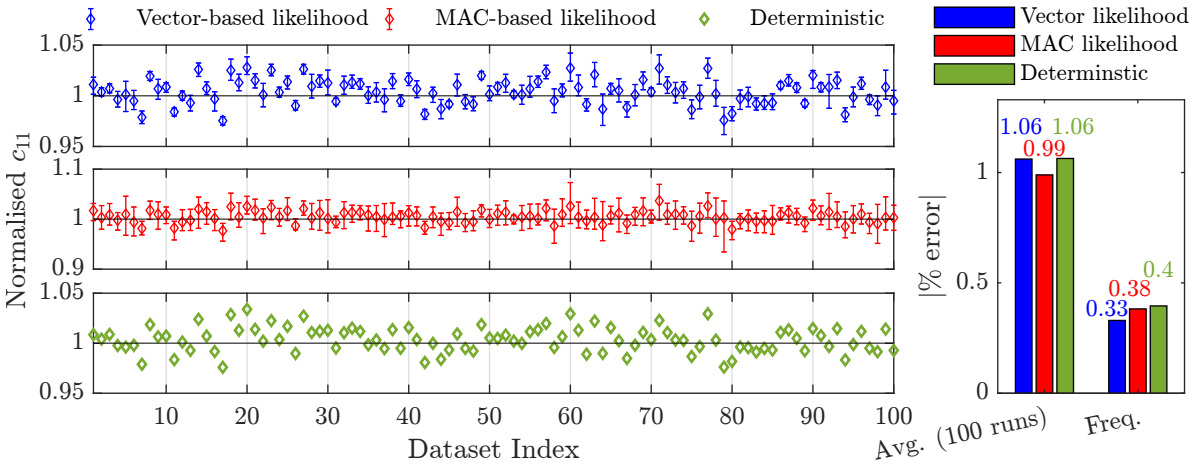
(i) Results from an unweighted objective function and undefined Bayesian measurement error

Initially, to assign equal importance to all modal parameters, the deterministic weighting factors were set to one and the Bayesian measurement error parameters were set to zero: $W_{f,i} = W_{\text{MAC},i} = 1;\ \boldsymbol{\Sigma}_D = 0;\ \&\ \sigma_{f,i}^D = \sigma_{\text{MAC},i}^D = 0$. However, later in this chapter, these parameters are defined based on either the uncertainty estimated by SSI-cov or a frequentist approach, in order to investigate their influence within the the model updating frameworks.

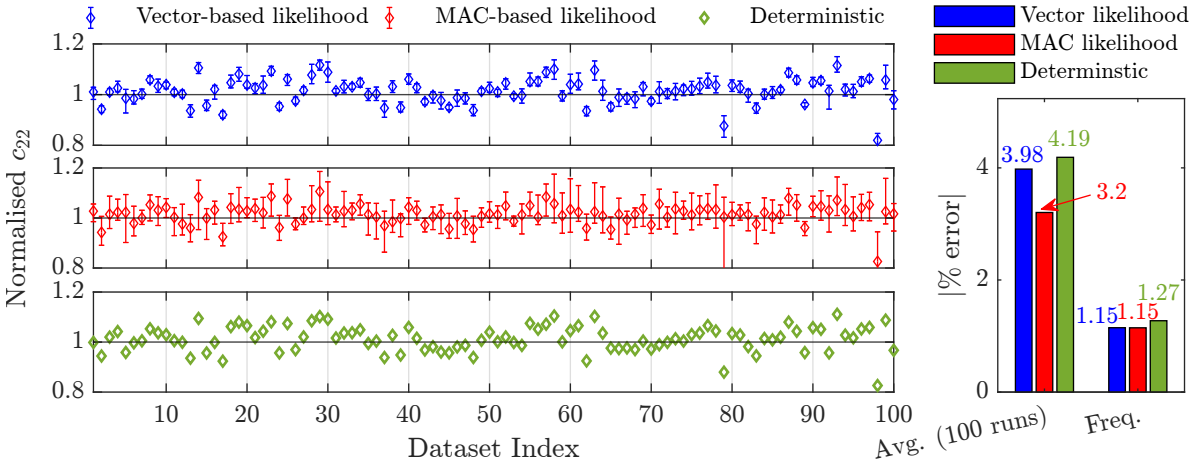
Deterministic vs Bayesian model updating

The updated foundation parameters for the deterministic case $\boldsymbol{\theta}_S^{\text{opt}}$ and the Bayesian *Maximum a Posteriori* (MAP) estimates $\boldsymbol{\theta}_S^{\text{MAP}}$ are shown in Figure 3.8. For the deterministic results, the true compliance values were used as the initial guess in the optimisation algorithm `fminsearch`, i.e., $\mathbf{C}_{\text{macro}}^{\text{init}} = \mathbf{C}_{\text{macro}}^{0}$. While this provides the deterministic approach with an unrealistic advantage, the figure demonstrates that the foundation parameters can be consistently estimated with minimal errors for all three model updating frameworks, with average errors in the compliance parameters less than 5 %. Moreover, a comparison between the absolute percentage errors shown in the bar plots in Figure 3.8 indicates that all three frameworks perform competitively in terms of their precision in estimating the foundation parameters. However, the deterministic method is significantly more efficient computationally, requiring approximately 10^2 model calls compared to 10^5–10^6 for the Bayesian frameworks.

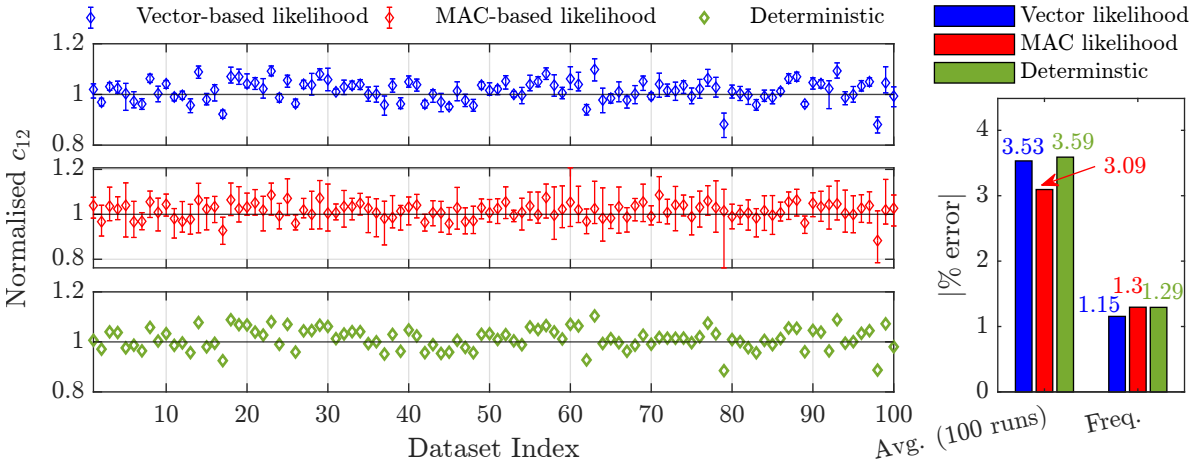
The bar plots in Figure 3.8 also compare the average absolute percentage error from 100 datasets (labelled 'Avg. (100 runs)') to the absolute percentage error from using a frequentist approach (labelled 'Frequentist'). Here it is observed that the errors in the compliance parameters are greater on average when only one dataset is available, as compared to using the average modal parameters across 100 datasets. This result can be explained by considering Figure 3.7a,



(a) Lateral compliance c_{11}



(b) Rotational compliance c_{22}



(c) Cross-coupling compliance c_{12}

Figure 3.8: Updated compliance estimates, normalised by their true values. Comparison between all three model updating frameworks where the error bars indicate a 95% confidence interval for the Bayesian frameworks. The corresponding % errors from averaging across 100 individual datasets (labelled 'Avg. (100 runs)') are compared to the frequentist approach (labelled 'Freq.') in bar plots on the right.

where it was shown that the error in the modal parameters from averaging all 100 datasets is smaller than that from any dataset individually. As a result, the foundation parameters are estimated with less error.

Although the deterministic approach performs comparably to the Bayesian frameworks when the initial guess is close to the true foundation values, Figure 3.9 demonstrates how its performance varies with different initial conditions. Here it can be seen that different results are obtained depending on whether $\mathbf{C}_{\text{macro}}^{\text{init}} = 0.5\mathbf{C}_{\text{macro}}^{0}$, $\mathbf{C}_{\text{macro}}^{\text{init}} = \mathbf{C}_{\text{macro}}^{0}$, or $\mathbf{C}_{\text{macro}}^{\text{init}} = 2\mathbf{C}_{\text{macro}}^{0}$. This implies that the deterministic objective function comprises many local minima. As a consequence, the algorithm can converge into one of those minima when the initial estimate used is in its near neighbourhood. Figure 3.9 also plots the objective function value at the updated parameter estimates: $J(\boldsymbol{\theta}^{\text{opt}})$. By comparing $J(\boldsymbol{\theta}^{\text{opt}})$, it can be concluded that the solution using $\mathbf{C}_{\text{macro}}^{\text{init}} = \mathbf{C}_{\text{macro}}^{0}$ yields the smallest objective function value, suggesting that this could be the global minimum. However, due to the large number of possible combinations for this foundation parameterisation, there exists no guarantee that another initial guess within the prior bounds would not yield a lower value. Furthermore, while the solution with the lowest objective function value produced the most accurate parameter estimates for this case study, this may not be the case in general (owing to modelling and identification errors). This undermines the level of trust that can be attributed to the deterministic results, motivating the use of Bayesian methods that jointly achieve global optimisation and uncertainty quantification when tasked with identifying the problem parameters (in this case the foundation stiffness properties).

The results of the Bayesian framework include the posterior probability distributions for the uncertain parameters, for which the 95 % confidence interval is graphically shown in the errors bars of Figure 3.8. However, Figure 3.10 and Figure 3.11 show the joint and marginal posterior distributions from the frequentist case where $\bar{\boldsymbol{\xi}} = E\left[\hat{\boldsymbol{\xi}}^{id}\right]$, for the vector-based and MAC-based likelihood functions, respectively. The width of the distributions in these figures allows the user to assess the identifiability of the results; in these examples, the narrow distributions highlight that the foundation parameters can be strongly identified from the modal data.

Additionally, the Bayesian approach provides insight to the error present in the frameworks

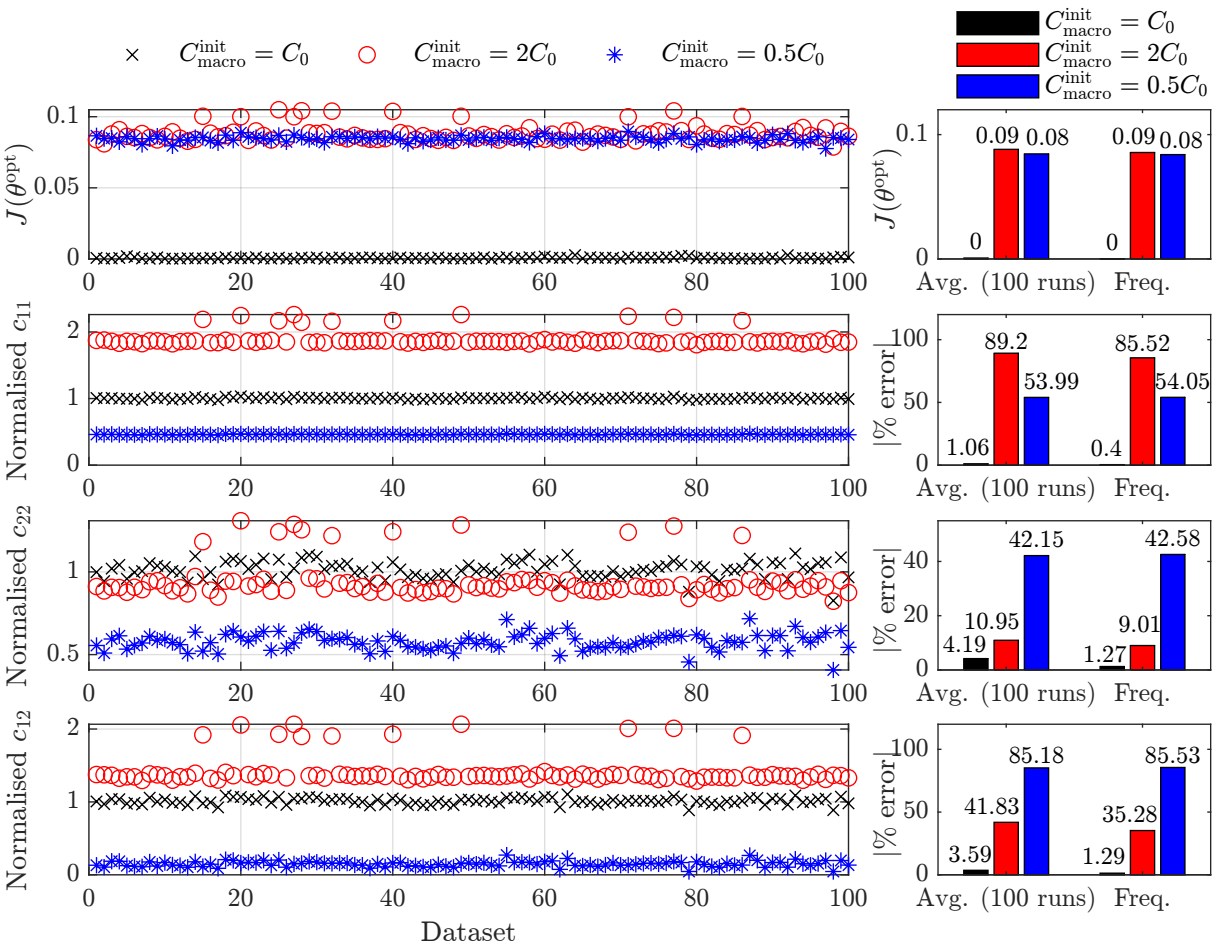


Figure 3.9: Updated compliance estimates when using the deterministic framework with different initial guesses. The % errors from averaging across 100 individual datasets (labelled 'Avg. (100 runs)') are compared to the frequentist approach (labelled 'Freq.') in bar plots.

by assessing the updated estimates for the model error parameters (σ_f, σ_ϕ, σ_f^G, & σ_{MAC}^G). These parameters reflect how well the model predicted modal parameters fit the measured modal data. In Figures 3.10 & 3.11, low values can be observed for these model error parameters, indicating a good fit between the model predictions and data. This was expected for this case study, as there was no model error present and minimal measurement error.

Vector vs MAC based likelihood functions

Figure 3.8 shows that in terms of MAP estimates, there is minimal difference between the two Bayesian frameworks. However, the error bars in Figure 3.8 represent the 95 % confidence intervals for the marginal distribution of each compliance parameter. Here it can be seen that the MAC-based likelihood function results in consistently wider posterior distributions for all datasets. This result can also be observed by comparing the posterior distributions for the fre-

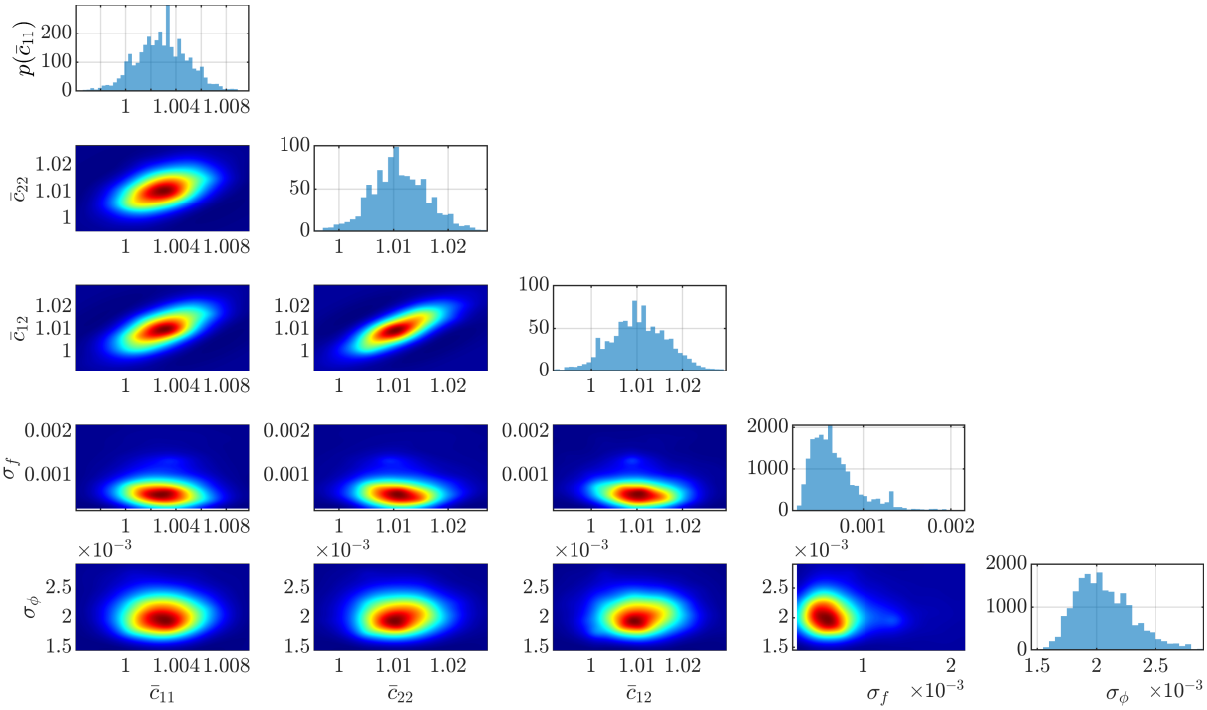


Figure 3.10: Posterior distributions for the **vector-based likelihood function** when the measurement uncertainty was undefined/neglected and the modal parameters' average from 100 datasets were used for model updating.

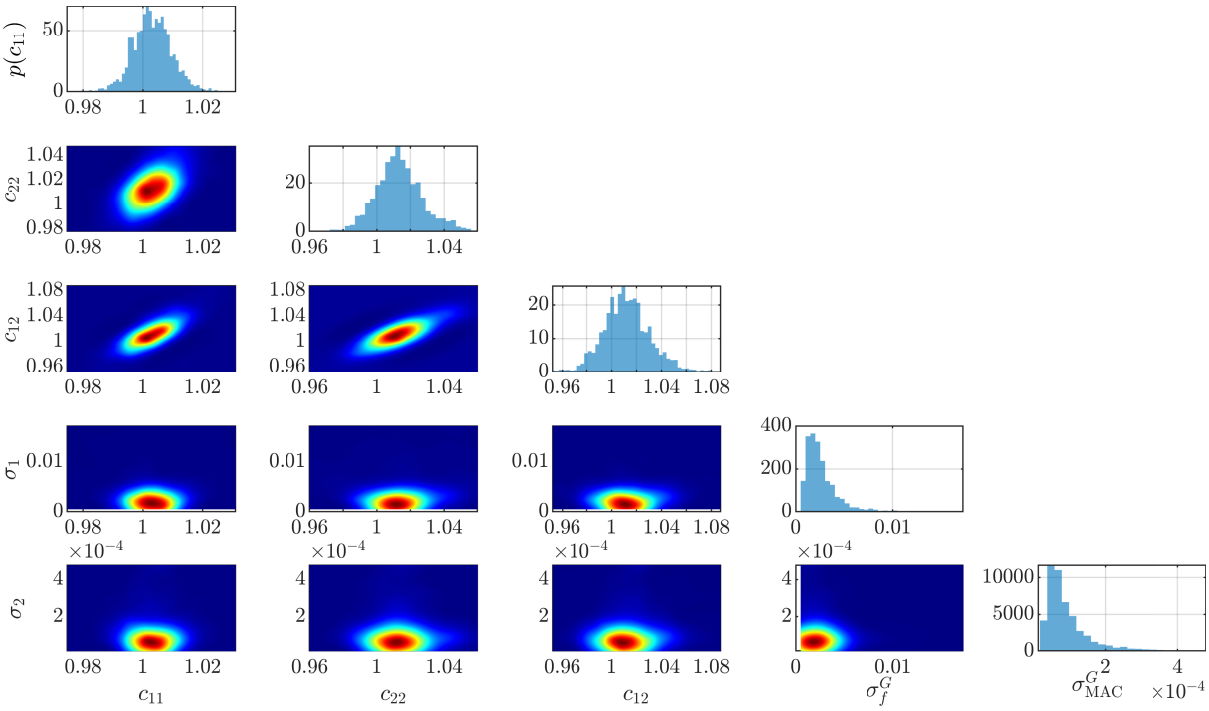


Figure 3.11: Posterior distributions for the **MAC-based likelihood function** when measurement uncertainty was undefined/neglected and the modal parameter average across 100 datasets was used for model updating.

quentist approach in Figures 3.10 & 3.11. This observation is linked to the number of terms in each framework's likelihood function. Although both frameworks use the same modes, the likelihood function for the MAC-based approach (Equations 3.13–3.15) only includes two terms for each mode: one frequency and one MAC value. In contrast, the vector-based approach has nine terms for each mode: one frequency and eight mode shape components (see Equations 3.10–3.11). As the number of terms in the likelihood function increases, the uncertainty in the posterior distributions decreases, resulting in narrower distributions for the vector-based framework.

Therefore, the vector-based approach yields updated foundation parameters with lower associated uncertainty. However, this may not always be beneficial. Due to the small bias that can be present in the measured modal parameters, the foundation estimates may also be biased. While this bias is minor in the context of uncertainty and safety factors in foundation design, an overly narrow posterior distribution may not capture the true foundation parameter values within a reasonable confidence interval. As a result, the vector-based approach often failed to capture the true values within a 95 % confidence interval. Specifically, out of the 100 datasets considered, the vector-based approach only captured the true values for all three compliance parameters 21 times. Whereas, the MAC-based approach captured all three true compliance values for 79 out of the 100 datasets, likely owing to the wider posterior distributions.

(ii) Results from a weighted objective function and defined Bayesian measurement error

The previous results were based on an unweighted deterministic objective function and Bayesian likelihood functions, where $W_{f,i} = W_{\text{MAC},i} = 1$ in Equation 3.7, $\boldsymbol{\Sigma}_D = 0$ in Equation 3.11, and $\sigma^D_{f,i} = \sigma^D_{\text{MAC},i} = 0$ in Equations 3.16 and 3.17. This approach effectively assigns equal weight to all terms in the objective or likelihood function during model updating. However, this method may inadvertently favour less accurate (or more uncertain) modal parameters, potentially resulting in greater errors in the foundation parameters. To address this issue, this section investigates the influence of incorporating estimated uncertainties in modal parameters, to define a weighted objective function or a likelihood function with fixed measurement error. The analysis uses the same datasets and model updating frameworks as those used in Section

3.4.4, with the modification that the deterministic weighting factors are now different for each term in the objective function and the Bayesian measurement error parameters are defined as non-zero: $W_{f,i} \neq 1$, $W_{\text{MAC},i} \neq 1$, $\mathbf{\Sigma}_D \neq 0$, $\sigma^D_{f,i} \neq 0$, and $\sigma^D_{\text{MAC},i} \neq 0$. These parameters were treated as fixed and are defined in the discussion below. For clarity, the approach introduced in this section will be referred to as the *weighted* approach to distinguish it from the *unweighted* approach previously considered.

Similar to Section 3.4.4 (i), model updating results were obtained both for the analysis of each of the 100 datasets individually and for the frequentist approach, where modal data were defined as the average across these 100 datasets. When each dataset was analysed individually, the estimated modal parameter uncertainty from SSI-cov was used. This provides a standard deviation for each frequency and mode shape component for the idth dataset:

$$\hat{\sigma}^{id}_{f_i} = \sqrt{\text{Cov}(\hat{f}_i)^{id}}; \tag{3.25}$$

$$\hat{\sigma}^{id}_{\phi_{i,j}} = \sqrt{\text{Cov}(\hat{\phi}_{i,j})^{id}}. \tag{3.26}$$

For the case where the modal parameters from all 100 datasets are averaged, the SSI-cov estimated uncertainty is ignored and the uncertainties from a frequentist approach are used instead:

$$\bar{\sigma}_{f_i} = E\left[\hat{f}_i^2\right] - E\left[\hat{f}_i\right]^2; \tag{3.27}$$

$$\bar{\sigma}_{\phi_{i,j}} = E\left[\hat{\phi}_{i,j}^2\right] - E\left[\hat{\phi}_{i,j}\right]^2. \tag{3.28}$$

For the vector-based Bayesian framework, these standard deviations can be used to define the diagonal terms of the measurement error covariance matrix $\mathbf{\Sigma}_D$. For the deterministic and MAC-based Bayesian frameworks, $\hat{\sigma}^{id}_{f_i}$ or $\bar{\sigma}_{f_i}$ were used to define $W_{f,i}$ and $\sigma^D_{f_i}$ as:

$$W_{f,i} = \frac{\hat{f}_i^{id}}{\hat{\sigma}^{id}_{f_i}} \quad \text{or} \quad W_{f,i} = \frac{E\left[\hat{f}_i^{id}\right]}{\bar{\sigma}^{id}_{f_i}}; \tag{3.29}$$

$$\sigma^D_{f_i} = \hat{\sigma}^{id}_{f_i} \quad \text{or} \quad \sigma^D_{f_i} = \bar{\sigma}_{f_i}. \tag{3.30}$$

However, since these frameworks rely on use of the MAC value, the uncertainty in each mode shape component had to be propagated to obtain a standard deviation for the MAC of each mode. This standard deviation was then used to compute $W_{\text{MAC},i}$ and $\sigma^D_{\text{MAC},i}$, following Algorithm 1.

Algorithm 1 Computation of MAC value weights and standard deviations

1: For each mode shape component, draw 10,000 samples from a normal distribution with a mean of $\hat{\phi}^{id}_{i,j}$ or $E\left[\hat{\phi}_{i,j}\right]$ and standard deviation $\hat{\sigma}^{id}_{\phi_{i,j}}$ or $\bar{\sigma}_{\phi_{i,j}}$;
2: Combine mode shape components to obtain 10,000 samples of the current mode shape;
3: For each sample, compute the MAC value between the sample and $\hat{\phi}^{id}_{i,j}$ or $E\left[\hat{\phi}_{i,j}\right]$ to obtain a histogram of MAC values for the current mode;
4: Find the truncated normal distribution $\mathcal{TN}(\text{MAC}; E\left[\text{MAC}\right], \sigma^2_{\text{MAC}}, 0, 1)$ that best fits the histogram by optimising for σ^2_{MAC};
5: Use the optimised σ^2_{MAC} to define $W_{\text{MAC},i}$ and $\sigma^D_{\text{MAC},i}$:

$$W_{\text{MAC},i} = \frac{1}{\sqrt{\sigma^2_{\text{MAC}}}}; \tag{3.31}$$

$$\sigma^D_{\text{MAC},i} = \sqrt{\sigma^2_{\text{MAC}}}. \tag{3.32}$$

Figure 3.12 provides an example for one dataset of the histograms obtained from propagating the sampled mode shape components to MAC values and compares this with the optimised truncated normal distribution. The figure shows that the truncated normal distribution provides a good fit to the histograms of MAC values, suggesting that the optimised values for σ^2_{MAC} are a reasonable indication of the measurement uncertainty in each MAC value.

Influence on the updated foundation parameters

Figure 3.13 shows the compliance error in the most probable estimates across all frameworks. It compares the average error in the compliance parameters calculated across individual datasets (labelled 'Avg. (100 runs)'), against the error when using the Frequentist-averaged modal parameters (labelled 'Frequentist'). For each case, results from the 'unweighted' approach—where measurement error was undefined—are compared against those from the 'weighted' approach, which incorporated either the SSI-cov estimated uncertainty (for individual datasets) or the frequentist uncertainty (for the averaged dataset). It can be seen that the weighted results reduce the errors in the compliance parameters for all methods. Overall, the best results were obtained

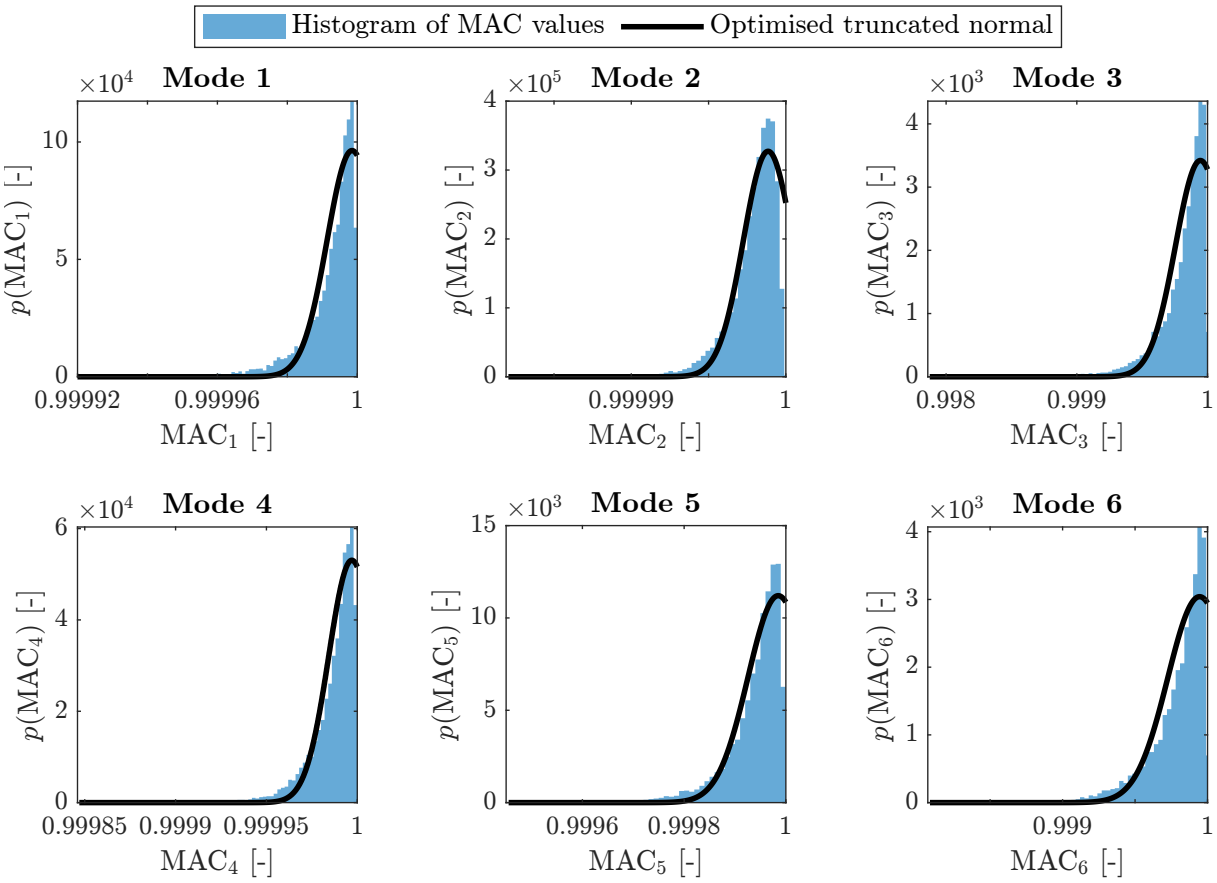


Figure 3.12: Histograms from propagating uncertainty from mode shape components to MAC values along with the optimised truncated normal for each mode.

when using a frequentist approach for both the modal parameters (Equation 3.21) and their uncertainty (Equations 3.27–3.28).

Although not shown here, the errors in the compliance parameters were not always improved for each of the individual datasets for the weighted case. This can be related back to Figure 3.7b and the fact that the uncertainty estimated by SSI-cov can vary significantly from one dataset to another, which may differ from that obtained using a frequentist approach. As a result, for some datasets—specifically, 16 %, 30 % and 29 % for the deterministic, MAC-based, and vector-based frameworks, respectively—incorporating the SSI-cov estimated uncertainty was found to increase the error in the compliance parameters. However, Figure 3.13 shows that, based on the average across 100 datasets, the compliance errors are reduced when the uncertainty from SSI-cov is used in the model updating frameworks.

The results for the deterministic framework were observed to show greater reduction in the compliance parameter error when using a weighted objective function, compared to the Bayesian

frameworks. One possible explanation for this could relate to the fact that the Bayesian frameworks already account for uncertainty by updating the model error parameters. Therefore, the explicit definition of the measurement uncertainty is less influential on the MAP estimates for the Bayesian frameworks than it is for the deterministic case. However, the weighted deterministic framework was still prone to finding false minima in the objective function, depending on the initial guess used in the optimisation process.

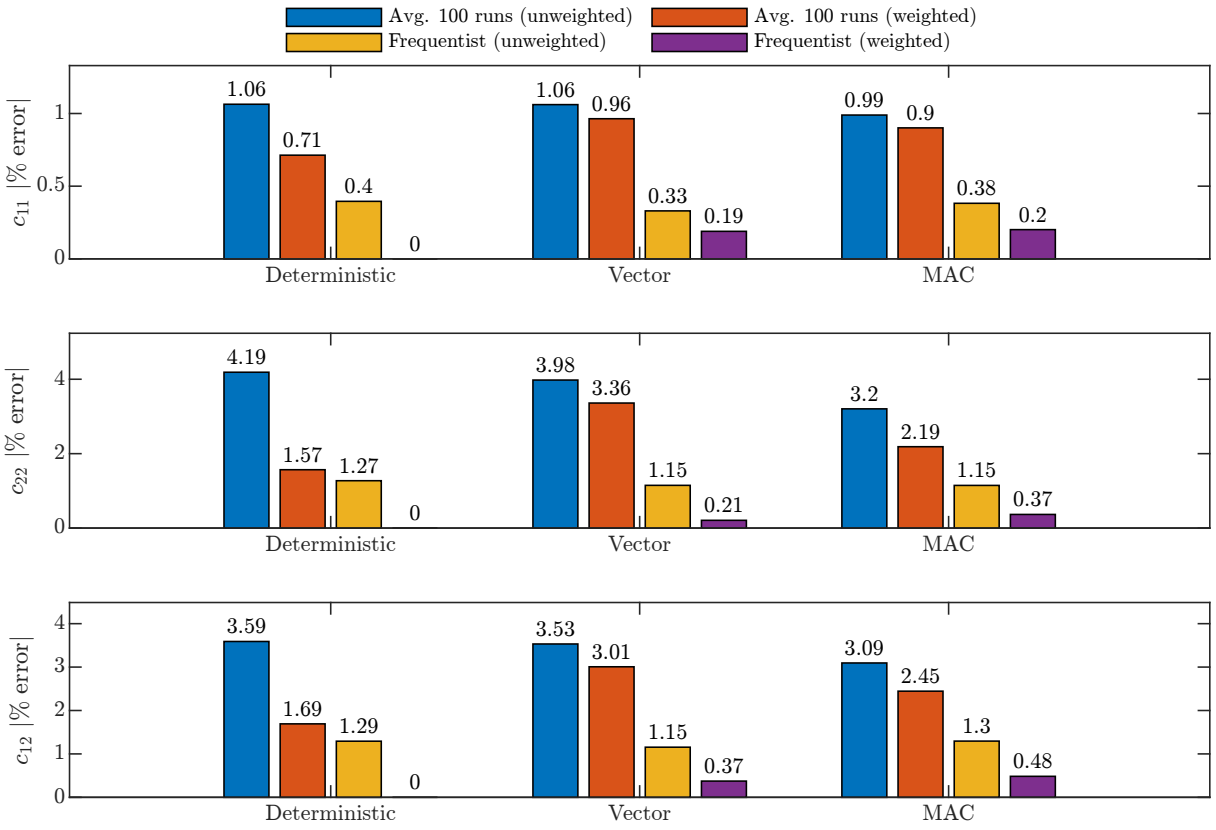


Figure 3.13: Most probable estimates for the compliance parameters for all methods. 'Avg. 100 runs' refers to the case when the most probable estimates are averaged across 100 individual datasets.

Influence on the model evidence and updated model error parameters

For the Bayesian methods, additional insight into the accuracy of the frameworks can be obtained by analysing the updated model error parameters estimates (σ_f & σ_ϕ or σ_f^G & σ_{MAC}^G). Moreover, one advantage of using iTMCMC to sample the posterior is that it provides an estimate for the model evidence, $p(\hat{\boldsymbol{d}}|\mathcal{M})$, as a by-product, which can be used for model class selection.

The updated model error parameters and model evidence are provided in Figure 3.14 for the

vector-based and MAC-based likelihood functions. It should be noted that, due to the different treatment of mode shapes between the vector-based and MAC-based likelihood functions, it is difficult to compare the model evidence between each likelihood function type. Similarly, due to the slightly different definitions of the model error parameters, these cannot be compared directly. However, within each likelihood function, the different cases can be considered for comparison.

When comparing results from the unweighted and weighted cases, the model evidence increases when the measurement error is provided, implying a better fit to the data for the weighted likelihood functions.

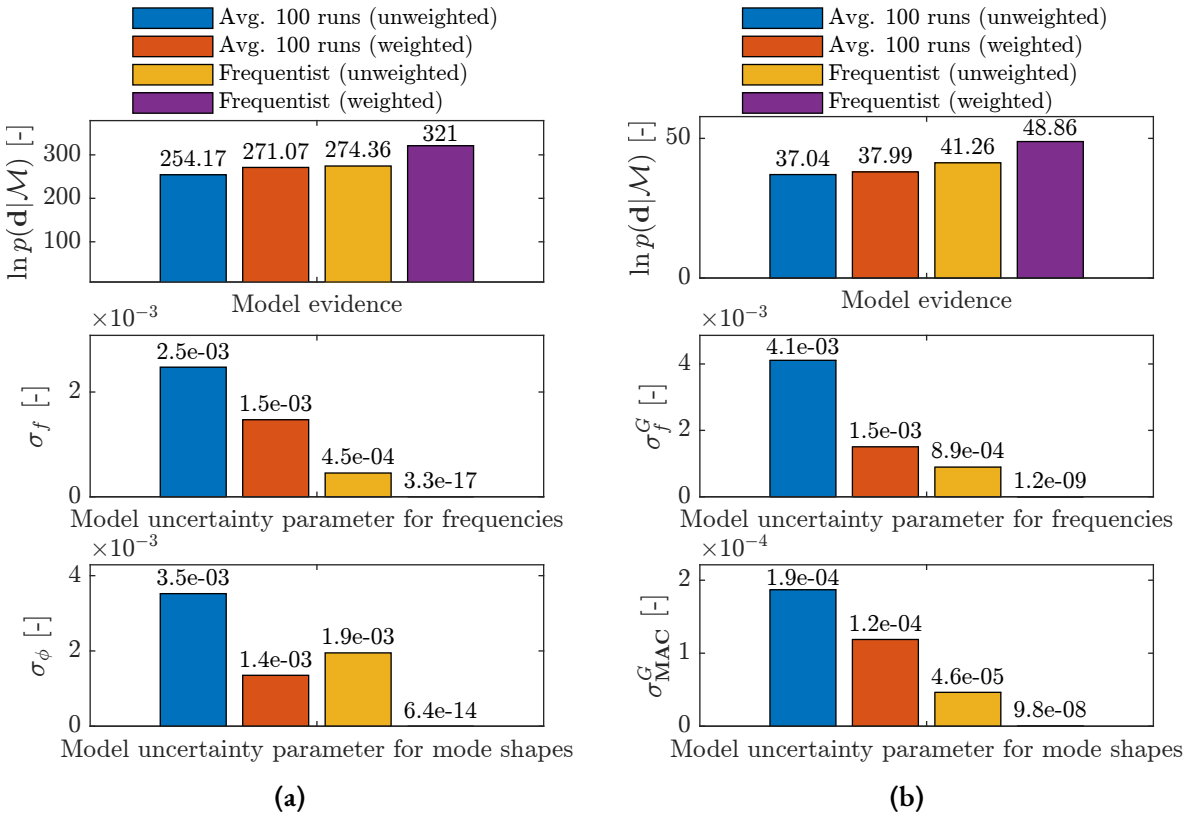


Figure 3.14: Model evidence and model error parameter MAP estimates for the Bayesian methods. 'Avg. 100 runs' refers to the case when the most probable estimates are averaged across 100 individual datasets. (a) Vector-based likelihood funciton (b) MAC-based likelihood function.

For all cases considered, Figure 3.14 indicates very low values for the updated model error parameters, implying minimal error in the model updating frameworks; an expected result due to the same model being used for data generation and model updating. However, when comparing results from the unweighted and weighted cases, it is observed that the model

error parameters are reduced for the weighted cases. This is because, when information on measurement uncertainty is provided, the model error parameters do not need to inflate to compensate for the lack of measurement error. In other words, the reduced model error parameters reflect less overall error in the frameworks, since the measurement uncertainty was already accounted for.

3.5 Summary

This chapter has compared various model updating frameworks for identifying the foundation parameters of offshore wind turbines, highlighting their benefits and limitations. A numerical case study was presented to assess the performance of a deterministic approach and two variants of a Bayesian framework. Realistic synthetic data were used, generated from a fully coupled hydro-aero-elastic offshore wind turbine model that accounts for the coupled dynamics of the tower and blades. Operational modal analysis was then performed using the SSI-cov method to obtain measured frequencies and mode shapes, which were subsequently compared with model-predicted values to update the compliance parameters of a macro-element spring foundation.

Addressing the research questions posed in the introduction, the following conclusions have been drawn:

Can the foundation parameters of an offshore wind turbine be accurately estimated using modal data?

> Yes, the updated estimates for the compliance parameters were consistently estimated with very small errors across all methods; the absolute percentage errors in the updated compliance parameters were less than 5%, which is considered very small in the context of offshore wind turbine foundation design. Furthermore, the Bayesian methods produced narrow posterior distributions, indicating that the foundation parameters were well-identified using modal data, even with a limited sensor setup.

What are the key differences between the various model updating frameworks in terms of most probable foundation estimates and the associated uncertainty?

The deterministic method performed competitively with the two Bayesian frameworks in terms of the errors in the updated compliance parameters. However, the deterministic method was susceptible to the initial guess provided by the user, indicating multiple minima in the objective function. This provided motivation for global optimisation methods for which the Bayesian frameworks fall into.

Both Bayesian frameworks yielded similar MAP estimates for the foundation parameters. However, the vector-based approach resulted in narrower marginal posterior distributions compared to the MAC-based approach, due to the larger number of terms in the likelihood function from individual mode shape components. While this provided smaller uncertainty in the foundation estimates, it often led to the true compliance parameters falling outside the 95 % confidence interval due to small measurement biases. In contrast, the wider posterior distributions from the MAC approach were found to capture the true values more frequently, whilst still identifying the foundation parameters well.

How does the uncertainty in the modal parameters, estimated using SSI-cov for an individual dataset, compare with the frequentist uncertainty across multiple datasets? Does incorporating the modal parameter's measurement uncertainty into the model updating frameworks improve the identifiability of the foundation parameters?

The modal parameters estimated by SSI-cov from individual datasets were identified with high accuracy. However, the slight variability in the bias of these modal parameters led to greater variability in the estimated foundation parameters. Whereas, averaging the modal parameters across 100 datasets reduced this measurement bias compared to individual datasets, resulting in more accurate foundation parameter estimates.

The uncertainty estimated from SSI-cov varied significantly across datasets, though it showed reasonable agreement with the uncertainty derived from a frequentist approach. As a result, incorporating this uncertainty into the model updating frameworks led to a slight reduction in compliance errors. However, the best

results were achieved when using the uncertainty from a frequentist approach.

Therefore, if multiple datasets are available, averaging the modal parameters and incorporating the frequentist uncertainty yields the most accurate foundation estimates. However, if only a single dataset is available, the uncertainty from SSI-cov can still provide a slight improvement in the foundation estimates on average. It should be noted that Hierarchical Bayesian model updating methods have been shown to provide better estimates of the parameter uncertainty when multiple datasets are available [98, 99]. However, such frameworks were outside the scope of this work.

Overall, this chapter has shown that limited modal data can be utilised to accurately estimate the foundation parameters of offshore wind turbines. The results presented here improve our understanding of the reliability and robustness of the model updating frameworks, ultimately contributing to safer and more efficient monitoring strategies within the industry. However, the numerical case study focused solely on measurement errors, whereas in practice, modelling errors are also expected. A significant source of such error may arise from the flexible blade properties, for which detailed and necessary design information is often unavailable due commercial confidentiality. Chapter 4 addresses this limitation by investigating the impact of modelling errors—introduced through assumptions about unknown blade properties—on the robustness of Bayesian model updating frameworks for foundation parameter identification.

CHAPTER 4

Modelling Challenges for Foundation Parameter Estimation

4.1 Introduction

Chapter 3 has shown the potential for foundation parameters to be identified from modal data and highlighted the advantages of using a Bayesian framework for this task. The examples accounted for typical measurement error present in practical applications—introduced through additive white noise in the acceleration response, as well as identification errors arising from operational modal analysis. However, the model was assumed to be perfect, which is an assumption that is rarely valid in practice. Therefore, this chapter builds upon the previous work by investigating the impact of modelling error on the identifiability of the macro-element foundation parameters.

A significant amount of design information is necessary to define the Simulink Model for identification of the foundation parameters. The geometric and material properties of the tower, monopile, and rotor nacelle assembly (RNA) are all required to accurately represent the dynamic behaviour of an offshore wind turbine. While developers typically have access to the design of the tower and monopile, they are often limited to lumped mass properties for the RNA, which may not be known with confidence.

To overcome this lack of information, assumptions must be made to model the blades as flexible components. A common approach is to scale known blade properties from a reference turbine to the specific turbine under consideration. However, this scaling inevitably introduces model error and its impact on the identifiability of the foundation parameters is not well reported. Moreover, the choice of reference turbine may also influence the results. Therefore,

this chapter investigates how different scaling methods from two reference turbines affect the identifiability of the macro-element foundation parameters within a Bayesian model updating framework. To isolate the effects of model error, this chapter does not consider any measurement error; i.e., modal data are not identified from operational modal analysis methods but instead are taken directly from the exact model predictions.

This chapter begins by introducing the baseline wind turbine model in Section 4.2, which provides the modal data for identification. Section 4.3 then outlines the various blade property scaling strategies considered, including their effects on the resulting modal parameters. Finally, Section 4.4 applies these scaling strategies to assess their impact on the identifiability of the foundation parameters.

4.2 Baseline Wind Turbine Model

To investigate the impact of modelling errors, an *exact* wind turbine model is first defined to serve as a baseline for comparison with the *scaled* blade models. For continuity, this baseline model was chosen to be the nearly identical to that used in Chapter 3 and described in Section 3.4.1. As a reminder, this model was based on the IEA Wind 15-MW reference turbine [10] with a monopile foundation, as shown in Figure 4.1 (reproduced for ease of reading). The same baseline macro-element stiffness matrix, $\mathbf{K}^0_{\text{macro}}$, defined in Equation 3.19, was also adopted in this chapter.

The only difference here to Chapter 3 relates to the number of reduced order elements used to model the blades. Previously, a single reduced order element with six external nodes for applying aerodynamic forces was used. In contrast, this chapter modelled each blade using ten reduced-order elements and eleven external nodes. Increasing the number of elements improves the accuracy of the blade model but comes at the cost of additional computational time. In Chapter 3, 100 forward simulations were required, so a simplified but faster model was adopted. However, since this chapter does not require multiple forward simulations for data generation, the more discretised blade model was favoured. Nevertheless, this baseline model used here, and the model used in Chapter 3 both implement the same IEA 15-MW blade property distributions [10] to define the elements' mass and stiffness matrices.

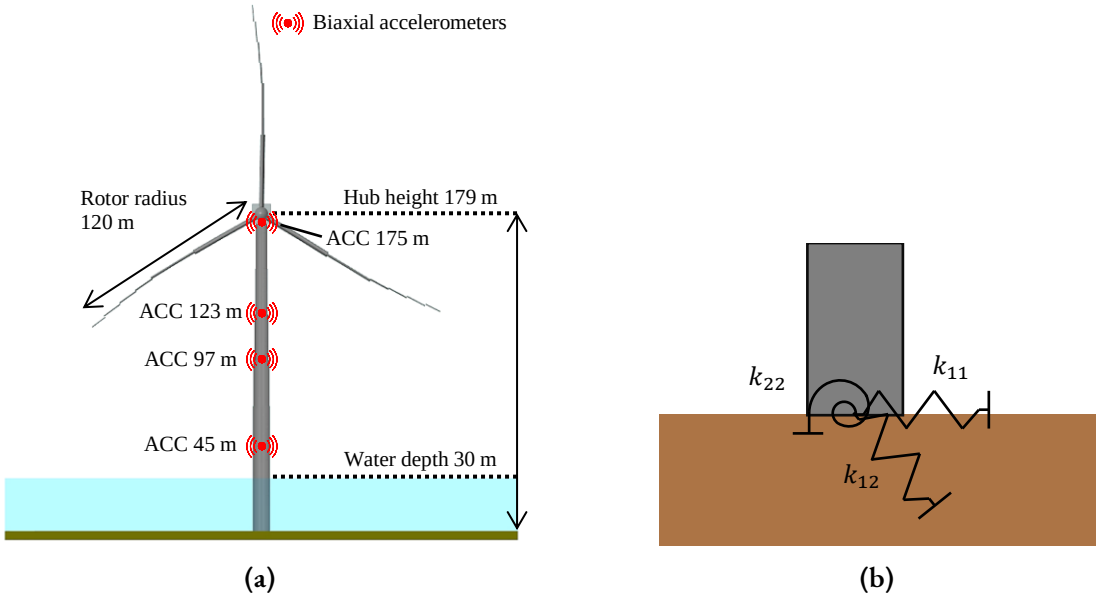


Figure 4.1: Baseline wind turbine model, reproduced from Chapter 3. (a) Baseline IEA 15-MW offshore wind turbine and accelerometer locations (b) Macro-element foundation representation.

4.3 Blade Scaling Methods

The geometry, mass, inertia and stiffness distributions with respect to length, required to model the blades, are listed in Table 4.1. To scale from a known reference turbine, two factors were introduced: a length scale factor λ_L and a mass scale factor λ_M. Then depending on the units of each blade distribution, a combined scale factor was derived as shown in Table 4.1.

Table 4.1: Blade property distributions used for flexible blade modelling, with corresponding units and scale factors.

Blade Property Distribution	Units	Scale Factor
Chord length	m	λ_L
Mass per unit length	kg/m	λ_M/λ_L
Mass moment of inertia per unit length	kg · m	$\lambda_M \cdot \lambda_L$
Axial stiffness EA	kg · m/s^2	$\lambda_M \cdot \lambda_L$
Bending stiffness EI	kg · m^3/s^2	$\lambda_M \cdot \lambda_L^3$

The length scale factor, λ_L, was defined as the ratio of the rotor diameters:

$$\lambda_L = \frac{D_{\text{target}}}{D_{\text{ref}}}, \tag{4.1}$$

where D_{target} and D_{ref} are rotor diameters of the target and reference turbines, respectively.

To define the mass scale factor, λ_M, two methods were considered, termed *Geometric Scaling* and *Blade Mass Scaling*. These approaches are outlined below and depend on the information available to the user.

1. **Geometric Scaling:**

   For this scaling method, it was assumed that only the rotor diameter of the target turbine is known. Assuming constant material density, the mass scale factor was defined geometrically in terms of λ_L:

   $$\lambda_M = \lambda_L^3. \tag{4.2}$$

   However, in practice, empirical trends have shown the blade mass typically scales with an exponent less than three [176–178]. As a result, the total blade mass is likely to be overestimated and this geometric scaling can only guarantee that the blade length is matched.

2. **Blade Mass Scaling:**

   In this approach, the total mass of the blades is also assumed to be known, in addition to the rotor diameter. Therefore, this approach guarantees matching the total blade mass and length. With this additional information, the mass scale factor was defined as:

   $$\lambda_M = \frac{M_{\text{target}}}{M_{\text{ref}}}, \tag{4.3}$$

   where M_{target} and M_{ref} are blade masses of the target and reference turbines, respectively.

4.3.1 Reference Turbines

Two reference turbines were used to scale up to the 15 MW baseline turbine: the NREL 5 MW reference turbine [8] and the IEA 10 MW reference turbine [9]. Based on these reference turbines and the two blade scaling methods, four wind turbines models were developed. The scale factors used to define the blade properties for each of these models are presented in Table 4.2.

Table 4.2: Definition of up-scaled models used for comparison with the baseline model to investigate the effects of blade modelling error.

Model Name	Scaling Method	Reference Turbine	λ_L	λ_M
Geom5	Geometric	NREL 5 MW	1.905	6.911
Geom10	Geometric	IEA 10 MW	1.212	1.781
BM5	Blade Mass	NREL 5 MW	1.905	3.776
BM10	Blade Mass	IEA 10 MW	1.212	1.397

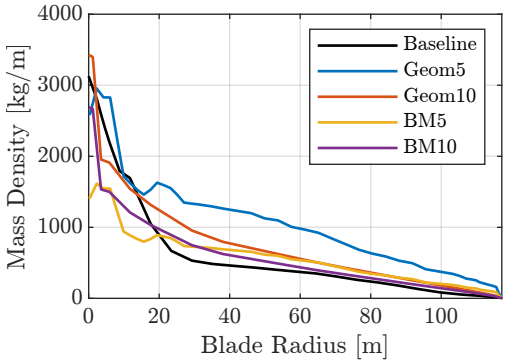
4.3.2 Influence on Blade Property Distributions

The blade property distributions for the baseline model and the four up-scaled blade models listed in Table 4.1 are compared in Figure 4.2. Overall, the up-scaled distributions more closely match the baseline when scaling from the 10 MW reference turbine than from the 5 MW. Additionally, the blade mass scaling approach appears to provide a closer fit to the baseline compared with the geometric method. However, since the cross-sectional properties at different locations along the blade in the reference distributions cannot be exactly transformed into those of the 'true' structure by using only two scaling parameters, there is no scaling factor that can result in the up-scaled models perfectly matching the baseline model.

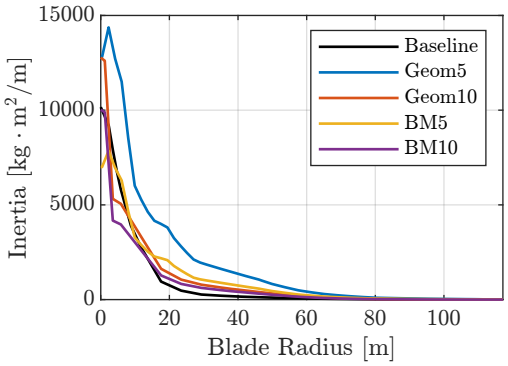
In fact, as shown in Figure 4.2, scaling can cause a property to be underestimated at one location along the length of a blade and overestimated at another. For example, the axial stiffness is overestimated at 20 m but underestimated at 60 m by all blade models. Similarly, within a given blade model, one property may match the baseline reasonably well while another deviates significantly. Therefore, all blade models contain varying levels of error which may introduce bias into the structure's modal parameters, which are used in the model updating frameworks to identify foundation parameters. The influence on the modal parameters is considered next.

4.3.3 Influence on Modal Parameters

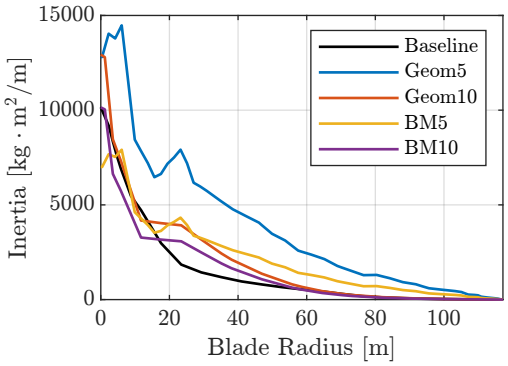
Before attempting to estimate the optimal foundation parameters within the Bayesian model updating frameworks, it is useful to investigate the bias that is introduced by each scaled blade model. This is explored within this subsection. The modal parameters of the full wind turbine model are computed using the true foundation parameters $\mathbf{K}^0_{\text{macro}}$ (Equation 3.19) for



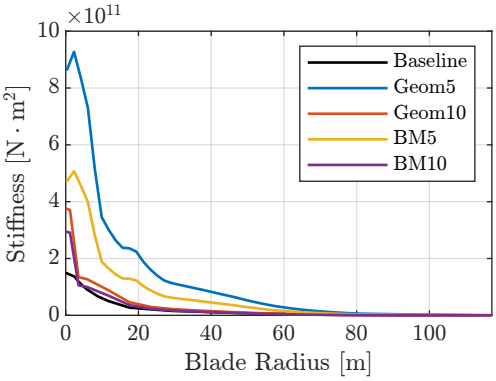
(a) Mass density



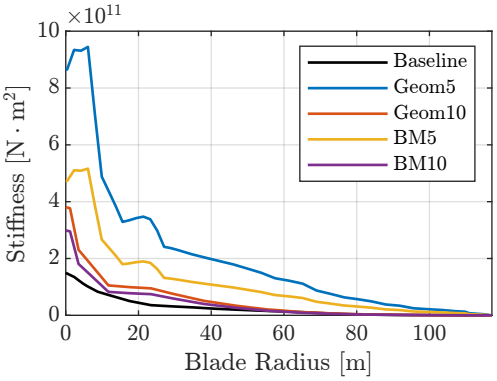
(b) Flapwise inertia



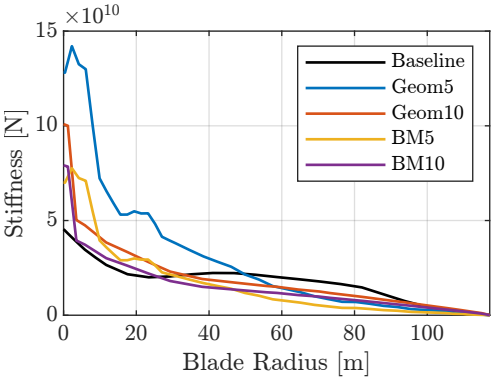
(c) Edgewise inertia



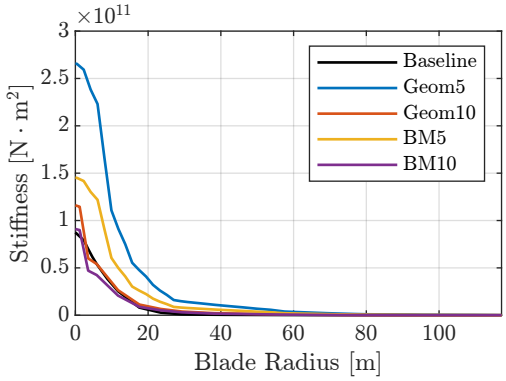
(d) Flapwise stiffness



(e) Edgewise stiffness



(f) Axial stiffness



(g) Torsional stiffness

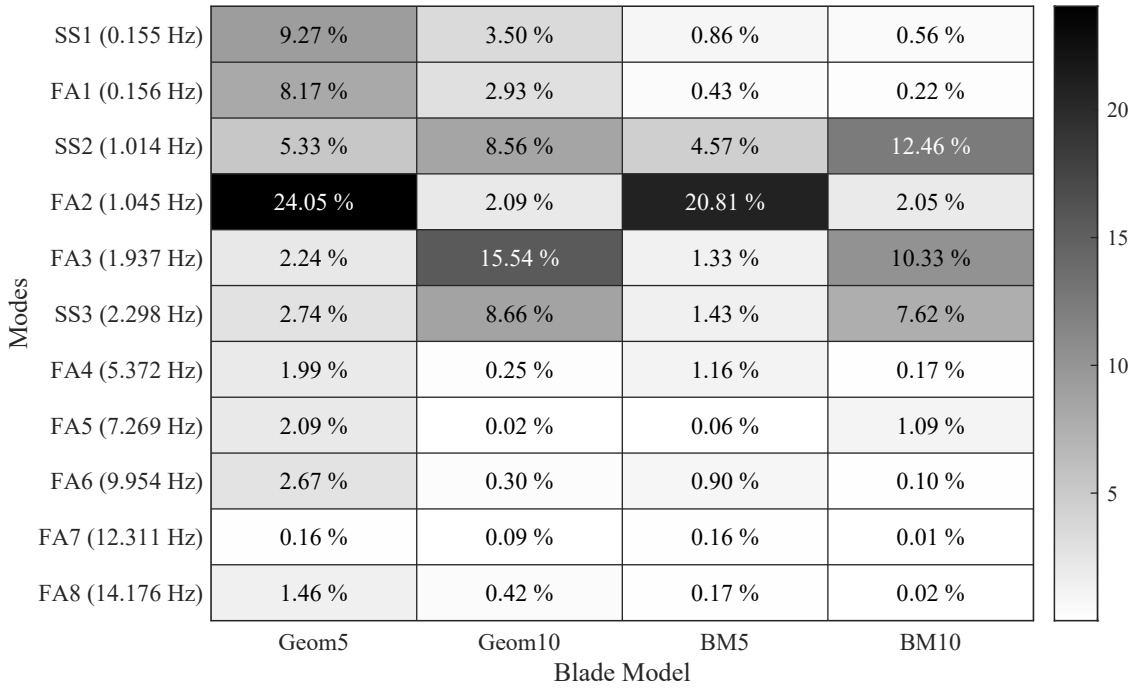
Figure 4.2: Blade property distributions for baseline and up-scaled model.

each model and compared to those obtained from the baseline turbine. It should be noted that each model must be linearised about a steady state operating point to obtain the modal parameters. In this chapter, this operating point corresponds to the steady state response to a constant 2 m/s wind field in still water. Since each model has different blade properties, the aerodynamic interaction with this wind field will vary, leading to a slightly different operating points for each case. Although this does affect the modal parameters, the impact was found to be minimal compared to the error introduced by scaling.

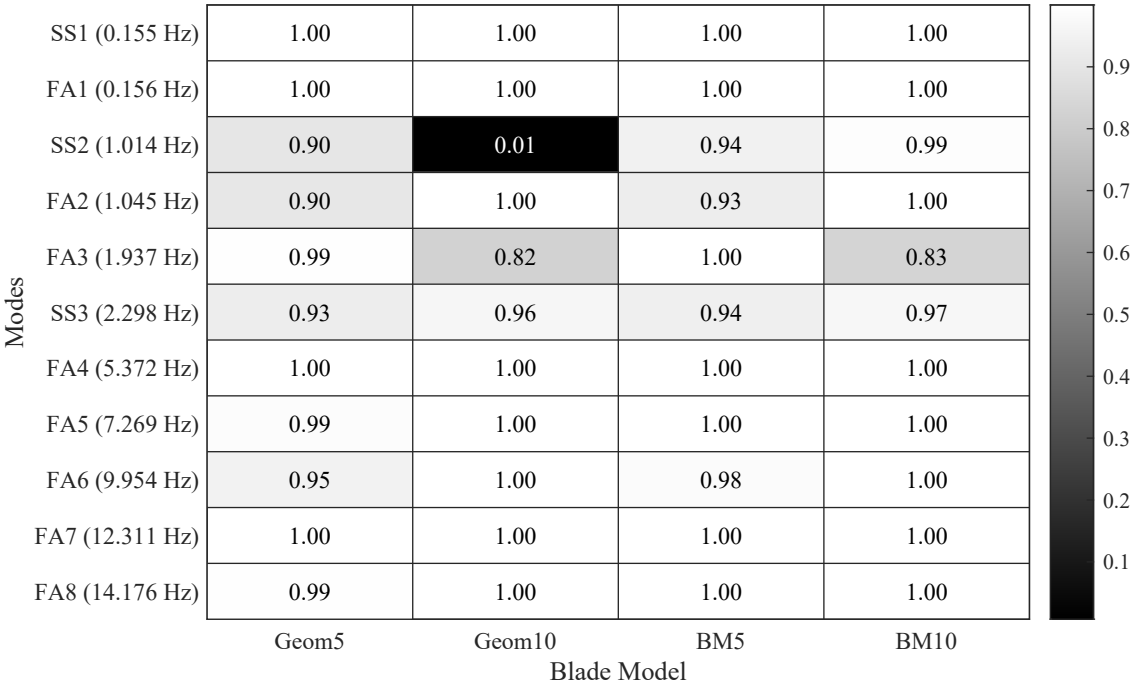
In Section 4.4, eleven modes are selected for parameter identification, corresponding to those with a significant contribution from the tower (as shown later in Figure 4.6). Therefore, only the error in these modal parameters is investigated here, as this represents the bias present within the model updating framework. The modes are labelled according to their dominant direction ('FA' = fore-aft or 'SS' = side-side), and numbered according to the order of the baseline turbine's frequencies. For example, the first fore-aft and first side-side modes are labelled 'FA1' and 'SS1', respectively.

Figure 4.3 presents heatmaps of the absolute percentage error in frequency and MAC value for each blade model relative to the baseline turbine, where darker shades indicate larger errors. Additionally, Figure 4.4 shows the mode shapes corresponding to the displacements of the tower for all four up-scaled blade models alongside the baseline model. From these figures, it can be observed that the second and third sets of modes (SS2, FA2, FA3, SS3) exhibit the largest errors across all models. This behaviour was expected, since these modes lie close in frequency to the blade modes, resulting in strong coupling between the tower and blades. This indicates that these modes are highly sensitive to both the reference turbine and blade scaling method. For the remaining modes, the errors are relatively small, suggesting that these modes are less sensitive to the blade model.

Comparing the absolute frequency errors across blade models reveals larger errors for the *Geom5* turbine. This was expected due to the consistent overestimation of blade property distributions shown in Figure 4.2. Another notable result is that the frequency error in the first pair of modes (SS1 and FA1) are reduced when using the blade mass scaled models (*BM5* and *BM10*). Again, this outcome was expected, as these modes rely less on the distributed



(a) Absolute % errors in frequency for turbines with different blade models.



(b) MAC values evaluated using tower mode shapes for turbines with different blade models.

Figure 4.3: Comparison of frequencies and mode shapes between the baseline turbine and turbines with the up-scaled blade models. Modal parameters were computed using the baseline macro-element foundation $\mathbf{K}^0_{\text{macro}}$.

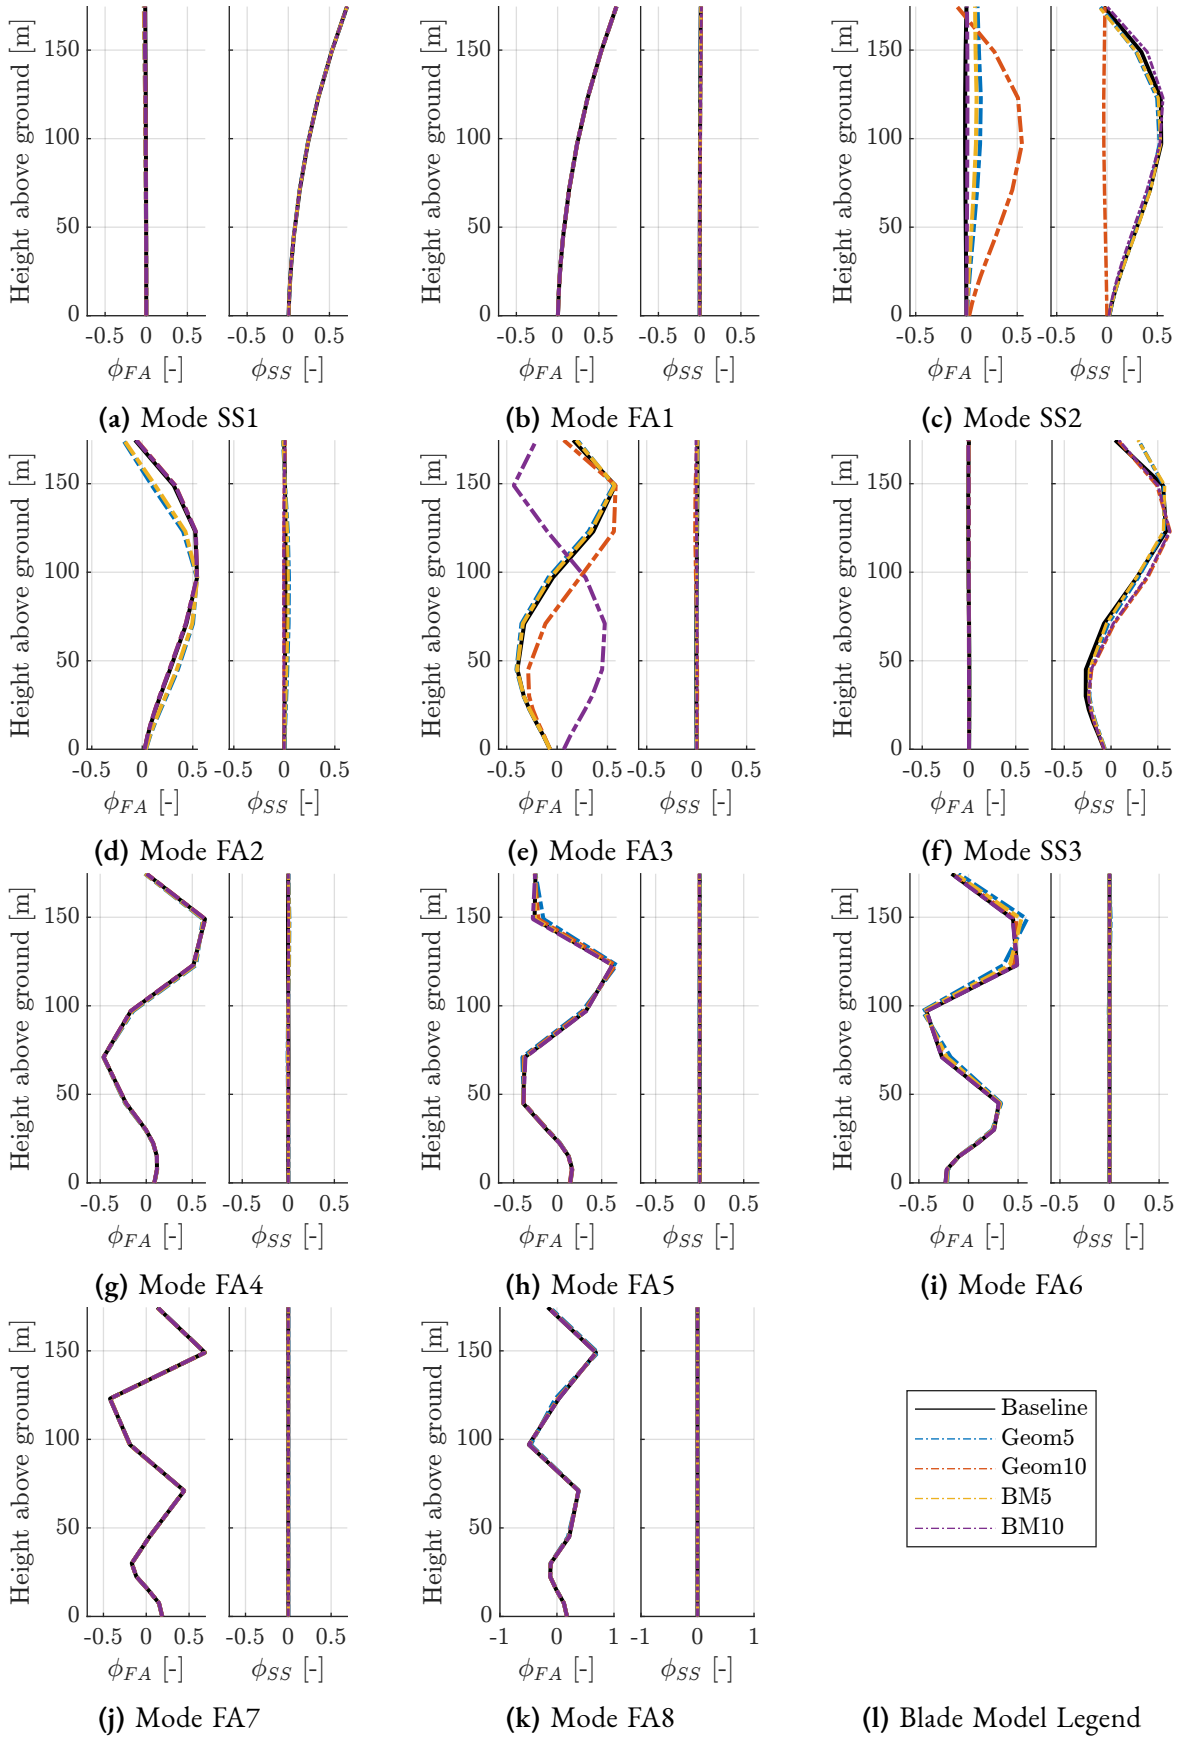


Figure 4.4: Tower mode shapes for turbines with different blade models. The foundation is defined using $\mathbf{K}^{0}_{\text{macro}}$.

properties and more on lumped mass of the RNA, which is only matched using the blade mass scaled models.

Overall, the relationship between the choice of reference turbine and scaling method with the resulting modal parameters is highly complex. There is no single approach that clearly outperforms the others, further indicating that these approaches are reasonable approximations in the absence of more detailed information. The modal parameters can be influenced in different ways depending on many modelling assumptions required by the user. As a result, all the up-scaled blade models considered here introduce non-negligible model bias, which may in turn impact the identifiability of the foundation parameters. The impact is investigated in the following section.

4.4 Identification of Foundation Parameters

In this section, the macro-element foundation parameters (see Figure 4.1b) are identified using the vector-based Bayesian model updating framework introduced in Section 3.3.2. Modal data are taken from the baseline wind turbine, and the four up-scaled blade models are used for identification. No measurement noise is added to the modal data, in order to focus on the effects of modelling error. The choice of mode set is discussed further in Section 4.4.2. It should be noted that the MAC-based likelihood function (also introduced in Section 3.3.2) was found to produce similar conclusions to the vector-based approach and to those already discussed in Chapter 3. Therefore, results are presented here for the vector-based approach only.

4.4.1 Macro-element Parameterisations and Prior Distributions

Following the vector-based framework, the compliance parameters (c_{11}, c_{22}, and c_{12}) of the macro-element foundation are updated alongside the likelihood function model error parameters (σ_f and σ_ϕ). Uniform distributions were assigned to each parameter, with lower and upper bounds listed in Table 4.3. The compliance bounds are defined relative to the baseline

macro-element foundation:

$$\mathbf{C}^0_{\text{macro}} = \begin{bmatrix} c^0_{11} & c^0_{12} \\ c^0_{12} & c^0_{22} \end{bmatrix} = \begin{bmatrix} 0.3033 & 0.0190 \\ 0.0190 & 0.0031 \end{bmatrix} \times 10^{-9}, \tag{4.4}$$

where the units of c^0_{11}, c^0_{22}, and c^0_{12} are m/N, rad/(Nm) and m/(Nm), respectively. These priors were selected to be as uninformative as possible, both to reflect the realistic uncertainty in the parameters and to assess if the compliance parameters can be identified from the likelihood function alone in the presence of model error.

However, the use of uniform priors leads to equal weighting across a wide range of macro-element configurations, including some that may be physically unrealistic. Therefore, a more informative prior was also considered. The assumption underlying this new prior, is that if the stiffness of one spring increases, then the others should increase as well. In other words, higher probability should be assigned to an equally scaled foundation such as $3 \times [c^0_{11},\ c^0_{22},\ c^0_{12}]^T$, than to an unevenly scaled foundation such as $[0.25c^0_{11},\ 4c^0_{22},\ c^0_{12}]^T$.

To implement this prior, the macro-element compliances were reparameterised using three new variables, γ, d_1 and d_2:

$$\begin{bmatrix} c_{11} \\ c_{22} \\ c_{12} \end{bmatrix} = \left(\gamma \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} + \frac{d_1}{\sqrt{2}} \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix} + \frac{d_2}{\sqrt{6}} \begin{bmatrix} 1 \\ -2 \\ 1 \end{bmatrix} \right) \times \begin{bmatrix} c^0_{11} \\ c^0_{22} \\ c^0_{12} \end{bmatrix}, \tag{4.5}$$

where γ uniformly scales the baseline compliance values, and d_1 and d_2 define the magnitude of deviation vectors, that are orthogonal to the initial compliance vector $[c^0_{11},\ c^0_{22},\ c^0_{12}]^T$. This prior formulation will be referred to as the γ-$\boldsymbol{d}$ parameterisation throughout this thesis, whereas the previous formulation will be referred to as the compliance parameterisation.

A uniform prior was assigned to γ, while d_1 and d_2 were defined by zero-mean normal distributions, as listed in Table 4.3. These priors allow the user to control how much confidence they place in the initial relative compliance values through the standard deviations assigned to d_1 and d_2 (σ_{d_1} and σ_{d_2}). For example, as these standard deviations tend to zero, the framework

Table 4.3: Prior distributions assigned to each parameter for both foundation parameterisations. $\mathcal{U}(a, b)$ denotes a uniform distribution with lower bound a and upper bound b. $\mathcal{N}(\mu, \sigma)$ denotes a normal distribution with mean μ and standard deviation σ.

Compliance Parameterisation		**γ-$\boldsymbol{d}$ Parameterisation**	
Parameter	Distribution	Parameter	Distribution
c_{11}	$\mathcal{U}(0.25c_{11}^0,\ 4c_{11}^0)$	γ	$\mathcal{U}(0.25, 4)$
c_{22}	$\mathcal{U}(0.25c_{22}^0,\ 4c_{22}^0)$	d_1	$\mathcal{N}(0,\ 0.01)$
c_{12}	$\mathcal{U}(0.25c_{12}^0,\ 4c_{12}^0)$	d_2	$\mathcal{N}(0,\ 0.01)$
σ_f	$\mathcal{U}(0, 0.1)$	σ_f	$\mathcal{U}(0,\ 0.1)$
σ_ϕ	$\mathcal{U}(0, 0.1)$	σ_ϕ	$\mathcal{U}(0,\ 0.1)$

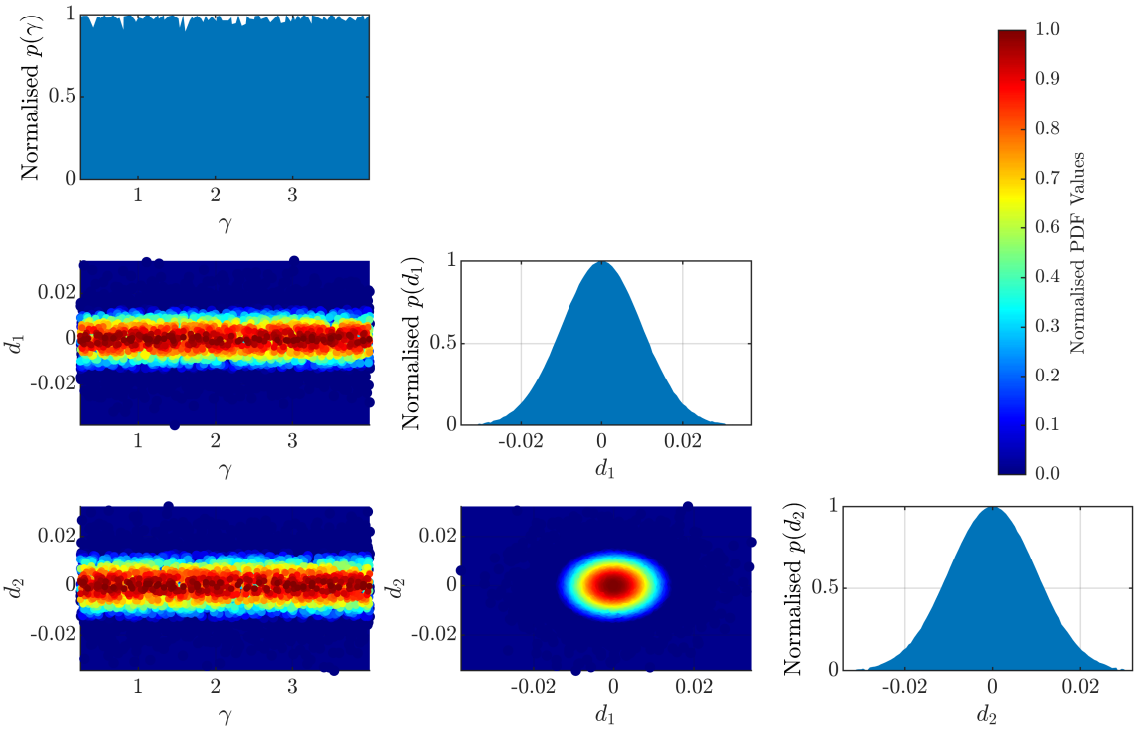
effectively reduces to the identification of a single scaling parameter γ. Conversely, as σ_{d_1} and σ_{d_2} tend to infinity, the prior becomes equivalent to the original compliance parameterisation.

It should also be noted that only the direction of the initial vector $[c_{11}^0,\ c_{22}^0,\ c_{12}^0]^T$ matters, not its magnitude, as this is scaled by the uniformly distributed parameter γ. For example, specifying $[c_{11}^0,\ c_{22}^0,\ c_{12}^0]^T = [1\ 2\ 3]^T$ with $\gamma = 1$ would assign equal prior probability to an identical foundation such as $[c_{11}^0,\ c_{22}^0,\ c_{12}^0]^T = [2\ 4\ 6]^T$ with $\gamma = 0.5$, assuming constant d_1 and d_2. To obtain the results presented in this chapter, the initial compliance vector was set to 0.8 times the baseline values shown in Equation 4.4—i.e., the direction of the compliance vector was assumed to be known exactly but with an incorrect magnitude. Consequently, a low standard deviation of 0.01 was assigned to both σ_{d_1} and σ_{d_2}.

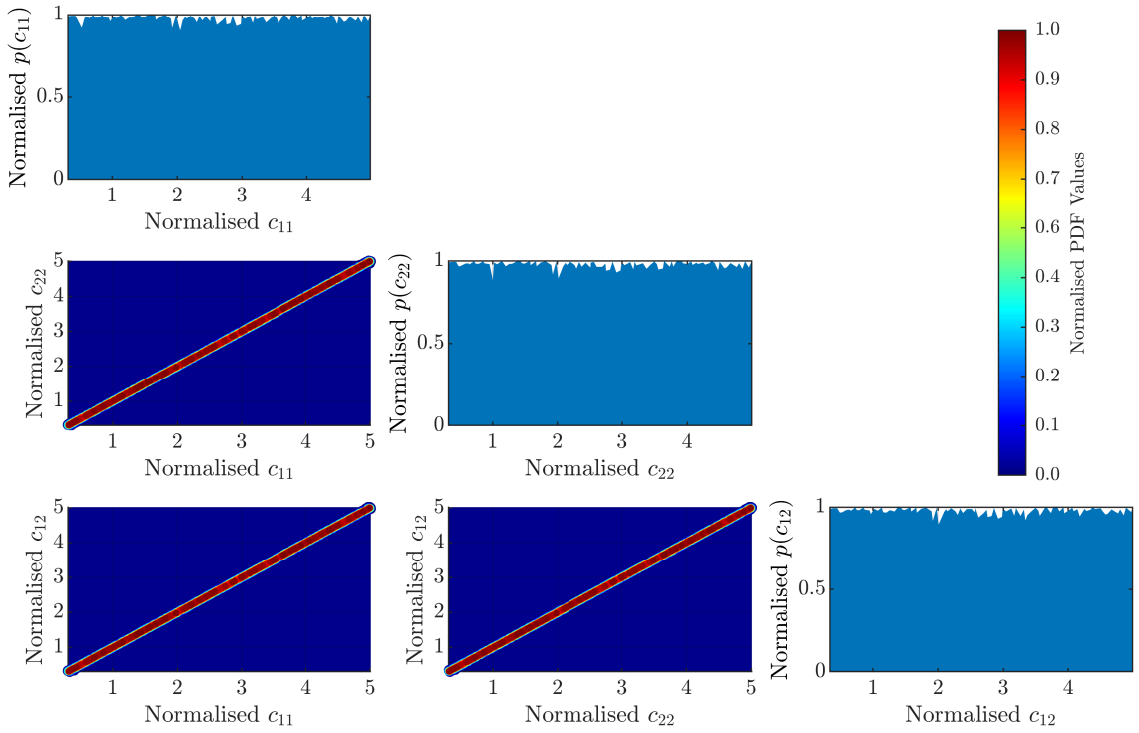
Figure 4.5 illustrates the prior distribution for the γ-$\boldsymbol{d}$ parameterisation, shown both in the original space (Figure 4.5a) and mapped to the compliance space (Figure 4.5b). The figure in the compliance space highlights how the γ-$\boldsymbol{d}$ prior formulation penalises foundation parameterisation that deviate orthogonally from the initial compliance vector—i.e., the regions in the upper left and lower right of the joint distributions in Figure 4.5b are assigned zero prior probability, and therefore zero posterior probability.

4.4.2 Modal Data: Mode Set Selection

The modal parameters used to define the likelihood function (Equation 3.11) were obtained by linearising the baseline wind turbine model about the operating point corresponding to



(a)



(b)

Figure 4.5: Prior distribution for γ-$\boldsymbol{d}$ macro-element parameterisation in the (a) original γ-$\boldsymbol{d}$ space and (b) mapped to the compliance space.

the steady state response under a constant 2 m/s wind field in still water. The main difference between the modal data here to those used in the previous chapter, is that the data were obtained directly from the model instead of using operational modal analysis on simulated data—i.e., the modal data are free of measurement and estimation error. However, to maintain consistency with the previous chapter and avoid introducing variation due to sensor placement, the mode shape components corresponding to the same four biaxial tower accelerometers shown in Figure 4.1a were used.

Due to the complexity of the Simulink model, there are a large number of states, resulting in many model-predicted modes. However, not all of these modes can be practically identified from typical tower accelerometers using operational analysis methods. Therefore, to determine a suitable subset of modes to be used for model updating, a forward simulation was run to generate synthetic data under turbulent wind and irregular wave loading. The wind field and JONSWAP wave spectrum were defined using the same hyperparameters as in the previous chapter: a mean hub height wind speed of 2 m/s with 10 % turbulence intensity, $H_s = 3.5\,\mathrm{m}$ and $T_p = 8\,\mathrm{s}$. From the resulting response, the average power spectral density of the four biaxial accelerometers was plotted, as shown in Figure 4.6. The modes corresponding to the dominant spectral peaks were considered identifiable and were selected as candidates for model updating. This process resulted in a subset of eleven modes, which were introduced in Section 4.3.3.

In Section 4.3.3, the error introduced by each scaled blade model on these eleven modes was presented. It was shown that the second and third mode pairs (SS2, FA2, FA3 and SS3) were particularly sensitive to blade modelling, resulting in significant error. Therefore, to investigate the impact of mode set selection on foundation parameter identification, two sets of modes were considered for model updating: one set including all eleven modes, and another excluding the four blade-model-sensitive modes. These mode sets are summarised in Table 4.4.

Table 4.4: Mode sets used for model updating to identify the foundation parameters.

Mode Set Name	Mode Set
Full	All eleven modes (see Figures 4.3 and 4.4)
Reduced	All eleven modes excluding modes SS2, FA2, FA3 and SS3.

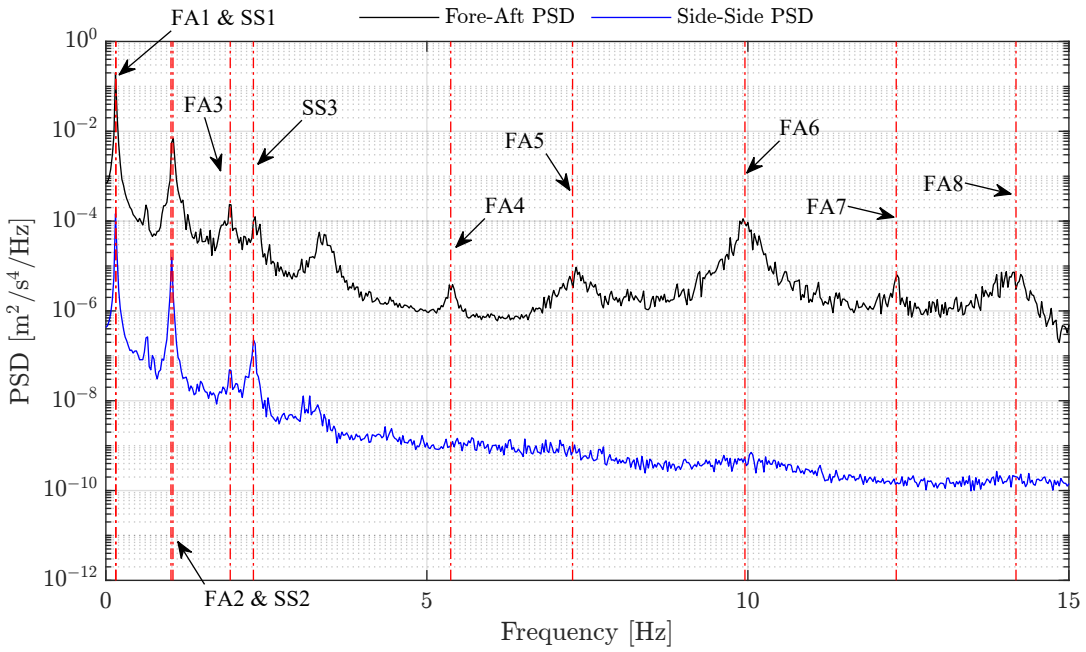


Figure 4.6: Averaged fore-aft and side-side accelerometer power spectral densities (PSDs) from a turbulent forward simulation. The modes selected for model updating are indicated by the vertical lines and labelled corresponding to Figures 4.3 and 4.4.

4.4.3 Results and Discussion

Using the baseline model's modal parameters as 'measurements', posterior distributions were obtained for all combinations across the four scaled models (Table 4.2), the two foundation parameterisations (Table 4.3) and two mode sets (Table 4.4). This resulted in 16 sets of posterior distributions for the uncertain parameters. These distributions are presented in Figure 4.7, visualised as bar plots of the MAP estimates (Figure 4.7a) and error bars representing the 95 % confidence intervals of each parameter's marginal distribution (Figure 4.7b). It should be noted, that the posterior samples of the γ-$\boldsymbol{d}$ parameterisations are transformed to the compliance space using Equation 4.5, to allow a direct comparison between the two foundation parameterisations.

The performance of the framework is primarily assessed based on how accurately the foundation parameters are identified, rather than the modal parameters. However, a correlation was observed whereby lower compliance errors corresponded to lower errors in the model-predicted modal parameters for the MAP estimates. For the interested reader, the frequency errors and MAC values for each case are tabulated in Appendix A.

The following sections discuss the influence of blade scaling method, reference turbine,

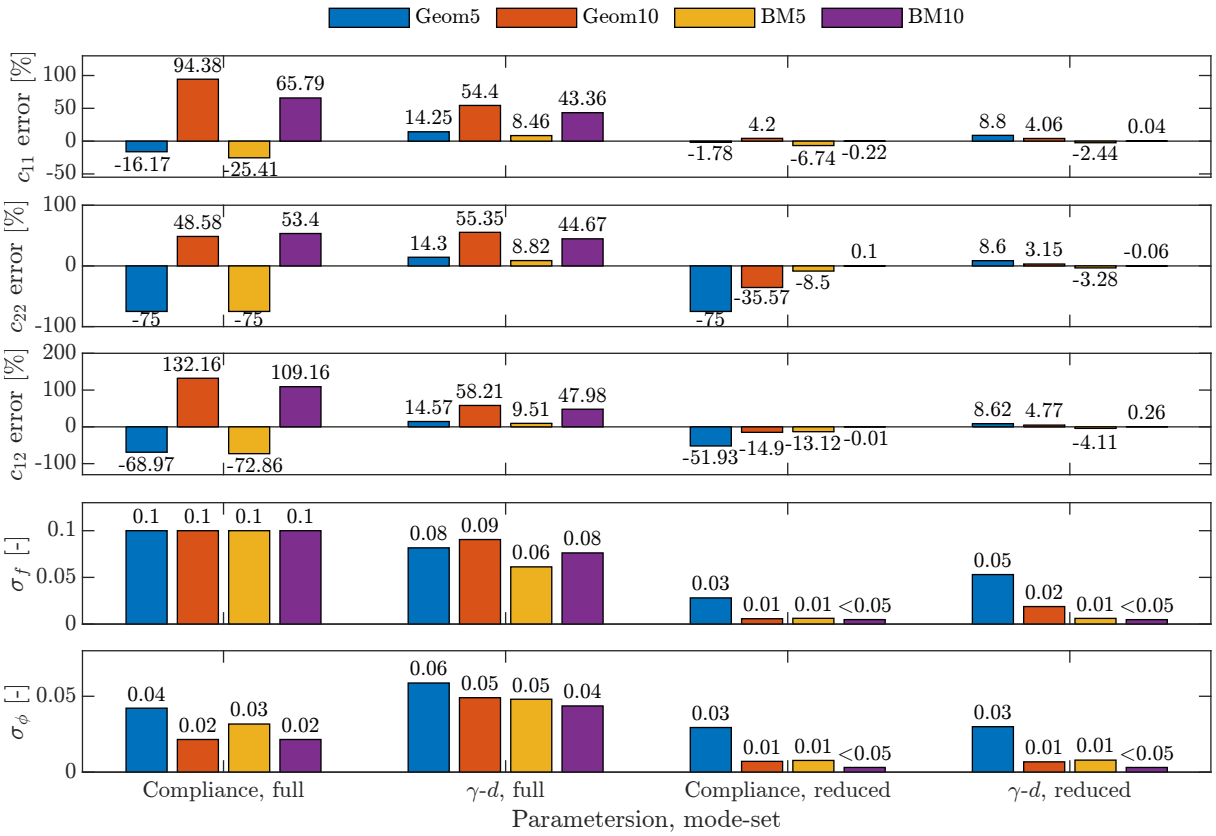
choice of mode set and parameterisation on the identifiability of the macro-element foundation.

(i) Influence of chosen mode set and foundation parameterisation

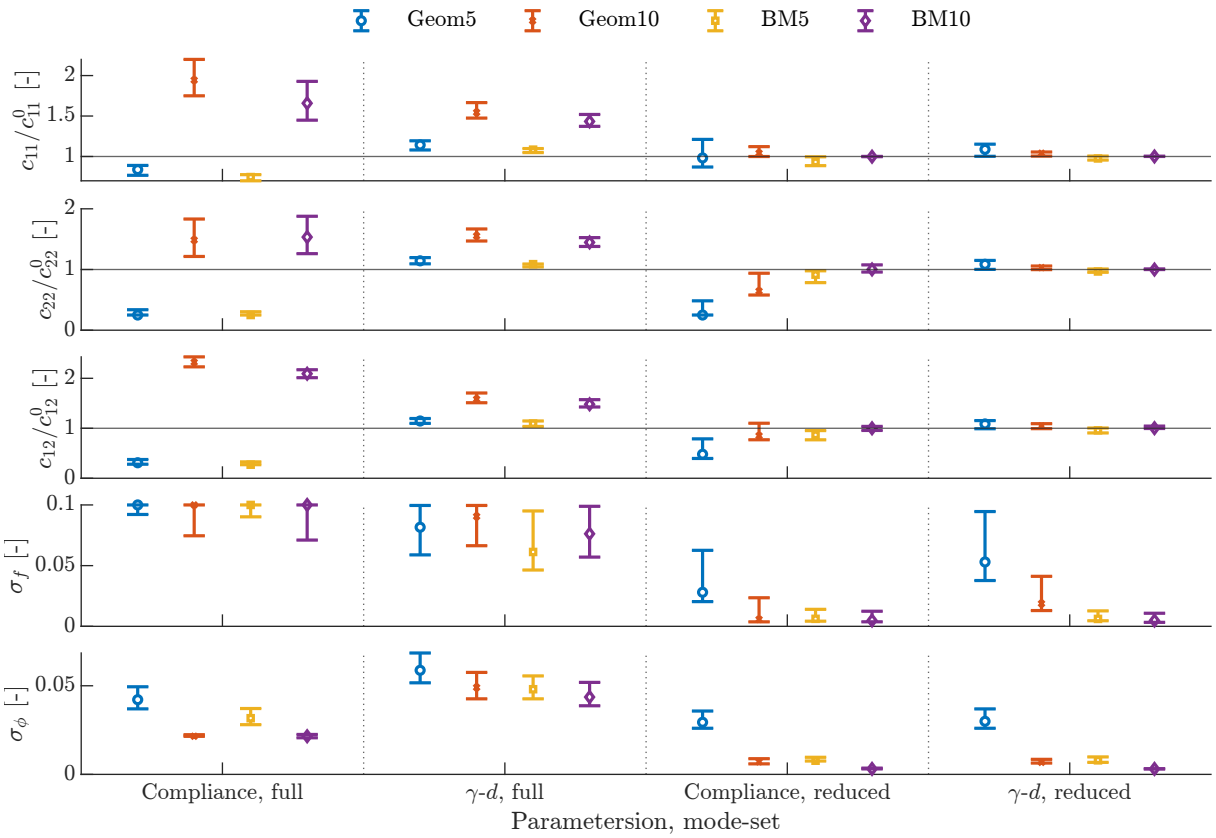
Figure 4.7 shows that the choice of mode set has a significant influence on the accuracy the posterior distributions and the robustness of the model updating framework. Results from the full mode set exhibit substantial variation and, in most cases, unacceptable errors in the identified foundation parameters. In contrast, the reduced mode set shows considerably improved results, highlighting the importance of careful mode set selection. In these examples, the improvement is due to the exclusions of the second and third pairs (modes SS2, FA2, SS3, FA3), which are the most sensitive to variations in blade properties and hence associated with the largest model errors (see Table 4.3). Therefore, these modes violate the zero-mean assumption used to define the likelihood function (Equation 3.11), and cause the estimated foundation parameters to compensate for model error when these modes are included. If the model error is such that this compensation leads to a good match between the model and measured modal parameters, then the framework results in overconfident but inaccurate posterior distributions. For example, the error bars in Figure 4.7b show that the marginal distributions fail to capture the true foundation parameters within a 95 % confidence interval for all cases using the full mode set.

Despite this, the inclusion of the model error parameters (σ_f and σ_ϕ) in the Bayesian framework provides valuable information to interpret these results. Figure 4.7a shows that the posterior distributions obtained from the full mode set resulted in significantly larger MAP estimates for these parameters—particularly σ_f, which reaches its upper bound of 0.1 across all models under the compliance parameterisation (see 'Compliance, full' result for σ_f in Figure 4.7a). This indicates a poor fit between the model and the measured data, signalling to the user that the inferred foundation parameters may not be reliable due to substantial model error.

On the other hand, when both σ_f and σ_ϕ are relatively low—as is the case for the results obtained with the reduced mode set—there is a general trend towards smaller errors in the identified compliance parameters. This demonstrates an example where lower estimates of σ_f



(a) MAP estimates for each parameter



(b) Error bars indicating the 95 % confidence interval for the marginal distributions of each parameter

Figure 4.7: Posterior distributions obtained when using the models defined by the scaling methods and reference turbines listed in Table 4.2 and mode sets defined in Table 4.4.

and σ_ϕ are associated with more accurate and trustworthy posterior results. However, for the results from the full mode set, this relationship becomes less obvious because σ_f may increase while σ_ϕ decreases across different models, making it more difficult to use these parameters as indicators for selecting an optimal model. For example, in the 'γ-d, full' results shown in Figure 4.7a, the *Geom5* model has lower compliance errors than the *Geom10* model, yet the former has a lower σ_f MAP estimate while the latter has a lower σ_ϕ MAP estimate. This complex interaction is likely a consequence of the significant model error introduced by the second and third mode pairs. Therefore, the effectiveness of σ_f and σ_ϕ as indicators of accuracy in the identified foundation parameters depends on the extent of blade-induced error in the selected mode set.

Figure 4.7 shows that the posterior distributions also depend on the choice of foundation parameterisation. The compliance parameterisation, defined with uniform priors for c_{11}, c_{22}, and c_{12}, allows a wide range of equally likely parameter combinations. While this was not problematic in Chapter 3, where the modal parameter bias from measurement error was minimal, the results obtained using the compliance parameterisation here show large variability—and in some case, potentially unphysical foundation estimates—due to increased model error.

In contrast, the γ-$\boldsymbol{d}$ parameterisation constrains the allowed combinations of compliance terms (see Figure 4.5). Consequently, the results using this parameterisation show reduced variability across blade models and more accurate foundation estimates. Although the γ-$\boldsymbol{d}$ parameterisation still shows some sensitivity to the choice of blade model when using the full mode set, the combination of the γ-$\boldsymbol{d}$ with the reduced mode set consistently resulted in acceptable compliance error—comparable to that of Chapter 3.

This suggests that the variability and magnitude of the model-error-induced bias in the identified parameters can be mitigated by increasing the informativeness of the prior. In other words, when model errors are suspected to be significant, their influence on the posterior can be controlled by adjusting the balance between the likelihood and prior.

(ii) Influence of Blade Scaling Method and Reference Wind Turbine

Figure 4.7 also highlights that the posterior distributions can be sensitive to the choice of blade scaling method and the reference turbine. Apart from the 'γ-d, reduced' results, the MAP estimates in the compliance parameters vary significantly depending on which up-scaled model is used. For example, in the 'Compliance, full' results, the MAP estimates for c_{11}, c_{22} and c_{12} using the *BM10* model are 2.22, 6.14, and 7.71 times those obtained from the *BM5* model, respectively. Moreover, when updating with the full mode set, there is no clear 'best choice' of scaling strategy to minimise the impact of blade model error on the identified foundation parameters. This is because the model errors introduced by the second and third pairs (SS2, FA2, SS3, FA3) are too large and dominate the results. As shown in Section 4.3.3, the errors in these modes depend on the specific characteristics of each blade model, with no consistent trend across reference turbines or scaling methods. Consequently, this lack of trend carries over into the errors in the identified foundation parameters, making it difficult to recommend an optimal blade modelling approach when modes that are highly sensitive to blade properties are used in the model updating framework.

However, when these biased modes are excluded using the reduced mode set, a clearer trend in the influence of blade scaling becomes apparent. In this case, the blade mass scaling approach leads to more accurate foundation estimates than the geometric scaling approach. This observations was expected, since blade mass scaling ensures the total mass is preserved, which is more likely to result in similar dynamic behaviour. Additionally, better performance is observed when scaling from the 10 MW to the 15 MW, compared with scaling from the 5 MW. Again, this aligns with expectations, as the 10 MW and 15 MW reference turbines are closer in size, meaning that key design decisions that affect the blade properties—such as geometric and material distributions—are more likely to be comparable. As a result, the modal behaviour is better preserved during the scaling process, reducing model error and improving the accuracy of the identified foundation parameters.

(iii) Model Evidence vs Average Compliance Error

As previously discussed, the model error parameters σ_f and σ_ϕ provide valuable insight for assessing the reliability of the foundation parameters. In addition to these parameters, certain sampling algorithms—such as TMCMC [174] or the improved-TMCMC [175]—output estimates of the model evidence. For a given dataset, the model evidence can be used for model selection [82]. Therefore, this metric has the potential to identify the most appropriate framework for identifying the foundation parameters.

To investigate this, Figure 4.8 plots the model evidence against the average compliance MAP error. The average compliance MAP error, $\varepsilon_c^{\text{MAP}}$, is defined as:

$$\varepsilon_c^{\text{MAP}} = 100 \times \frac{1}{3}\left(\left|\frac{c_{11}^{\text{MAP}} - c_{11}^0}{c_{11}^0}\right| + \left|\frac{c_{22}^{\text{MAP}} - c_{22}^0}{c_{22}^0}\right| + \left|\frac{c_{12}^{\text{MAP}} - c_{12}^0}{c_{12}^0}\right|\right). \tag{4.6}$$

Figure 4.8 groups results from all four blade models depending on the mode set and foundation parameterisation used. Each marker for a given colour in these plots corresponds to a different model, and for each group, the linear line of best fit is shown, along with its slope and coefficient of determination R^2.

Considering first the results from the full mode set with the compliance parameterisation (Figure 4.8a, blue group), a positive trend is observed between the model evidence and foundation error. This suggests that the model evidence is not a reliable metric for model selection in this case. One possible explanation for this is that the compliance parameterisation, which imposes fewer restrictions on the parameters due to the assumed priors, allows the framework to identify a parameter combination that achieves a good fit between the model and the measured data by compensating for model errors. As a result, the model evidence can increase despite large errors in the foundation parameters.

For the γ-$\boldsymbol{d}$ parameterisation with the full mode set (Figure 4.8b, blue group), the extent to which the compliance parameters can compensate for model error is reduced. However, the model error is still significant enough to dominate the results, and consequently, no clear relationship is observed between model evidence and foundation error in this case.

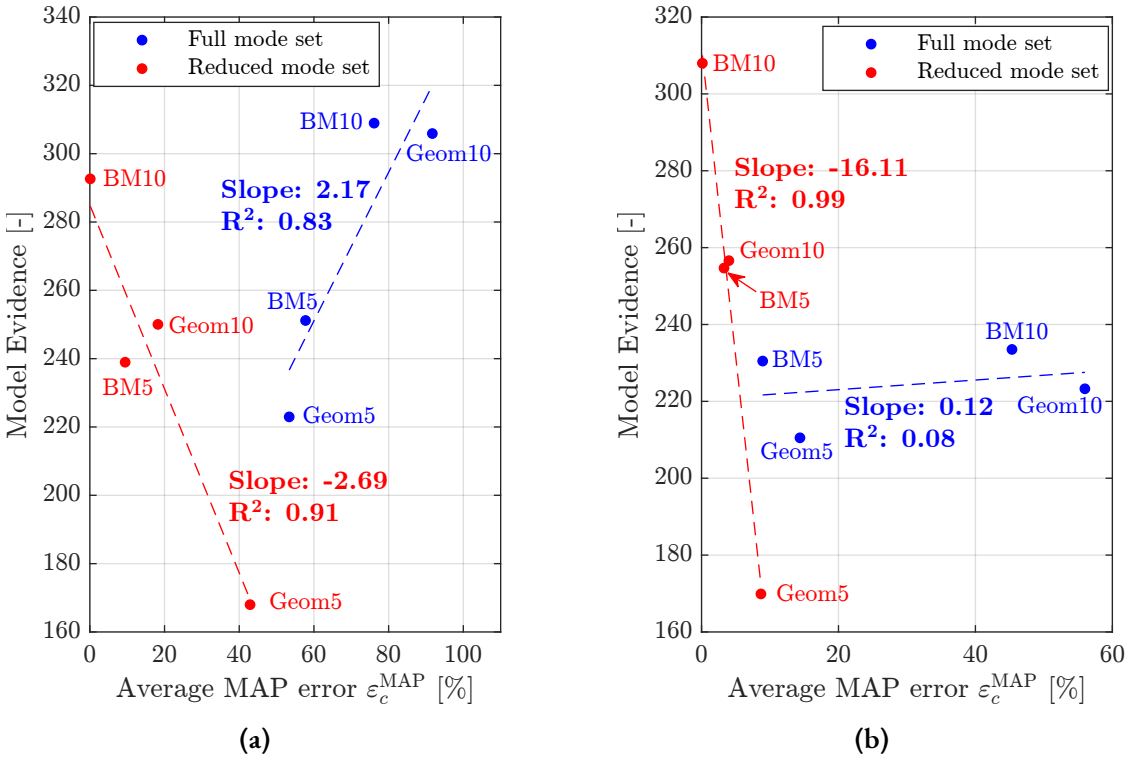


Figure 4.8: Trend between model evidence and average MAP estimate error in the compliance parameters. The dashed lines show the line of best fit for each mode set. (a) Compliance parameterisation (b) γ-$\boldsymbol{d}$ parameterisation.

When updating with the reduced mode set (red groups in Figure 4.8), both parameterisations shown a strong negative correlation between model evidence and foundation error. This indicates that, in the absence of highly biased modes, higher model evidence corresponds to improved accuracy in the identified foundation parameters. Therefore, this metric provides similar information to the model error parameters (σ_f and σ_ϕ) and can be used in combination with them to interpret the results.

It is important to note that the model evidence tends to be higher when more modes—and therefore more terms in the likelihood function—are included. This explains why the full mode set generally produces higher model evidence values than the reduced mode set, particularly for the compliance parameterisation in Figure 4.8a

Given the small number of data points per group (four), the trends shown in Figure 4.8 should be interpreted with caution. Nonetheless, these results provide valuable insight into the potential of the model evidence as a metric for model selection in the context of foundation parameter identification in the presence of blade modelling errors.

4.5 Summary

This chapter has investigated the impact of blade modelling error on the identifiability of the foundation parameters. Two scaling methods were considered: a geometric scaling approach and a blade mass scaling approach. Blade properties from two reference turbines—the NREL 5 MW and the IEA 10 MW— were upscaled to the 15 MW baseline structure.

The four resulting blade-scaled models were found to introduce varying degrees of model error in the predicted modal parameters. The second and third pairs of modes (SS2, FA2, FA3, SS3) were particularly sensitive to the blade modelling error. This sensitivity was attributed to their dynamic behaviour being strongly influence by both the tower and the first blade modes. In contrast, the first mode pair and higher-order modes were sufficiently separated from the dominant blade dynamics, and therefore exhibited reduced sensitivity to the scaling methods.

Results from using each scaled model to identify the foundation parameters demonstrated that the posteriors distributions can be highly sensitive to assumptions made about the blade model. The scaling method and reference turbine were both shown to influence the estimated foundation parameters. In some cases, overconfident posteriors failed to capture the true compliance values due to compensation for modelling errors. Therefore, in the absence of additional information—beyond blade radius and mass—the results should be interpreted with caution. Nevertheless, the Bayesian framework provides an estimate of the model error through the updating of σ_f and σ_ϕ, offering insight into the accuracy of the model and the reliability of the results.

When the degree of model error in the Bayesian framework was reduced—by discarding the second and third pairs of modes—the results were found to be more consistent, and the choice of scaling method had less influence on the identified foundation parameters. Furthermore, with this reduced mode set, the model evidence could be used to select the most appropriate scaling strategy. In this case, the blade mass scaling approach using the IEA 10 MW reference turbine performed best, in line with expectations.

A new parameterisation of the macro-element foundation was introduced, defined by three parameters: γ, which scales an initial *a priori* compliance matrix, and two additional parameters,

d_1 and d_2, which control how much the compliance terms can explore the space in directions orthogonal to γ. By assigning Gaussian priors to d_1 and d_2, the prior distribution becomes more informative than the uniform distributions used in the original compliance parameterisation. This new approach led to improved accuracy in the identified foundation parameters, but relied on a good estimate of the initial relative spring values c_{22}^0/c_{11}^0 and c_{12}^0/c_{11}^0. This improvement was attributed to the reduced influence of the likelihood—and therefore of the blade model error—on the posterior, as a result of the increased weight of the prior.

Overall, this chapter has demonstrated that accurate blade modelling is critical for reliable identification of the foundation parameters. While the effects of blade modelling error can be mitigated through careful selection of the mode set, scaling strategy, and parameterisation, uncertainty quantification via a Bayesian model updating framework is essential for assessing the confidence that can be placed in the results.

CHAPTER 5

A Substructuring Approach for Input-State Estimation

5.1 Introduction

As discussed in Chapter 1, the Augmented Kalman Filter (AKF) has been successfully applied for input-state estimation of offshore wind turbines in previous studies. However, the accuracy of the estimated results depend on the chosen model, for which various complexities exist, making this a challenging task. For example, Chapter 4 highlighted the challenge of accurately modelling the rotor nacelle assembly, and its impact on the dynamic behaviour of offshore wind turbines.

Tatsis *et al.* [131] presented a solution to this challenge by developing a general substructuring framework for input-state estimation. Their work includes a numerical case study using an onshore wind turbine, where a 'cut' is applied at the the top of the tower and the AKF is used to estimate the tower substructure's response. The advantage of this approach is that the rotor nacelle assembly does not need to be modelled, thus requiring no design details from the manufacturer. However, it does introduce the task of estimating the interface inputs instead of directly modelling the rotor nacelle assembly.

For offshore wind turbines, a second complexity stems from the uncertainty surrounding the soil-structure interaction. While the Bayesian model updating frameworks developed in Chapter 3 can be used to reduce this uncertainty, Chapter 4 showed that the results can be highly sensitive to blade modelling assumptions. Moreover, the linear-elastic spring assumption for the macro-element foundation may not hold under high wind conditions or operating rotor speeds.

Therefore, the primary focus of this chapter is to avoid these limitations by considering a suitable substructure which eliminates the need to model both the rotor nacelle assembly and the soil-structure interaction. To obtain an appropriate model, the structure is 'cut' at two locations: one at the top of the tower and another at the ground level, as illustrated in Figure 5.1. For situations where wave loading is present, the unmeasured wave excitations on the tower are also considered as inputs to be identified. The resulting substructure comprises the non-embedded portion of the monopile and the tower. However, for simplicity, it will be referred to as the 'tower substructure' in this work. Since the considered substructure requires fewer modelling assumptions, it is expected to be associated with significantly reduced model uncertainty, thus offering an advantage over modelling the entire wind turbine system.

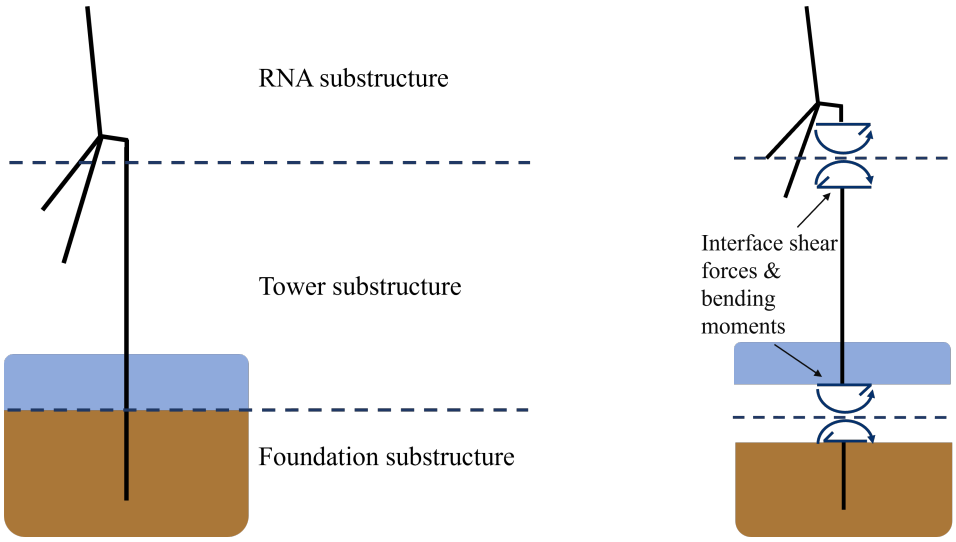


Figure 5.1: Overview of substructuring method for an offshore wind turbine. The structure is 'cut' at the dashed lines, which introduces three substructures: the Rotor Nacelle Assembly (RNA), the tower, and the foundation. This figure provides a simplified 2D representation, whereas the model used in this study is 3D.

To the best of our knowledge, this approach has not been previously attempted. However, it is worth noting that while virtual sensing was successfully applied to the monopile-tower substructure of a 3.3 MW offshore wind turbine by Henkel *et al.* [179], this method requires either a numerical model for the soil-structure interaction or measurements at subsoil locations—both of which are assumed to be unavailable in this chapter.

The chapter is structured as follows. First, in Section 5.2, the mathematical model for a substructure is presented, followed by an overview of the AKF algorithm in Section 5.3.

Section 5.4 presents two numerical case studies involving a 15 MW offshore wind turbine with a monopile foundation to evaluate the performance of the proposed approach. The Simulink model was used to generate realistic synthetic data under operational conditions, as described in Section 5.4.1. However, a separate finite-element model, introduced in Section 5.4.2, was used to represent the substructure for the inverse problem. The first case study considers wind loading only (Section 5.4.4), whereas the second considers combined wind and wave loading (Section 5.4.5). Input-state estimation results are presented for both case studies, with examples demonstrating how the estimates can be subsequently applied for fatigue assessment, identification of the foundation parameters, and estimation of wind-wave misalignment.

5.2 Mathematical Modelling of Substructures

Offshore wind turbines are complex dynamical systems, with inputs arising from aerodynamics, hydrodynamics, and the controller that govern their response. Despite this complexity, the system can be divided into various substructures to alleviate specific challenges. For example, Figure 5.1 shows an offshore wind turbine separated into three substructures: the RNA substructure, the tower substructure and the foundation substructure. By focussing only on the tower substructure, beam-type finite elements can be used, which require no detailed knowledge of the rest of the system. However, the coupled interaction with the adjacent substructures must be taken into account in the analysis to accurately capture the overall system's behaviour. This section outlines the mathematical framework for developing a reduced-order model for the tower substructure. The presented theory can be applied generally to other systems when substructuring is of interest. Furthermore, since this substructure is unconstrained, rigid body modes are present, and a method for handling these modes is also described in this section.

The second order equation of motion for a linear time invariant system with two substructures, denoted with superscripts α and β, can be represented in the following form:

$$\begin{bmatrix} \mathbf{M}^{\alpha} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^{\beta} \end{bmatrix} \begin{bmatrix} \ddot{\boldsymbol{u}}^{\alpha}(t) \\ \ddot{\boldsymbol{u}}^{\beta}(t) \end{bmatrix} + \begin{bmatrix} \mathbf{C}^{\alpha} & \mathbf{0} \\ \mathbf{0} & \mathbf{C}^{\beta} \end{bmatrix} \begin{bmatrix} \dot{\boldsymbol{u}}^{\alpha}(t) \\ \dot{\boldsymbol{u}}^{\beta}(t) \end{bmatrix} + \begin{bmatrix} \mathbf{K}^{\alpha} & \mathbf{0} \\ \mathbf{0} & \mathbf{K}^{\beta} \end{bmatrix} \begin{bmatrix} \boldsymbol{u}^{\alpha}(t) \\ \boldsymbol{u}^{\beta}(t) \end{bmatrix} = \begin{bmatrix} \mathbf{S}_p^{\alpha} & \mathbf{S}_p^{\beta} \end{bmatrix} \begin{bmatrix} \boldsymbol{p}^{\alpha}(t) \\ \boldsymbol{p}^{\beta}(t) \end{bmatrix}, \tag{5.1}$$

where $\mathbf{M}$, $\mathbf{C}$, and $\mathbf{K}$ are the mass, damping and stiffness matrices respectively and $\boldsymbol{u}(t)$ is the

displacement vector at time t. The inputs applied to the system are denoted by $\boldsymbol{p}(t)$ and $\mathbf{S}_p$ is a selection matrix, which maps the input signal(s) to the degrees of freedom (DOFs) at which they act. Equation 5.1 represents the uncoupled equation of motion, where each substructure contains the DOFs relevant to that substructure, including the shared DOFs at the interface between adjacent substructures. Note that in the form presented here, $\boldsymbol{p}(t)$ contains both the externally applied loads and the internal forces. The interface forces are required so that compatibility between the two components is ensured [180].

To relate the notation within the context of this chapter, α may denote the primary substructure of interest (e.g. the tower substructure, as shown in Figure 5.1), while β represents the substructure(s) that are to be disregarded (e.g. the RNA and foundation substructures, as shown in Figure 5.1). Since only substructure α is being modelled, the superscripts α and β in Equation 5.1 will be omitted going forward to simplify the notation.

5.2.1 Reduced-Order Model

Equation 5.1 presents the equation of motion in physical coordinates. In this form, the model can be associated with a large number of system states, making input-state estimation a challenging task. However, in practice, the system's response is often dominated by only a limited number of modes. Therefore, it is often advantageous to apply model reduction techniques to approximate the original full-order model. While several methods exist for model reduction [181], this work uses a free-interface approach from the class of component mode synthesis techniques [154, 155], which allows for the reduction to be applied at the substructure level.

To apply this reduction, a transformation is introduced to convert Equation 5.1 from physical coordinates to modal coordinates:

$$\boldsymbol{u}(t) = \boldsymbol{\Phi}\boldsymbol{z}(t), \tag{5.2}$$

where $\boldsymbol{z}(t) \in \mathbb{R}^{n_m \times 1}$ is the displacement vector in modal coordinates and $\boldsymbol{\Phi} \in \mathbb{R}^{n_{dof} \times n_m}$ is the reduction basis. n_m and n_{dof} are the number of modes in the reduction basis and the number of degrees of freedom for the considered substructure, respectively.

After substituting Equation 5.2 into Equation 5.1, the following equation of motion in modal coordinates can be obtained:

$$\tilde{\mathbf{M}}\ddot{\boldsymbol{z}}(t) + \tilde{\mathbf{C}}\dot{\boldsymbol{z}}(t) + \tilde{\mathbf{K}}\boldsymbol{z}(t) = \boldsymbol{\Phi}^T\mathbf{S}_p\boldsymbol{p}(t), \tag{5.3}$$

where the reduced-order system matrices are computed as: $\tilde{\mathbf{M}} = \boldsymbol{\Phi}^T\mathbf{M}\boldsymbol{\Phi}$, $\tilde{\mathbf{C}} = \boldsymbol{\Phi}^T\mathbf{C}\boldsymbol{\Phi}$ and $\tilde{\mathbf{K}} = \boldsymbol{\Phi}^T\mathbf{K}\boldsymbol{\Phi}$.

(i) Reduction basis

For the unconstrained substructure considered in this chapter, the reduction basis $\boldsymbol{\Phi}$ consists of the retained free-interface normal modes, $\boldsymbol{\phi}_n \in \mathbb{R}^{n_{dof}\times n_r}$, the attachment modes, $\boldsymbol{\phi}_a \in \mathbb{R}^{n_{dof}\times n_a}$, and rigid body modes, $\boldsymbol{\phi}_{rb} \in \mathbb{R}^{n_{dof}\times n_{rb}}$, where n_r, n_a, and n_{rb} are the number of retained normal modes, attachment modes and rigid body modes, respectively. Consequently, the size of the reduction basis can be calculated as $n_m = n_r + n_a + n_{rb}$ and:

$$\boldsymbol{\Phi} = \begin{bmatrix} \boldsymbol{\phi}_{rb} & \boldsymbol{\phi}_n & \boldsymbol{\phi}_a \end{bmatrix}. \tag{5.4}$$

The retained free-interface normal modes, $\boldsymbol{\phi}_n$, attachment modes, $\boldsymbol{\phi}_a$, and rigid body modes, $\boldsymbol{\phi}_{rb}$ are defined as follows.

The rigid body modes, $\boldsymbol{\phi}_{rb}$, represent the substructure's motion without deformation [154], and their number correspond to the unconstrained DOFs of the substructure. In free-interface methods, the normal modes are obtained by placing no constraints at the interface DOFs. These modes capture the internal dynamics of the substructure, and by retaining only those that correspond to the dominant frequency range, the model size can be reduced without compromising accuracy [181]. To compute $\boldsymbol{\phi}_{rb}$ and $\boldsymbol{\phi}_n$, a free-free model of the substructure is first assembled, followed by an eigenvalue analysis:

$$\mathbf{K}\boldsymbol{\Psi} = \mathbf{M}\boldsymbol{\Psi}\boldsymbol{\Sigma}, \tag{5.5}$$

where $\boldsymbol{\Psi}$ is a matrix whose columns correspond to the eigenvectors and $\boldsymbol{\Sigma}$ is a diagonal matrix

containing the associated eigenvalues. The matrix $\mathbf{\Psi}$ includes both the rigid body modes and all free-interface normal modes. Therefore, $\boldsymbol{\phi}_{rb}$ is defined as the eigenvectors associated with zero eigenvalues and $\boldsymbol{\phi}_n$ comprises the retained subset of remaining eigenvectors.

The attachment modes, $\boldsymbol{\phi}_a$, are defined as the static displacement response to an imposed unit force at an interface DOF, with zero force applied elsewhere. These are obtained by solving the static equilibrium equation:

$$\mathbf{K}\boldsymbol{\phi}_a = \begin{bmatrix} \mathbf{0} \\ \mathbf{I}_s \end{bmatrix}, \tag{5.6}$$

where $\mathbf{I}_s \in \mathbb{R}^{n_a \times n_a}$ is an identity matrix corresponding to the interface DOFs, and the zero matrix is of size $(n_{dof} - n_a) \times n_a$ to ensure compatibility. These modes complement $\boldsymbol{\phi}_{rb}$ and $\boldsymbol{\phi}_n$ by capturing the static displacement induced by unit interface forces between coupled substructures [155, 182]. By including attachment modes in the reduction basis, the reduced-order model preserves the effect of interface forces, ensuring an accurate representation of the dynamic coupling between adjacent substructures.

To obtain $\boldsymbol{\phi}_a$, the stiffness matrix $\mathbf{K}$ needs to be inverted and therefore must be non-singular. However, in the case of unconstrained components, such as the tower substructure considered here, $\mathbf{K}$ will be a singular matrix and cannot be inverted.

To overcome this difficulty, the Residual Inertia-Relief Attachment (RIRA) modes $\boldsymbol{\phi}_{\text{RIRA}}$ can be used instead. A brief overview of their computation is presented next. However, for further information, the interested reader is referred to the review of substructuring methods in [183] or the paper by Vettori *et al.* [184], in which the following methodology is applied to estimate the response of a box assembly with a removable component.

1. Obtain the rigid body modes $\boldsymbol{\phi}_{rb}$ and the retained free-interface normal modes $\boldsymbol{\phi}_n$;
2. Calculate the reduced-order stiffness matrix: $\mathbf{\Lambda} = \boldsymbol{\phi}_n^T \mathbf{K} \boldsymbol{\phi}_n$;
3. Calculate the inertia-relief projector: $\mathbf{P}_r = \mathbf{I} - \mathbf{M}\boldsymbol{\phi}_{rb}\boldsymbol{\phi}_{rb}^T$;
4. Arbitrarily choose n_a DOFs to constrain, so that $\mathbf{K}$ becomes non-singular. Then compute

the Inertia-Relief Attachment modes $\boldsymbol{\phi}_{\text{IRA}}$:

$$\mathbf{K}\boldsymbol{\phi}_{\text{IRA}} = \mathbf{P}_r \begin{bmatrix} \mathbf{0} \\ \mathbf{I}_s \end{bmatrix} \tag{5.7}$$

5. Finally, calculate $\boldsymbol{\phi}_{\text{RIRA}}$:

$$\boldsymbol{\phi}_{\text{RIRA}} = \mathbf{P}_r^T \boldsymbol{\phi}_{\text{IRA}} - \boldsymbol{\phi}_n \boldsymbol{\Lambda}^{-1} \boldsymbol{\phi}_n^T \begin{bmatrix} \mathbf{0} \\ \mathbf{I}_s \end{bmatrix} \tag{5.8}$$

5.2.2 State space model

A state space model can be used to replace the second order differential equation of motion (Equation 5.3) with two first order differential equations, so that the mathematics are compatible with the Augmented Kalman Filter. This results in a system described by two equations:

$$\dot{\boldsymbol{x}}(t) = \mathbf{A}\boldsymbol{x}(t) + \mathbf{B}\boldsymbol{p}(t) \tag{5.9}$$

$$\boldsymbol{y}(t) = \mathbf{G}\boldsymbol{x}(t) + \mathbf{J}\boldsymbol{p}(t). \tag{5.10}$$

Equation 5.9 is the state equation, with state vector $\boldsymbol{x}(t) \in \mathbb{R}^{n_s \times 1}$ and the number of states n_s. $\mathbf{A} \in \mathbb{R}^{n_s \times n_s}$ is known as the state matrix and $\mathbf{B} \in \mathbb{R}^{n_s \times n_p}$ is the input matrix, with n_p being the number of inputs to the system. After defining the state vector as $\boldsymbol{x}(t) = [\boldsymbol{z}(t) \;\; \dot{\boldsymbol{z}}(t)]^T$, Equation 5.3 can be rearranged so that $\mathbf{A}$ and $\mathbf{B}$ take the following form:

$$\mathbf{A} = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\tilde{\mathbf{M}}^{-1}\tilde{\mathbf{K}} & -\tilde{\mathbf{M}}^{-1}\tilde{\mathbf{C}} \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \mathbf{0} \\ \tilde{\mathbf{M}}^{-1}\boldsymbol{\Phi}^T\mathbf{S}_p \end{bmatrix}. \tag{5.11}$$

Equation 5.10 is known as the output equation, where $\boldsymbol{y}(t) \in \mathbb{R}^{n_o \times 1}$ is the output vector

and n_o is the number of measured quantities:

$$\boldsymbol{y}(t) = \begin{bmatrix} \mathbf{S}_d & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_v & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{S}_a \end{bmatrix} \begin{bmatrix} \boldsymbol{u}(t) \\ \dot{\boldsymbol{u}}(t) \\ \ddot{\boldsymbol{u}}(t) \end{bmatrix}. \tag{5.12}$$

$\mathbf{S}_d$, $\mathbf{S}_v$ and $\mathbf{S}_a$ are selection matrices for the displacements, velocities and accelerations respectively, and are used to map the physical degrees of freedom of the system to the measured degrees of freedom. $\mathbf{G} \in \mathbb{R}^{n_o \times n_s}$ is the output matrix and $\mathbf{J} \in \mathbb{R}^{n_o \times n_p}$ is the feedthrough matrix. After substituting Equation 5.12 and rearranging Equation 5.3, it can be shown that $\mathbf{G}$ and $\mathbf{J}$ take the following form:

$$\mathbf{G} = \begin{bmatrix} \mathbf{S}_d\boldsymbol{\Phi} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}_v\boldsymbol{\Phi} \\ -\mathbf{S}_a\boldsymbol{\Phi}\tilde{\mathbf{M}}^{-1}\tilde{\mathbf{K}} & -\mathbf{S}_a\boldsymbol{\Phi}\tilde{\mathbf{M}}^{-1}\tilde{\mathbf{C}} \end{bmatrix}, \quad \mathbf{J} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{S}_a\boldsymbol{\Phi}\tilde{\mathbf{M}}^{-1}\boldsymbol{\Phi}^T\mathbf{S}_p \end{bmatrix}. \tag{5.13}$$

Equations 5.9 and 5.10 represent the continuous time state space model. However, because the measurements are typically only available at discrete times, the continuous system needs to be transformed into its equivalent discrete-time representation:

$$\boldsymbol{x}_k = \mathbf{A}_d\boldsymbol{x}_{k-1} + \mathbf{B}_d\boldsymbol{p}_{k-1} \tag{5.14}$$

$$\boldsymbol{y}_k = \mathbf{G}\boldsymbol{x}_k + \mathbf{J}\boldsymbol{p}_k, \tag{5.15}$$

where $\mathbf{A}_d$ and $\mathbf{B}_d$ are the discrete versions of $\mathbf{A}$ and $\mathbf{B}$, respectively. k denotes the current time step, such that $t = k/f_s$ for a sampling frequency of f_s. If a zero-order hold scheme is used for the discretisation, then $\mathbf{A}_d$ and $\mathbf{B}_d$ can be computed as:

$$\mathbf{A}_d = \exp\left(\mathbf{A}/f_s\right) \quad \text{and} \quad \mathbf{B}_d = \mathbf{A}^{-1}\left[\mathbf{A}_d - \mathbf{I}\right]\mathbf{B}. \tag{5.16}$$

Note that exp in this equation denotes the matrix exponential.

5.3 Augmented Kalman Filter for Input-State Estimation

The Augmented Kalman Filter (AKF) [118] is a recursive Bayesian filter that allows for the joint estimation of inputs and states. In the AKF, the unknown inputs $\boldsymbol{p}_k$ are included in an augmented state vector $\boldsymbol{x}_k^a$:

$$\boldsymbol{x}_k^a = \begin{bmatrix} \boldsymbol{x}_k \\ \boldsymbol{p}_k \end{bmatrix}. \tag{5.17}$$

Having defined Equation 5.17, Equations 5.14–5.15 can be developed to take the following stochastic form:

$$\boldsymbol{x}_k^a = \mathbf{A}_d^a \boldsymbol{x}_{k-1}^a + \boldsymbol{w}_{k-1} \tag{5.18}$$

$$\boldsymbol{y}_k = \mathbf{G}^a \boldsymbol{x}_k^a + \boldsymbol{v}_k, \tag{5.19}$$

where $\boldsymbol{w}_k$ and $\boldsymbol{v}_k$ are independent zero-mean random variables. They represent the process noise and measurement noise, respectively, with normal probability distributions:

$$p(\boldsymbol{w}) \sim \mathcal{N}(0,\, \mathbf{Q}^a) \tag{5.20}$$

$$p(\boldsymbol{v}) \sim \mathcal{N}(0,\, \mathbf{R}), \tag{5.21}$$

where $\mathbf{R} \in \mathbb{R}^{n_o \times n_o}$ is the measurement covariance matrix and $\mathbf{Q}^a \in \mathbb{R}^{(n_s+n_p)\times(n_s+n_p)}$ is the augmented process noise covariance matrix. The augmented state matrix $\mathbf{A}_d^a$ and output matrix $\mathbf{G}^a$ then become:

$$\mathbf{A}_d^a = \begin{bmatrix} \mathbf{A}_d & \mathbf{B}_d \\ \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad \mathbf{G}^a = \begin{bmatrix} \mathbf{G} & \mathbf{J} \end{bmatrix}. \tag{5.22}$$

It should be noted that Equation (5.22) assumes a discrete-time random walk model for the unknown dynamics of the inputs, where:

$$\boldsymbol{p}_k = \boldsymbol{p}_{k-1} + \boldsymbol{w}_{k-1}^{\boldsymbol{p}}, \tag{5.23}$$

with superscript $^{\boldsymbol{p}}$ denoting the components of $\boldsymbol{w}_{k-1}$ which relate to the inputs. While this

random walk assumption was found to be sufficient in this work, it could be replaced by a Gaussian Process Latent Force Model (GPLFM), which allows for more advanced modelling of the input dynamics [142].

$\mathbf{R}$ reflects the uncertainty in the measurements and is often assumed to be a diagonal matrix, where each diagonal element represents the variance in the corresponding sensor. As a starting point, this is often defined based on the known uncertainty provided by the sensor specification sheet.

$\mathbf{Q}^a$ reflects the uncertainty in both the unknown states and inputs. It is usually assumed that the uncertainty in the states and inputs are independent, so that the following form can be adopted:

$$\mathbf{Q}^a = \begin{bmatrix} \mathbf{Q} & \mathbf{0} \\ \mathbf{0} & \mathbf{S} \end{bmatrix}, \tag{5.24}$$

where $\mathbf{Q} \in \mathbb{R}^{n_s \times n_s}$ is the covariance matrix for the unknown states and $\mathbf{S} \in \mathbb{R}^{n_p \times n_p}$ is the covariance matrix for unknown inputs. In contrast with $\mathbf{R}$, there is typically very little information available about $\mathbf{Q}$ and $\mathbf{S}$. As a result, these parameters tend to require tuning when using the AKF for input-state estimation. Although more sophisticated methods (such as those reviewed in Chapter 1) could be applied for this task, a trial and error method was chosen in this work. This decision was made because satisfactory results were achieved in a short time frame, making the more complex tuning methods unnecessary.

With the augmented state space model described as above, the AKF algorithm (see Algorithm 2) can be used to recursively estimate the mean augmented state vector $\hat{\boldsymbol{x}}_k^a$ and its associated covariance matrix $\mathbf{P}_k$. This algorithm iteratively updates $\hat{\boldsymbol{x}}_k^a$ and $\mathbf{P}_k$ in two stages: a time update and a measurement update. During the time update stage, the process model (Equation 5.14) is employed to predict the state vector mean and covariance matrix for the current time step, based on the updated values from the previous time step. These predicted values form the *a priori* estimate of the state vector and covariance matrix, which are then passed to the measurement update stage. The measurement update stage utilises the output equation (Equation 5.15) to compute the *a posteriori* estimate of the state vector and covariance matrix.

In summary, the time update stage provides a prior belief for the unknown states, which are then updated using the measurements, to obtain a posterior belief.

Algorithm 2 AKF

Given: $\mathbf{A}_d^a,\ \mathbf{G}^a,\ \mathbf{Q}^a,\ \mathbf{R}$
Initialise: $\hat{\boldsymbol{x}}_0^a,\ \mathbf{P}_0$
for $k = 1, 2, \ldots$ **do**
(1) Time Update:
$\hat{\boldsymbol{x}}_k^{a-} = \mathbf{A}_d^a \hat{\boldsymbol{x}}_{k-1}^a$ ▷ Predict state vector
$\mathbf{P}_k^- = \mathbf{A}_d^a \mathbf{P}_{k-1} \mathbf{A}_a^{aT} + \mathbf{Q}^a$ ▷ Predict state covariance matrix

(2) Measurement Update:
$\mathbf{L}_k = \mathbf{P}_k^- \mathbf{G}^{aT} (\mathbf{G}^a \mathbf{P}_k^- \mathbf{G}^{aT} + \mathbf{R})^{-1}$ ▷ Compute Kalman gain $\mathbf{L}_k$
$\hat{\boldsymbol{x}}_k^a = \hat{\boldsymbol{x}}_k^{a-} + \mathbf{L}_k (\boldsymbol{y}_k - \mathbf{G}^a \hat{\boldsymbol{x}}_k^{a-})$ ▷ Update state vector
$\mathbf{P}_k = (\mathbf{I} - \mathbf{L}_k \mathbf{G}^a) \mathbf{P}_k^{-1}$ ▷ Update state covariance matrix
end for

Result: Augmented state mean $\hat{\boldsymbol{x}}^a$ and covariance $\mathbf{P}$

5.4 Performance Evaluation using Numerical Case Studies

In this section, numerical case studies are presented to assess the accuracy of the proposed substructuring approach on a 15 MW offshore wind turbine with a monopile foundation. Results are presented for two case studies; one with wind loading only and a second with both wind and wave loading.

5.4.1 Synthetic data generation

The Simulink Model from Chapter 2 was used to generate synthetic data for the case studies presented here. Similar to Chapters 3 & 4, the model was defined based on NREL's 15 MW reference offshore wind turbine with a monopile foundation [10]. The main features of this model are depicted in Figure 5.2.

The model setup in this chapter is identical to the baseline wind turbine model introduced in Chapter 4, which is based on the model used in Chapter 3, but with a more accurate blade representation. Therefore, the implementation details for the hub, nacelle, tower, monopile, and macro-element foundation—defined by $\mathbf{K}_{\text{macro}}^0$ (Equation 3.19)—can be found in Section 3.4.1. For details on the blade implementation, refer to Section 4.2.

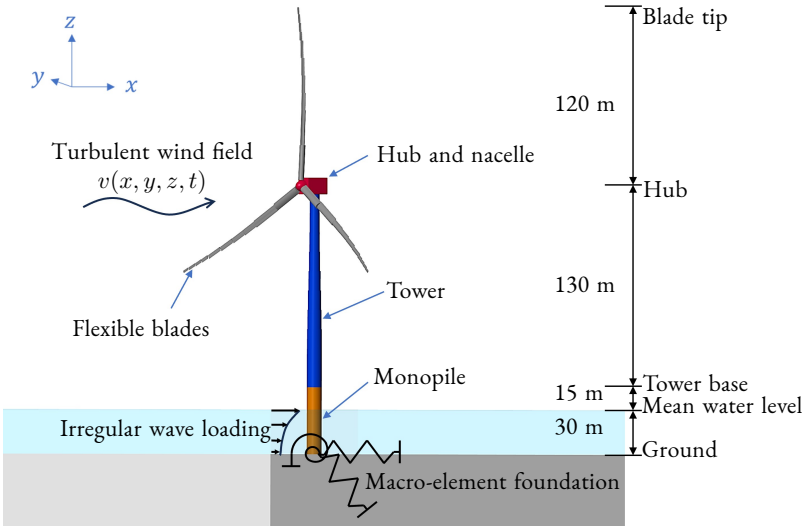


Figure 5.2: Simulink model setup used for generating synthetic data

It is worth emphasising that the soil-structure interaction and the rotornacelle assembly modelling required here is solely for generating data for the numerical case studies. For the reverse problem, a separate model—introduced in Section 5.4.2—is used, which only requires the tower and monopile to be modelled.

To generate data with wind loading, a 3D velocity wind field covering the rotor-plane was generated using Turbsim [162] and passed to the aerodynamics module in the Simulink model (see Section 2.3.4 for implementation details).

For the case study with wind and wave loading, the Morison equation was used to compute the hydrodynamic forces and moments acting on submerged elements. Linear wave theory was used to generate unidirectional irregular wave kinematics, based on a JONSWAP spectrum (see Section 2.3.5 for implementation details).

The hyperparameters for calculating the turbulent wind and irregular wave loading are defined later, when introducing each case study. However, both case studies consider the turbine in an operational state, where the ROSCO controller was actively regulating the generator torque and blade pitch angle.

5.4.2 Linear finite element model

To provide separation between the forward and reverse problems, a different model was used for input-state estimation. For this model, 1D Euler-Bernoulli beam elements were used to assemble finite-element mass and stiffness matrices for the tower substructure, which was discretised into 19 element using the mesh shown in Figure 5.3a. The geometric and material properties are provided in Table 5.1. Note that these values are the same as those used for the tower and monopile substructure in the Simulink model to generate data.

Since this substructure is obtained by applying a cut just below the nacelle and another at the ground level, the user is not required to make assumptions about the RNA or soil-structure interaction. To introduce structural damping into the model, a proportional damping approach was employed [104], with a 0.8% damping ratio for the first mode. The damping matrix was then computed in proportion to the mass and stiffness matrices. It was also assumed that the response in the fore-aft (FA) and side-side (SS) directions were dominant, so that the axial and torsional degrees of freedom could be neglected.

Table 5.1: List of user-defined inputs required to model the tower substructure.

Model Input	Value/reference for 15 MW wind turbine
Diameter profile	Provided in the reference document [10]
Wall thickness profile	Provided in the reference document [10]
Young's modulus	210 GPa
Poission's ratio	0.3
Material density	7850 kg/m^3 for monopile or 8500 kg/m^3 for tower

Once the mass **M**, stiffness **K**, and damping **C** matrices were assembled, the approach outlined in Section 5.2.1 was applied to obtain a reduced-order model. In the case studies considered here, a reduced-order model with four retained normal modes was used, corresponding to the first two bending modes in the fore-aft and side-side directions. This resulted in a reduction basis comprising 16 modes: 4 rigid body modes, 4 retained normal modes, and 8 attachment modes. The rigid body modes and retained free-interface normal modes are illustrated in Figure 5.4. Each of the 8 attachment modes correspond to one of the interface forces to be estimated, which include shear forces and bending moments at both the ground level and

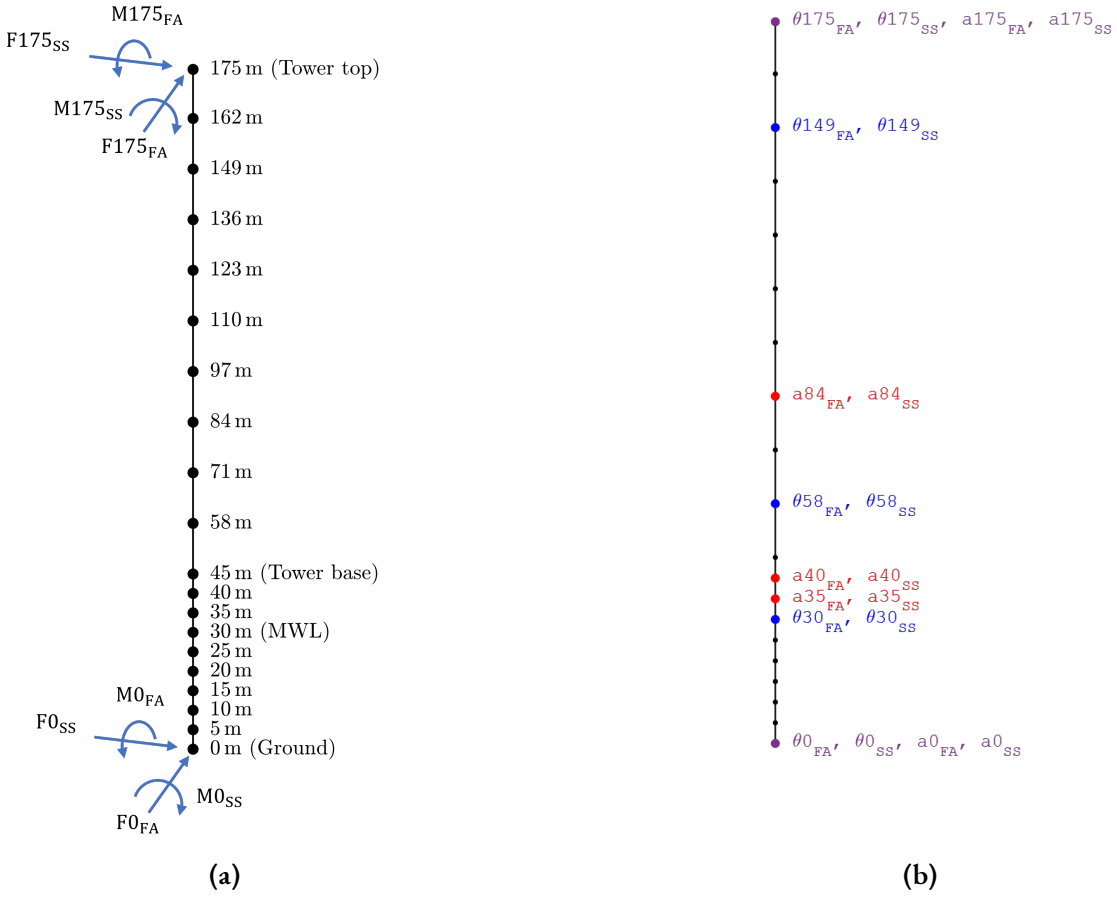


Figure 5.3: (a) Linear finite element model with interface forces and moments used for input-state estimation, (b) 'Sparse: θ-a' sensor setup.

tower top, as shown in Figure 5.3a.

(i) Sensor setup

Three different sensor setups were considered for the numerical case studies and are summarised in Table 5.2. In this chapter, unless stated otherwise, all inputs, states and measurements are denoted using the notation $\alpha\beta_\gamma$, where α refers to the signal type (F: shear force, M: bending moment, a: acceleration, x: displacement, θ: rotation), β refers to the signal's location in metres above ground level and γ denotes the signal's direction (FA: fore-aft, SS: side-side).

Two of the considered setups assume that sensors are installed at every node in the FE model. These sensor configurations are labelled 'Dense: x-a' and 'Dense: θ-a'. They consist of either accelerometers and displacements, or accelerometers and inclinometers, respectively. Displacement sensors are mostly here used as a reference; in reality, displacement information

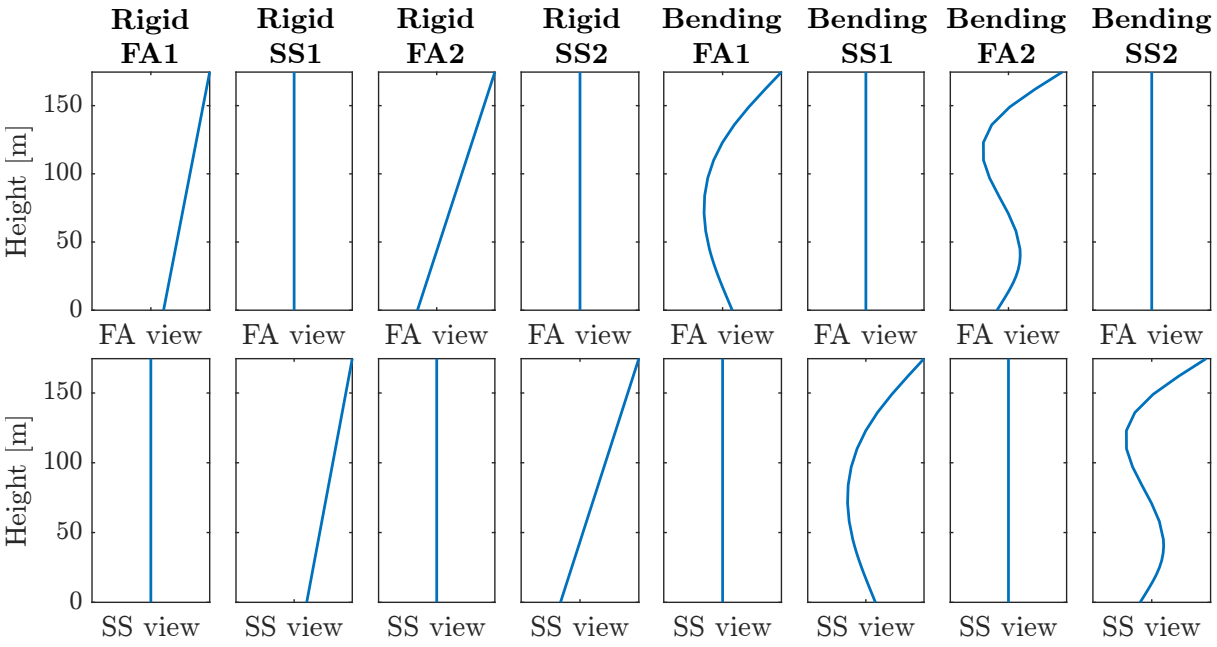


Figure 5.4: Graphical representation of the rigid body modes and retained normal modes for the tower substructure. Each column represents an individual mode, with the top and bottom rows displaying the projections in the fore-aft (FA) and side-side (SS) directions, respectively.

is rare and costly to extract. To the contrary, acceleration measurements, usually delivered by means of MEMs technologies at low cost are the common option. In addition, tilt-meter (inclinometer) measurements can also be made available via such low cost technologies, delivering a quantity that is linked to deformation, and not its derivatives (velocity or acceleration). The latter is highly useful for improving filter estimates in virtual sensing schemes.

Although it is unlikely to be provided with such a heavily instrumented structure in practice, these dense sensor configurations were initially considered so that a comparison could be

Table 5.2: Summary of sensor configurations considered in the numerical case studies. Note that at each sensor location, the response in both the fore-aft and side-side directions are used in the measurement set.

Sensor setup label	Number/Type of sensors	Sensor Locations
Dense: `x-a`	40 displacement sensors at:	every node (see Fig. 5.3a)
	40 accelerometers at:	every node (see Fig. 5.3a)
Dense: θ-`a`	40 inclinometers at:	every node (see Fig. 5.3a)
	40 accelerometers at:	every node (see Fig. 5.3a)
Sparse: θ-`a`	10 inclinometers at:	0, 30, 58, 149, 175 m (see Fig. 5.3b)
	10 accelerometers at:	0, 35, 40, 84, 175 m (see Fig. 5.3b)

made between displacement and inclinometer measurements, without the influence of sensor placement affecting the results.

A more realistic sensor configuration is further considered and is referred to as the 'Sparse: θ-a' setup. For this sensor setup, only inclinometers and accelerometers are considered, since these types of sensors are more likely to be installed in practice. Moreover, it was desired to keep the number of sensors to a minimum. Following the guidance in [185], to achieve stable input estimation, the number of displacement sensors (or inclinometers) has to be greater than or equal to the number of inputs. In the case studies considered here, the maximum number of inputs to be estimated is ten (for Case Study 2). As a result, the sparse setup includes ten inclinometers.

To improve the invertibility conditions for input estimation, accelerometers were also included, with their count limited to ten. Furthermore, due to the higher installation cost associated with submerged sensors, an additional restriction was imposed, which only allowed one submerged location in the sensor configuration.

With the above restrictions defined, the ten accelerometers and ten inclinometers were distributed as shown in Figure 5.3b and Table 5.2. Note that at each sensor location, the response in both the fore-aft and side-side directions are included in the measurement set. To replicate practical noise inherent in real sensors, zero mean white noise with standard deviation equal to 1% of the true signal was added to each measurement.

5.4.3 Comparison metrics

In order to evaluate the performance of the proposed approach, the following metrics were used:

1. The Time Response Assurance Criterion ($TRAC$) for two signals in the time domain, $\sigma_1(t)$ and $\sigma_2(t)$:

$$TRAC = \frac{\left[\sigma_1(t)^T \sigma_2(t)\right]^2}{\left[\sigma_1(t)^T \sigma_1(t)\right]\left[\sigma_2^T(t) \sigma_2(t)\right]} \tag{5.25}$$

2. The Frequency Response Assurance Criterion ($FRAC$) for two signals in the frequency

domain, $\sigma_1(f)$ and $\sigma_2(f)$:

$$FRAC = \frac{\left[\sigma_1(f)^T \sigma_2(f)\right]^2}{\left[\sigma_1(f)^T \sigma_1(f)\right]\left[\sigma_2^T(f)\sigma_2(f)\right]} \tag{5.26}$$

3. The Mean Absolute Error between two signals in the time domain, $\sigma_1(t)$ and $\sigma_2(t)$:

$$MAE = \frac{1}{n_t}\sum_{i}^{n_t} |\sigma_1(t_i) - \sigma_2(t_i)|, \tag{5.27}$$

where n_t is the number of data points and i indicates the time step.

The $TRAC$ value allows for a comparison of two signals in the time domain and can take a value between 0 and 1, where a value of 1 indicates that two signals are perfectly in phase and value of 0 suggests that two signals are out of phase. Similar to the $TRAC$, the $FRAC$ takes a value between 0 and 1 and provides insight into the similarity of two signals in the frequency domain. In this chapter, all three metrics are used in combination to evaluate the performance of the estimated results. However, various scenarios exist where these metrics alone can lead to misleading conclusions. Therefore, these criteria were always complimented by a visual inspection of the results presented.

(i) Calculating accumulated fatigue damage

In addition to the metrics described above, the estimated response was also used to calculate the accumulate fatigue damage. In this chapter, the damage was estimated based on the stress response and the widely used rainflow counting method [186], which determines the number of cycles at corresponding mean stress and stress range levels. Miner's rule was then subsequently used to compute the accumulate fatigue damage, D:

$$D = \sum_{i=1}^{k} \frac{n_i}{N_i}, \tag{5.28}$$

where n_i is the number of cycles in the stress time history at the ith stress range and N_i is the number of cycles to failure for the ith stress range.

The number of cycles to failure, N_i, for a given stress range can be determined using

Basquin's law,

$$\log(N_i) = \log(\overline{a}) - m\log(\Delta\sigma_0), \tag{5.29}$$

where $\Delta\sigma_0$ denotes the stress range for a mean stress of zero, m is the Wöhler exponent, and $\overline{a}$ is a constant that is determined from the intercept of Equation 5.29 on the $\log(N)$ axis. Fatigue failure was assumed to occur due to circumferential butt welds failing in the axial direction. The fatigue design standards for offshore steel structures [187] suggest using $m = 5$, and $\log(\overline{a}) = 13.617$ for this type of weld.

For situations when the stress range doesn't centre around a mean of zero, Goodman's rule is often applied to find the equivalent zero mean stress range, $\Delta\sigma_{0,eq}$:

$$\Delta\sigma_{0,eq} = \Delta\sigma_\sigma / \left[1 - \frac{\sigma_{nz}}{\sigma_{UTS}}\right]. \tag{5.30}$$

Here, $\Delta\sigma_\sigma$ is the stress range for the non zero mean stress σ_{nz}, and σ_{UTS} is the ultimate strength of the material under consideration. In this chapter, $\sigma_{UTS} = 500\,\text{MPa}$, since the substructure is assumed to be composed of S355 steel. As the stresses along the tower are inferred, the corresponding fatigue damage calculation will be shown along the tower and will be compared to the 'actual' fatigue life calculated from the simulations.

5.4.4 Case study 1: wind loading only

To generate synthetic data for this wind-only case study, a 700 second forward simulation was first carried out using the Simulink model. In this simulation, Turbsim was used to generate a turbulent wind field based on a mean wind speed of 8 m/s at hub height and 10 % turbulence intensity. The aerodynamic forces calculated based on this wind field predominantly apply loading in the fore-aft direction. However, due to combined effects of the 3D turbulent wind and the aeroelastic interaction, as well as the dynamics introduced by the controller, the substructure also experiences excitation in the side-side direction. Considering these loading conditions, the substructure's input vector $\boldsymbol{p}$ comprises the 8 interface forces and bending

moments depicted in Figure 5.3a:

$$\boldsymbol{p} = \begin{bmatrix} \mathtt{F0}_{\mathtt{FA}} & \mathtt{M0}_{\mathtt{FA}} & \mathtt{F0}_{\mathtt{SS}} & \mathtt{M0}_{\mathtt{SS}} & \mathtt{F175}_{\mathtt{FA}} & \mathtt{M175}_{\mathtt{FA}} & \mathtt{F175}_{\mathtt{SS}} & \mathtt{M175}_{\mathtt{SS}} \end{bmatrix}^T \tag{5.31}$$

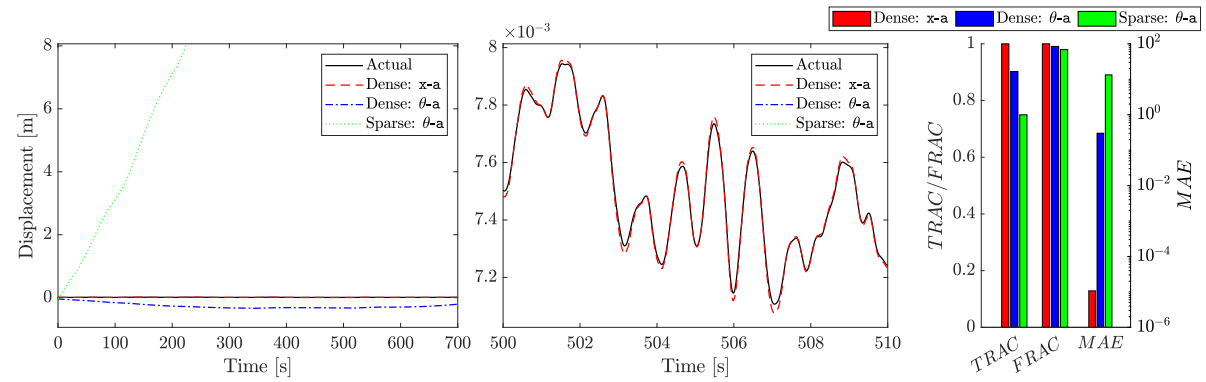
It should be noted that the interface forces and moments at the tower top are due to the combined effects of the external wind pressure, the vibration of the blades, the rotation of the rotor and the actions of the controller.

(i) Case Study 1 results - state estimation

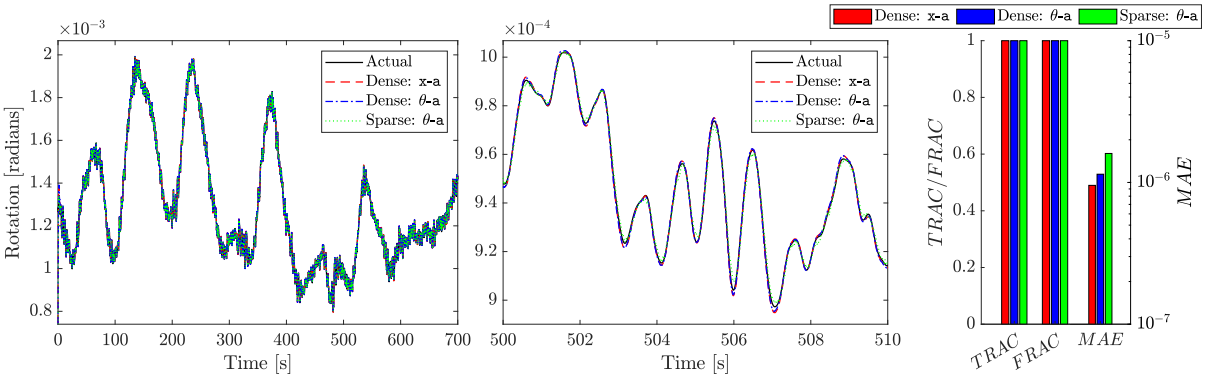
The performance of the proposed approach for state estimation is assessed by comparing the actual response with the AKF-estimated response at four representative DOFs. Due to an interest in gaining further insight into the foundation behaviour, the ground level displacement and rotation were included in the set to be analysed. However, as this node is collocated with sensors, the displacements and rotations were also compared at a submerged location where no sensors were installed (15 m). In summary, the selected DOFs for comparison include the displacements and rotations at the ground level (0 m) and halfway to the mean water level (15 m): $\{\mathtt{x0}_{\mathtt{FA}}, \theta\mathtt{0}_{\mathtt{FA}}, \mathtt{x15}_{\mathtt{FA}}, \theta\mathtt{15}_{\mathtt{FA}}\}$.

Figure 5.5 compares the AKF results with the actual results for all three sensor setups. Focusing first on the results where displacement sensors are used to supplement the accelerometers (labelled 'Dense: x-a'), the figures show that high accuracy has been achieved in the time domain response for the displacements and rotations at both locations. This conclusion is reinforced by the $TRAC$ and $FRAC$ scores being close to 1, indicating that the response has been estimated well in both the time domain and frequency domain.

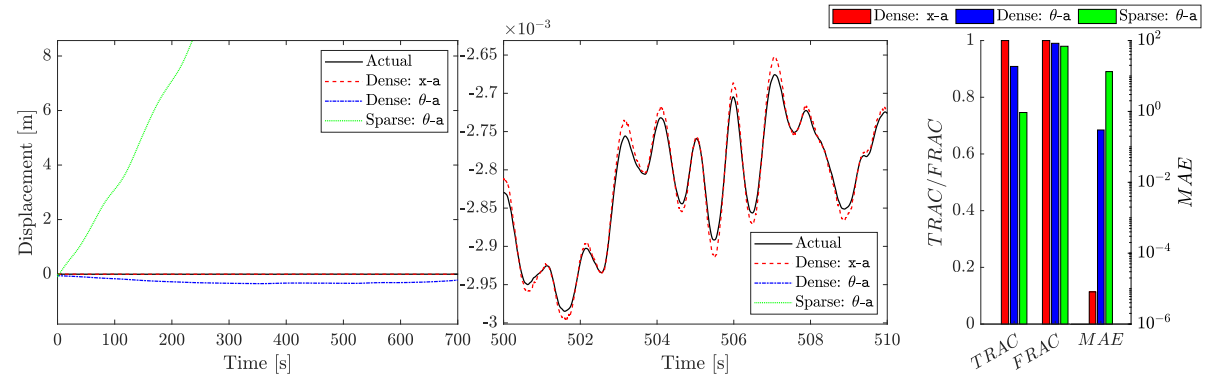
Now considering the results from the setup with all inclinometers and accelerometers (labelled 'Dense: θ-a'), it can be seen that the rotations were identified with similar accuracy to the case with displacement sensors. However, the time domain response for the displacements show that the AKF significantly diverges from the true response, so much so that the true response cannot be seen in the magnified figures. This is an artefact of the unconstrained substructure, since with only measurements from rotations and accelerations, the AKF has no



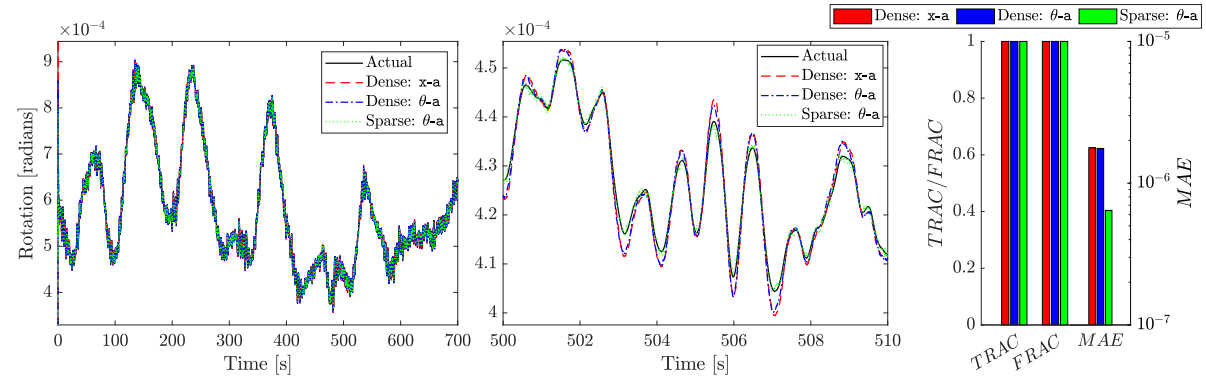
(a) Fore-Aft **displacement** at **15 m** ($\texttt{x15}_{\texttt{FA}}$)



(b) Fore-Aft **rotation** at **15 m** ($\theta\texttt{15}_{\texttt{FA}}$)



(c) Fore-Aft **displacement** at **0 m** ($\texttt{x0}_{\texttt{FA}}$)



(d) Fore-Aft **rotation** at **0 m** ($\theta\texttt{0}_{\texttt{FA}}$)

Figure 5.5: Comparison between actual response and AKF estimated response at selected DOFs for the three sensor setups (Case Study 1). *From left to right:* entire response duration, magnified excerpt, $TRAC/FRAC/MAE$ values.

way to penalise the error in the absolute position of the wind turbine, which leads to the observed drift in the figures. Therefore, if displacement sensors are not used, the displacement response of this unconstrained structure is unobservable. This drift also contributes to a drop in the $TRAC$ values and an increase in the MAEs which can be seen in Figures 5.5a and 5.5c. However, these figures also show that the corresponding $FRAC$ values remain high and similar to the values obtained for the sensor setup with displacements, suggesting that the response was still estimated well in the frequency domain. One possible explanation for this could be the inability to estimate the substructure's rigid body motion. To strengthen this argument, Figure 5.6 compares the estimated response in modal coordinates for the sensor setup 'Dense: θ-a'. Figures 5.6a–5.6d show the estimated and actual rigid body modes, whereas Figures 5.6e–5.6p show the remaining modes (normal and attachment). These figures show that the normal and attachment modes have been estimated with high accuracy and it is only the rigid body modes that exhibit the same drift that is observed in the displacement response from Figures 5.5a and 5.5c. Therefore, it can be concluded that despite the inability to estimate the absolute displacement when inclinometers are used, the contribution of the response from the internal and attachment modes is still captured with high accuracy, comparable to that achieved with displacement sensors.

Fig. 5.5 further indicates that using the sparse sensor setup yields similar performance and observations as the 'Dense: θ-a' sensor setup. It is worth noting that the displacement DOFs exhibit smaller $TRAC$ values and larger MAEs, primarily due to a drift with larger magnitude. However, this does not affect the estimated performance in the frequency domain. Interestingly, based on Figure 5.5d (particularly observing the MAEs), it could be argued that the sparse sensor setup slightly outperforms the dense sensor setups. The minor differences observed between the sparse and dense configurations are due to the different observations available for each setup. However, these results show that with well-placed sensors which ensure observability of the desired states, a reduced number of sensors can still achieve comparable accuracy to the dense configurations, reinforcing the robustness of the method.

Although Figure 5.5 presents only the DOFs in the fore-aft direction, state estimation in the side-side direction showed similar trends, as reflected in the performance metrics in Table

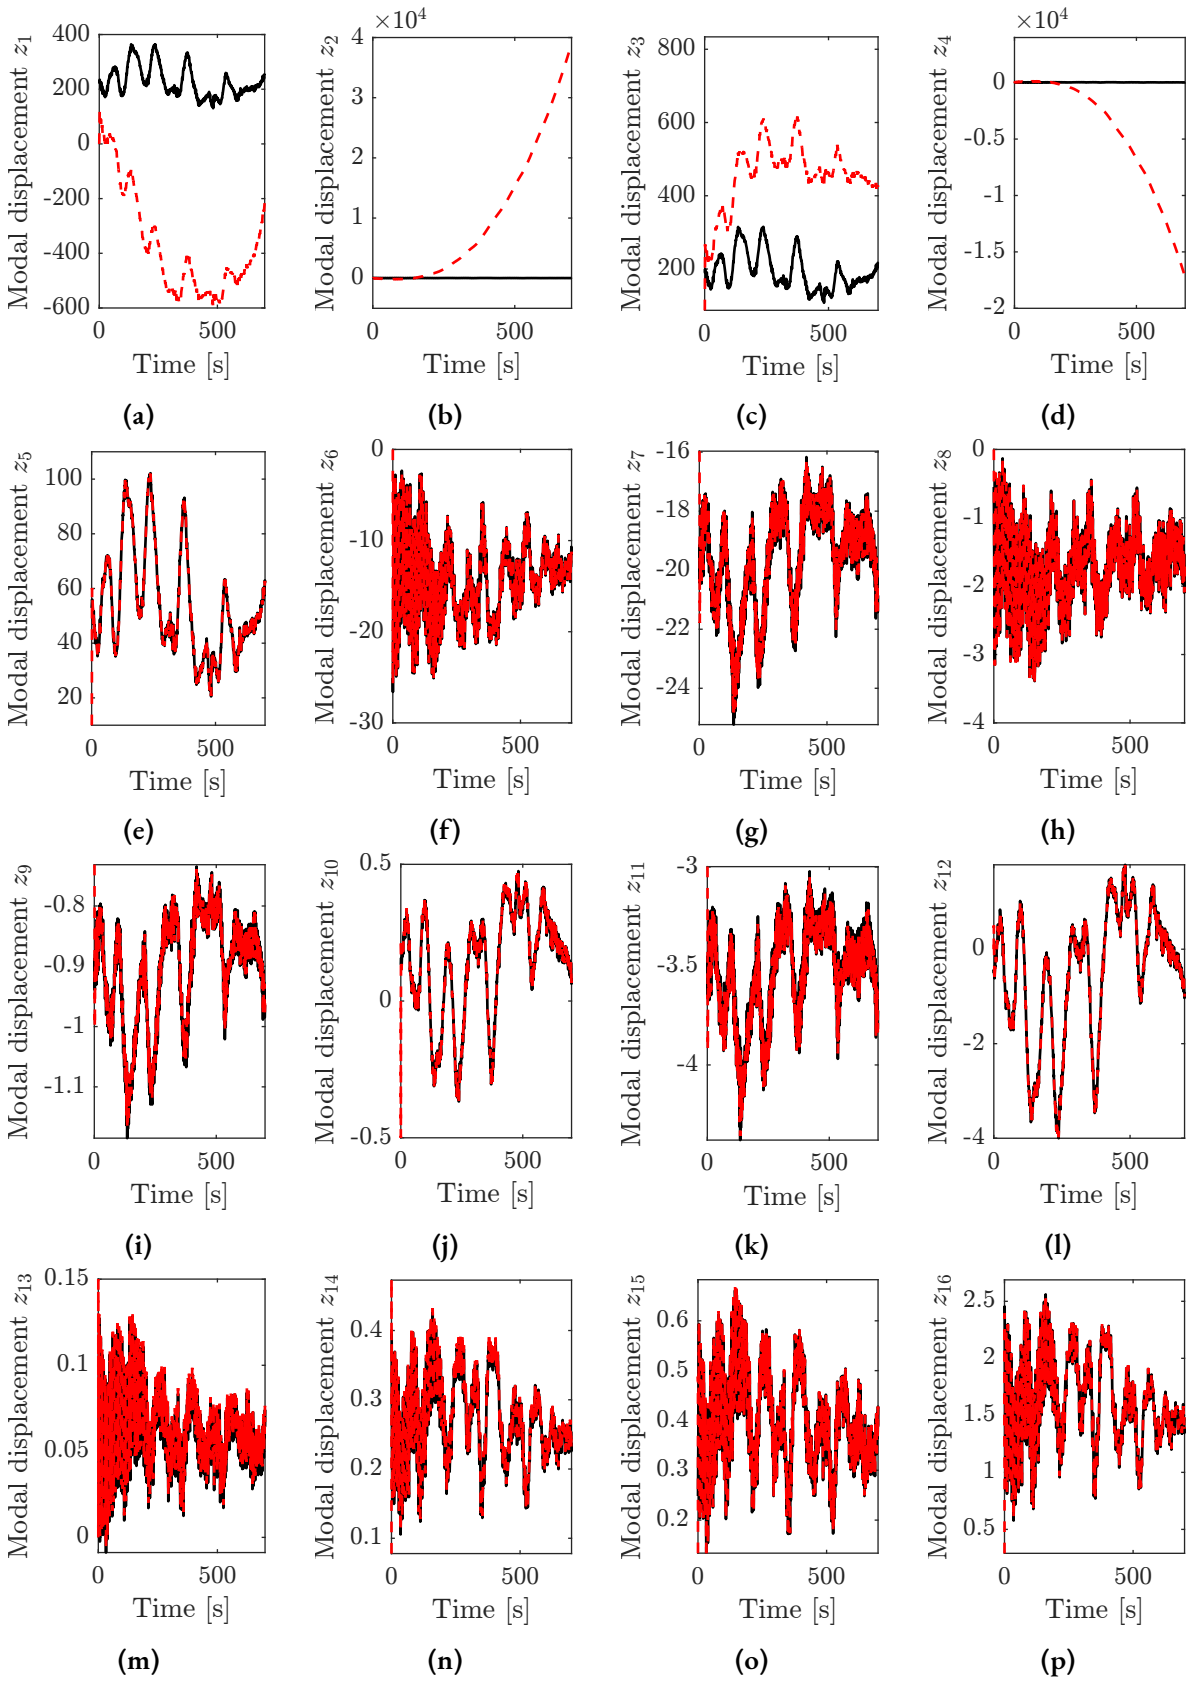


Figure 5.6: Estimated modal displacements when using 'Dense: θ-a' sensor setup (red-dashed lines) vs actual response (solid-black line). Sub-figures (a)-(d): *rigid body modes*, sub-figures (e)-(h): *normal modes*, sub-figures (i)-(p): *attachment modes*.

B.1 in Appendix B. Furthermore, the damage distributions introduced in Section 5.4.4 (ii) provide additional insight into state estimation performance for all nodes of the substructure, in both the fore-aft and side-side directions.

(ii) Case Study 1 results - fatigue damage

Once the displacements and rotations have been estimated, they can be used to calculate the damage for the corresponding node in the FE model, using the approach outlined in Section 5.4.3 (i). This was carried out for every node in the FE model to obtain a distribution of the damage along the height of the substructure. According to Miner's rule (Equation 5.28), failure is assumed to occur when the accumulated damage value reaches 1. Therefore, the damage values presented here represent the contribution over a ten-minute period towards this reference threshold. However, this threshold is an approximation and does not necessarily indicate actual failure, as it does not account for load sequence effects or other nonlinear damage accumulation mechanisms [188]. Nonetheless, this damage model provides a convenient framework for comparing the predicted and actual fatigue estimates.

Minor differences can be observed between the different sensor setups in Figure 5.7. For example, between 120 and 150 metres above ground level, the Sparse: θ-a setup outperforms the dense configurations in the fore-aft direction but shows slightly reduced accuracy in the side-side direction. As discussed in the previous section, these differences mainly arise because different observations result in varying levels of performance. However, since the normal and attachment modes are observable for all three configurations, the different sensor setups perform comparably for fatigue estimation. This is a case where the use of the additional sensors does not result in significant improvement of the estimated quantities. As expected, even if those additional sensors appear to be redundant in the cases studied, their use doesn't practically lead to deteriorating the estimated quantities. The additional information they offer could nonetheless be beneficial when dealing with other cases where there would be greater modelling errors and higher noise levels.

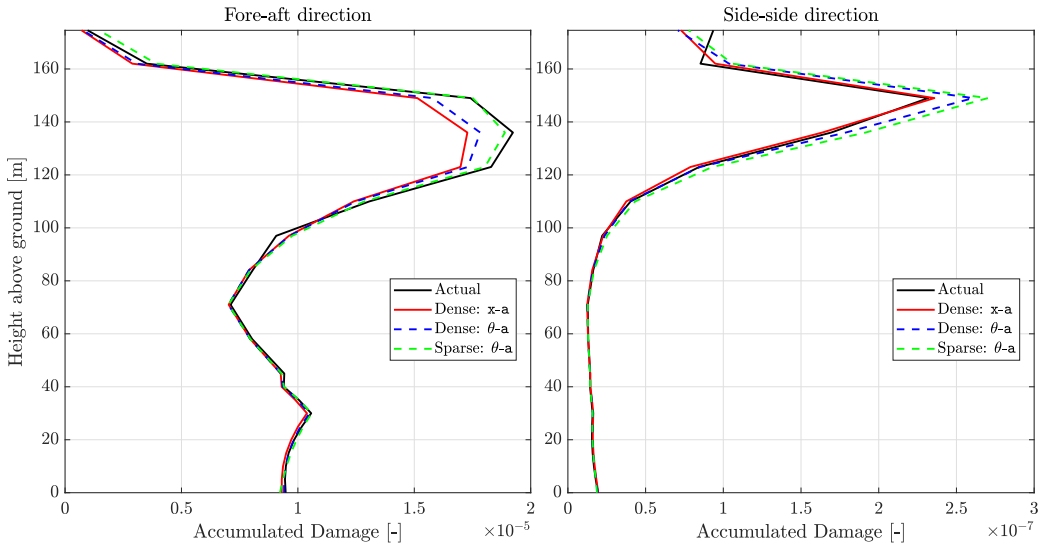
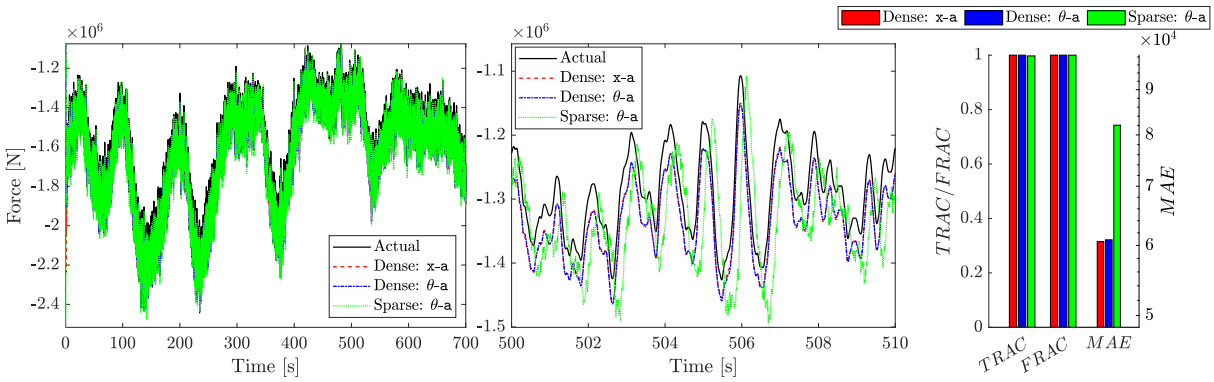


Figure 5.7: Accumulated damage at each node in the finite-element model (Left figure: fore-aft damage distribution. Right figure: side-side damage distribution)

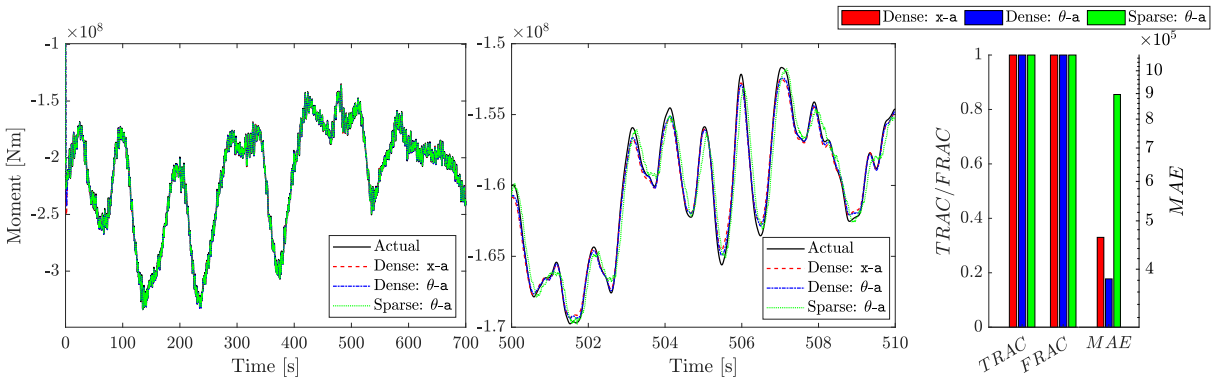
(iii) Case Study 1 results - input estimation

Figures 5.8 and 5.9 compare the estimated inputs at the ground level interface and the tower top interface respectively. Overall, these figures show that the AKF has been able to accurately estimate the main trends in the time domain for all sensor setups, which is further indicated by $TRAC$ values close to 1 for almost all interface inputs. Moreover, the high $FRAC$ values indicate that the response in the frequency domain has also been accurately estimated.

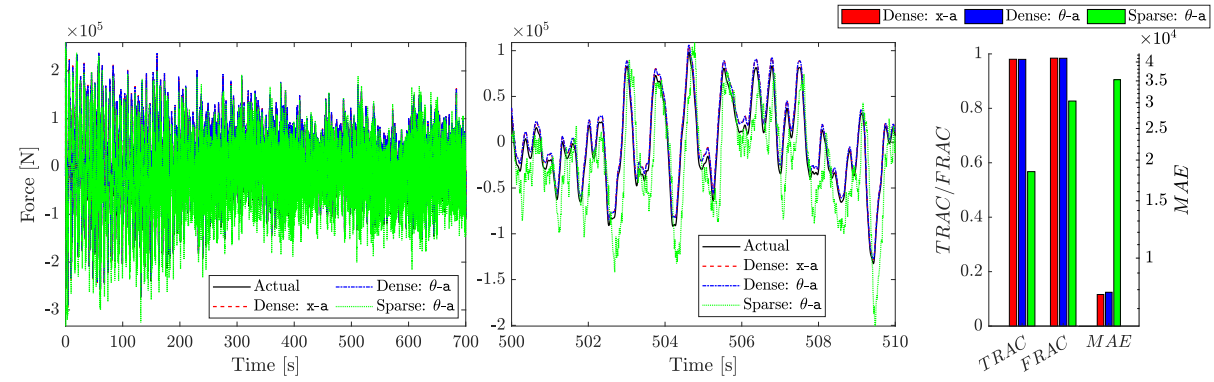
Comparing results between sensor setups 'Dense: x-a' and 'Dense: θ-a', it can be seen that the two sensor setups perform competitively. This suggests that for the task of input estimation, it is reasonable to substitute displacement sensors with inclinometers without sacrificing performance. Now considering the results from the 'Sparse: θ-a' setup, the figures show that all inputs have also been estimated with high performance. However, compared with the dense sensor setups, it can be seen that the MAEs increase for the ground level inputs (Figure 5.8). This is mainly due to a time lag in the estimated inputs (most noticeable in Figure 5.8a). This time lag is influenced by various factors, but in this case, one significant cause is the number and placement of accelerometers. This affects the terms in the direct feedthrough matrix **J** and consequently impacts the input-output coupling. This potentially motivates the use of optimal sensor placement strategies to design sensor setups for which this time lag is minimised. However, such an investigation is outside the scope of this thesis.



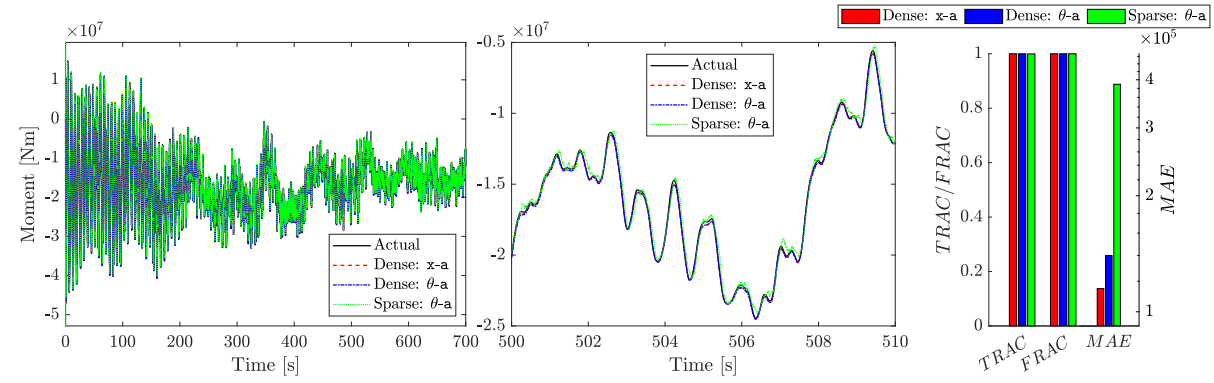
(a) Fore-Aft shear force ($\mathtt{FO}_{\mathtt{FA}}$)



(b) Fore-Aft bending moment ($\mathtt{MO}_{\mathtt{FA}}$)

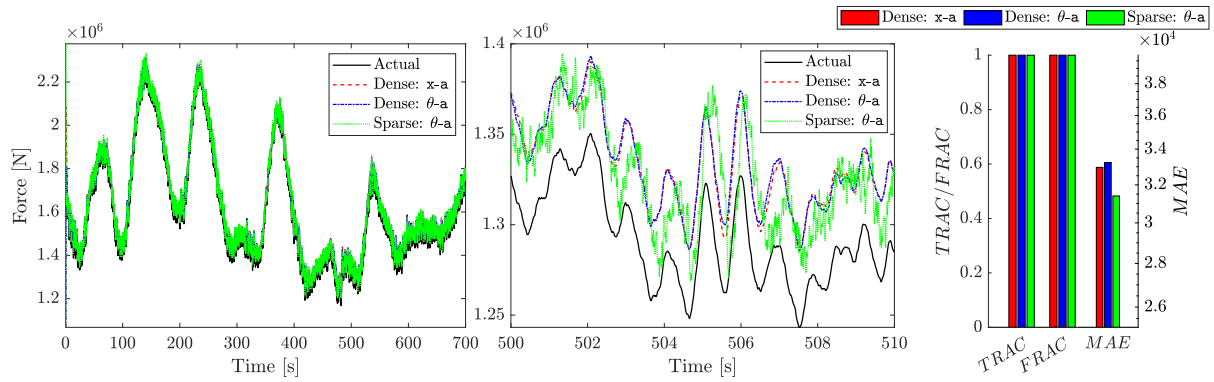


(c) Side-Side shear force ($\mathtt{FO}_{\mathtt{SS}}$)

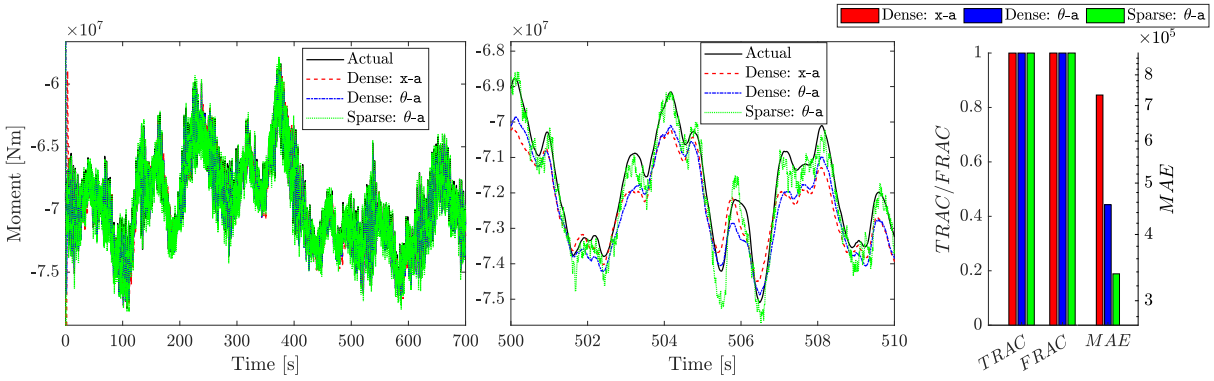


(d) Side-Side bending moment ($\mathtt{MO}_{\mathtt{SS}}$)

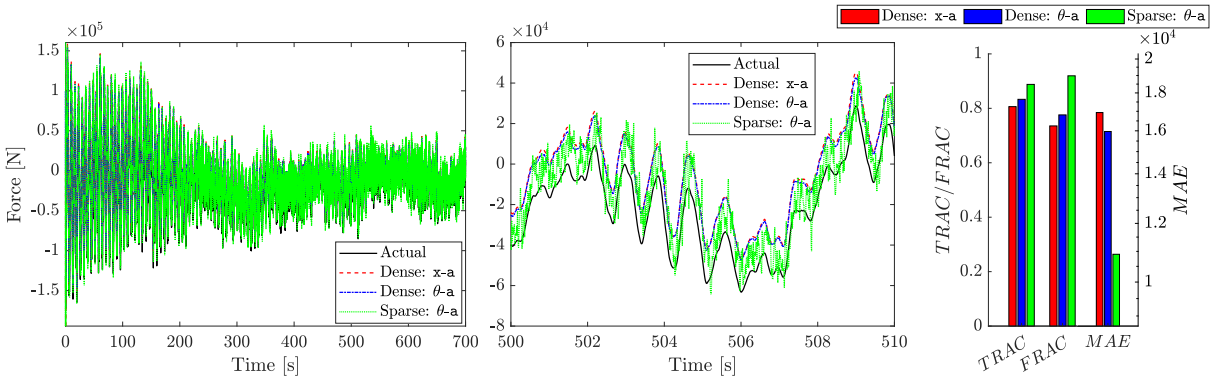
Figure 5.8: Shear forces and bending moments at the **ground level interface** for the three sensor setups (Case Study 1). *From left to right:* entire response duration, magnified excerpt, $TRAC/FRAC/MAE$ values.



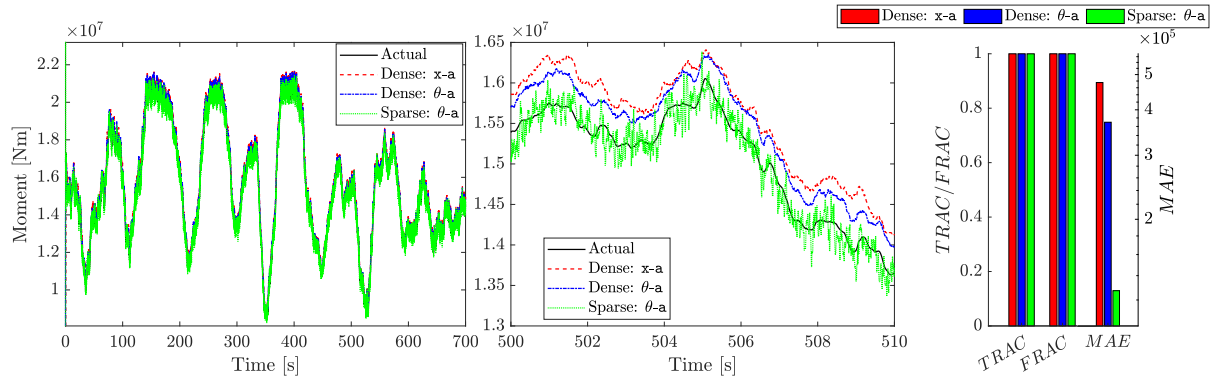
(a) Fore-Aft shear force ($\texttt{F175}_{\texttt{FA}}$)



(b) Fore-Aft bending moment ($\texttt{M175}_{\texttt{FA}}$)



(c) Side-Side shear force ($\texttt{F175}_{\texttt{SS}}$)



(d) Side-Side bending moment ($\texttt{M175}_{\texttt{SS}}$)

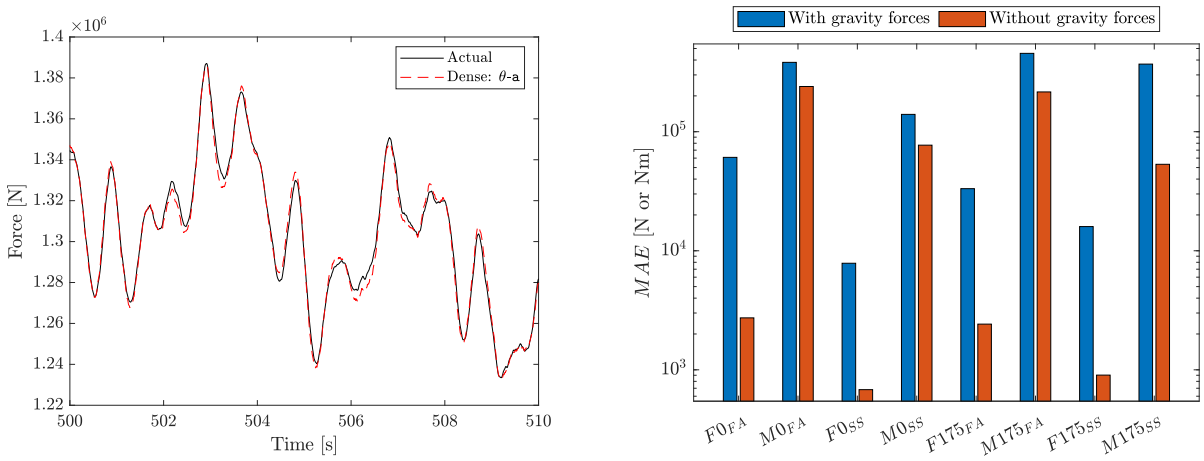
Figure 5.9: Shear forces and bending moments at the **tower top interface** for the three sensor setups (Case Study 1). *From left to right:* entire response duration, magnified excerpt, $TRAC/FRAC/MAE$ values.

An additional observation that can be made from Figures 5.8 and 5.9, is a DC offset in the estimated response for certain inputs and sensor setups. This is most noticeable for the fore-aft shear force at the tower top interface (see Figure 5.9a). One explanation for this offset relates to p-delta effects [189]. These effects arise due to the offset between the RNA's centre of mass and the top of the tower, which introduces a non-linear contribution to the lateral response from gravity forces. This non-linearity is captured by the Simulink model used to generate data. However, the linear FE model used for identification in the AKF does not account for these nonlinear p-delta effects. Consequently, the AKF has to compensate for this in the estimated inputs, resulting in the observed DC offset. To verify that this was the dominant cause of the offset, a less realistic forward simulation was run using the Simulink model to obtain measurement data without gravity forces included. Figure 5.10a presents an excerpt of the estimated fore-aft shear force at the tower top for the dataset without gravity. A comparison between Figure 5.10a and Figure 5.9a reveals a significant reduction in the offset to a negligible level. Similar observations were made for the remaining 7 inputs, so their time domain responses are not displayed here. However, Figure 5.10b presents a comparison of the MAEs across all inputs between datasets which either include or exclude gravity forces. This figure shows a reduction in the MAE for all estimated inputs when gravity forces are omitted from the simulated data. Therefore, it is reasonable to conclude that the p-delta effects were a significant cause of this offset. It is interesting to note however, that this DC offset in the inputs does not impact the performance for state and fatigue estimation.

(iv) Case Study 1: Estimating the foundation parameters

A further application of the input-state estimation results is to identify properties related to the unmodelled soil-structure interaction. As an example, this section demonstrates an approach for estimating the macro-element compliance parameters (c_{11}^0, c_{22}^0, c_{12}^0 from Equation 4.4), which were used to generate the data in this case study.

The macro-element compliances relate the estimated states to the estimated inputs at the



(a) **(b)**

Figure 5.10: Results for input estimation with synthetic data that excludes gravity forces (Case Study 1) (a) Fore-aft shear force at the tower top interface: $\texttt{F175}_{\texttt{FA}}$. Comparison between actual response and AKF-estimated response when using the 'Dense: θ-a' sensor setup. (b) MAEs in estimated inputs when the measured data do not include gravity forces, compared with the original results where gravity forces are included. Both sets of results are obtained using the 'Dense: θ-a' sensor setup.

ground level:

$$\begin{bmatrix} \texttt{x0}_{\texttt{FA}} \\ \theta\texttt{0}_{\texttt{FA}} \\ \texttt{x0}_{\texttt{SS}} \\ \theta\texttt{0}_{\texttt{SS}} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} & 0 & 0 \\ c_{12} & c_{22} & 0 & 0 \\ 0 & 0 & c_{11} & -c_{12} \\ 0 & 0 & -c_{12} & c_{22} \end{bmatrix} \begin{bmatrix} \texttt{F0}_{\texttt{FA}} \\ \texttt{M0}_{\texttt{FA}} \\ \texttt{F0}_{\texttt{SS}} \\ \texttt{M0}_{\texttt{SS}} \end{bmatrix}. \tag{5.32}$$

Given that both the ground-level states and the inputs can be successfully estimated, then the compliance parameters can be subsequently identified. While several approaches could be adopted to achieve this—such as linear regression— this study uses a classical Kalman filter [114].

To identify the compliance parameters using a Kalman filter, c_{11}, c_{22} and c_{12} were considered as states to be estimated and Equation 5.32 was transformed into the the following discrete-time

state-space model:

$$\underbrace{\begin{bmatrix} c_{11} \\ c_{12} \\ c_{22} \end{bmatrix}_k}_{\text{state vector}} = \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{\text{state matrix}} \underbrace{\begin{bmatrix} c_{11} \\ c_{12} \\ c_{22} \end{bmatrix}_{k-1}}_{\text{state vector}}$$

$$\underbrace{\begin{bmatrix} \mathtt{x0_{FA}} \\ \theta\mathtt{0_{FA}} \\ \mathtt{x0_{SS}} \\ \theta\mathtt{0_{SS}} \end{bmatrix}_k}_{\text{output vector}} = \underbrace{\begin{bmatrix} \mathtt{F0_{FA}} & \mathtt{M0_{FA}} & 0 \\ 0 & \mathtt{F0_{FA}} & \mathtt{M0_{FA}} \\ \mathtt{F0_{SS}} & -\mathtt{M0_{SS}} & 0 \\ 0 & -\mathtt{F0_{SS}} & \mathtt{M0_{SS}} \end{bmatrix}_k}_{\text{output matrix}} \underbrace{\begin{bmatrix} c_{11} \\ c_{12} \\ c_{22} \end{bmatrix}_k}_{\text{state vector}}, \tag{5.33}$$

where the subscript k indicates the current time step.

In other words, the estimated ground-level displacements and rotations are considered as observations, and the estimated ground-level forces and moments are used to define the time-varying output matrix. The state matrix is chosen as the identity matrix since the compliance values are assumed to be constant over time (i.e., their time derivatives are zero).

The state-space model of Equation 5.33 can now be used within a linear time-varying Kalman filter algorithm to recursively estimate the compliance parameters. This was carried out using the estimated results from each of the three sensor setups considered in this case study: 'Dense: x-a', 'Dense: θ-a' and 'Sparse: θ-a'.

For all sensor setups, the initial state vector was set to twice the true compliance values: $2 \times [c_{11}^0 \; c_{12}^0 \; c_{22}^0]^T$. The initial state covariance matrix was assumed to be diagonal, with entries equal to 1% of the corresponding true compliance values: $0.01 \times \text{diag}\left([c_{11}^0 \; c_{12}^0 \; c_{22}^0]\right)$. The process noise covariance matrix **Q** and the measurement noise covariance matrix **R** were tuned

manually and set as follows:

$$\mathbf{Q} = \begin{bmatrix} 10^{-20} & 0 & 0 \\ 0 & 10^{-25} & 0 \\ 0 & 0 & 10^{-25} \end{bmatrix}$$
$$\mathbf{R} = \begin{bmatrix} 10 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{5.34}$$

Figure 5.11 shows the time histories for the compliance parameters estimated using the results from the 'Dense: x-a' sensor setup. The compliance values quickly converge to stable solutions with minimal error: over the final 300 seconds, the average percentage errors were $-6.72\,\%$, $-0.04\,\%$ and $-0.42\,\%$ for c_{11}, c_{22}, and c_{12}, respectively. The slightly higher error in the lateral spring compliance (c_{11}) is likely due to the ground-level response being less sensitive to this parameter. Nevertheless, these errors are comparable to those obtained in Chapter 3 using the Bayesian model updating frameworks (with measurement error only), and are not affected by rotor nacelle assembly modelling errors discussed in Chapter 4, since the substructuring approach cuts below the nacelle.

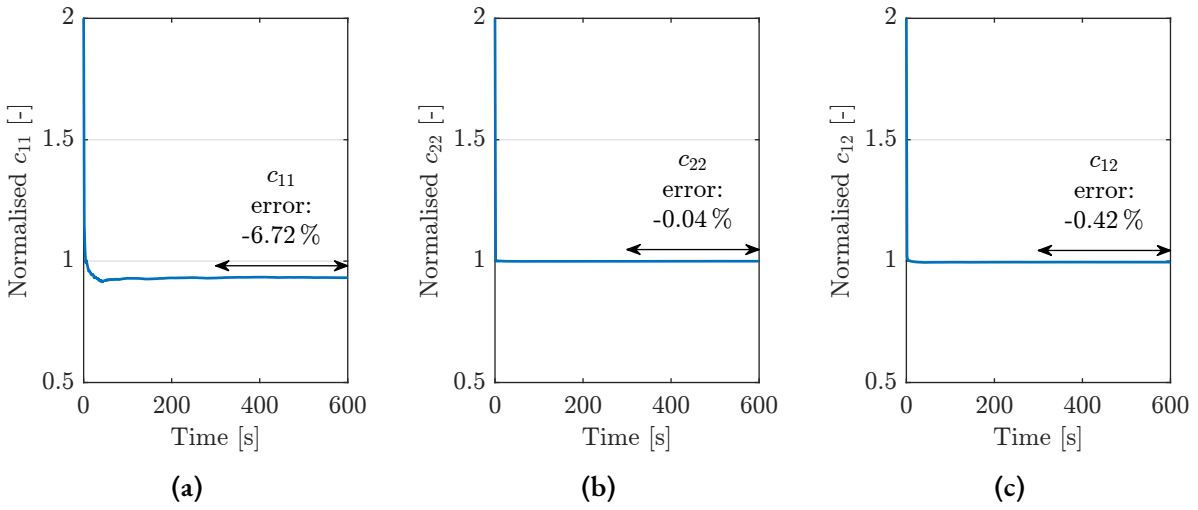


Figure 5.11: Estimated macro-element parameters from the 'Dense: x-a' sensor setup. The annotated percentage errors represent the average over the final 300 seconds. (a) Lateral compliance c_{11}; (b) Rotational compliance c_{22}; (c) Cross-coupling compliance c_{12}. Compliance values are normalised by their true values: c_{11}^0, c_{22}^0, and c_{12}^0.

However, the 'Dense: x-a' sensor setup is not practical in real-world applications due to the number of sensors and the use of displacement measurements. Figure 5.12 presents results from the sensor setups that use only inclinometers and accelerometers ('Dense: θ-a' and 'Sparse: θ-a'). In both cases, the compliance parameters are unidentifiable, with estimates that drift over time. This was an expected result, as accurate identification of the compliance parameters here relies on both the ground-level displacements and rotations being observable from the sensor setup. As discussed previously, the proposed substructuring approach cannot estimate the rigid body modes without a displacement sensor. As a result, the estimated compliance parameters exhibit the same drift seen in the displacement responses shown in Figure 5.5.

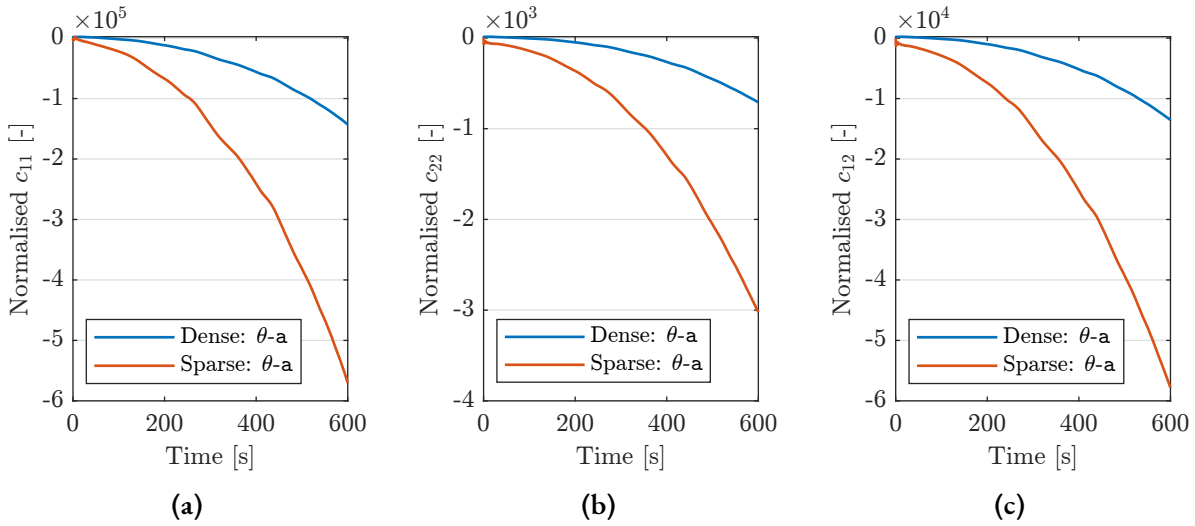


Figure 5.12: Estimated macro-element parameters from the sensor setups without displacements (i.e., only inclinometers and accelerometers). (a) Lateral compliance c_{11}; (b) Rotational compliance c_{22}; (c) Cross-coupling compliance c_{12}. Compliance values are normalised by their true values: c_{11}^0, c_{22}^0, and c_{12}^0.

Interestingly, it is still possible to identify c_{22} and c_{12} without displacement measurements. By removing the estimated displacement terms from the state-space model in Equation 5.33 (i.e., eliminating the first and third rows of the output equation), c_{22} and c_{12} then depend only on the ground-level rotation response and inputs, which were previously shown to be well estimated using inclinometers and accelerometers alone. Figure 5.13 presents the estimated c_{22} and c_{12} following this procedure and using the 'Sparse: θ-a' sensor setup. The results show that c_{22} is estimated with high accuracy, and c_{12} with reasonable accuracy.

In summary, displacement sensors are necessary to identify all three macro-element param-

eters. However, the rotational compliance c_{22} and the cross coupling compliance c_{12} can still be identified using only inclinometers and accelerometers.

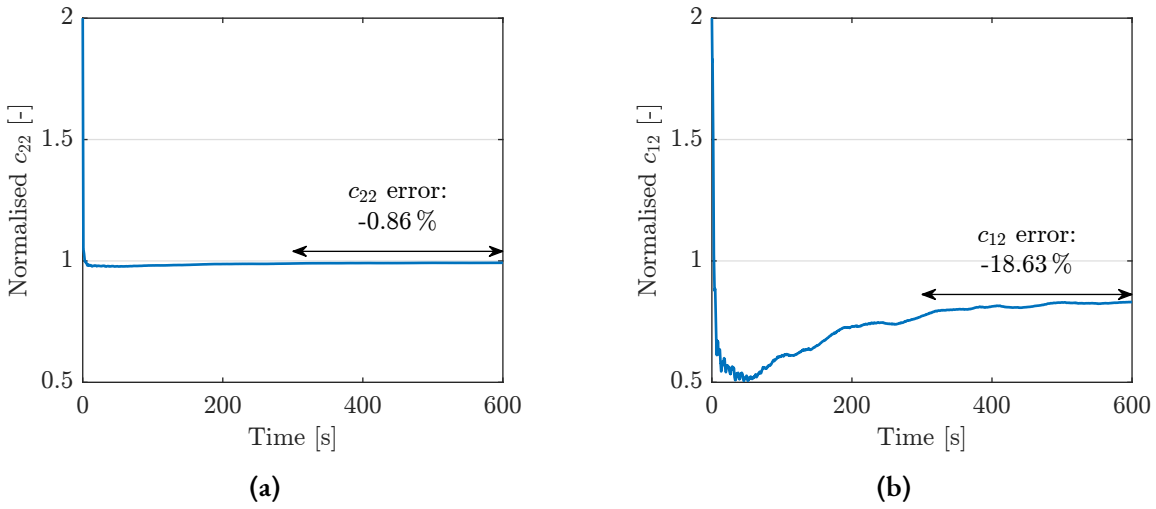


Figure 5.13: Estimated macro-element parameters from the 'Sparse: θ-a' sensor setup when estimated displacements are ignored. The annotated percentage errors represent the average over the final 300 seconds. (a) Rotational compliance c_{22} (b) Cross-coupling compliance c_{12}. Compliance values are normalised by their true values: c_{22}^0, and c_{12}^0.

5.4.5 Case Study 2: wind and wave loading

In the previous case study, still water conditions were assumed, so that no forces resulted from hydrodynamic wave loading. Building upon these results, this case study used simulated data that considered both aerodynamic and hydrodynamic loading. To generate this synthetic data, the same turbulent wind field from Case Study 1 was used, with the wind field aligned in the fore-aft direction. For the hydrodynamics, unidirectional wave loading was applied, which was calculated based on a JONSWAP spectrum defined by a significant wave height of 4.5 m and a peak wave period of 11 seconds. To explore the impact of misaligned wind and waves, two simulations were run: one with a wave yaw angle of 0° and another with a wave yaw angle of 30°. Here, a yaw angle of 0° indicates that the waves are aligned in the fore-aft direction and a yaw angle of 90° indicates that the waves are aligned in the side-side direction, as illustrated in Figure 5.18.

As a result of the wave loading, forces and moments act on each submerged node in the FE model. Attempting to estimate each of these forces and moments individually would result in an excessive number of inputs to be estimated for the given sensor setups. However, this

chapter's focus is not on estimating wave inputs but on interface forces and substructure states. Therefore, two approaches were implemented to account for the wave loads, while maintaining high accuracy in the quantities of interest.

In the first approach, the input vector was identical to that of Case Study 1, which included the 8 interface inputs:

$$\boldsymbol{p}_1 = \begin{bmatrix} \mathtt{FO}_{\mathtt{FA}} & \mathtt{MO}_{\mathtt{FA}} & \mathtt{FO}_{\mathtt{SS}} & \mathtt{MO}_{\mathtt{SS}} & \mathtt{F175}_{\mathtt{FA}} & \mathtt{M175}_{\mathtt{FA}} & \mathtt{F175}_{\mathtt{SS}} & \mathtt{M175}_{\mathtt{SS}} \end{bmatrix}^T. \tag{5.35}$$

In other words, with this set of inputs, the AKF does not attempt to estimate the wave excitation, which is non-zero for Case Study 2. This approach was included primarily to investigate the necessity of estimating unmeasured wave inputs.

In the second approach, the force applied on each submerged node followed a triangular distribution. Here, the wave force at the ground level was constrained to be 0 and linearly increased to a value of `Fw` at the mean water level. This approximation is visually represented in Figure 5.14.

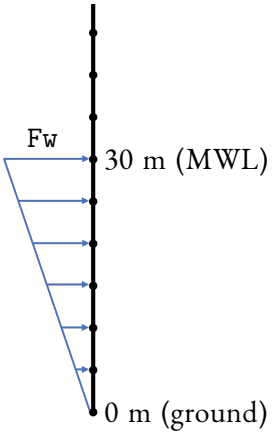


Figure 5.14: Triangular wave loading approximation

It should be noted that in the simulations the spatial distribution of the wave forces along the tower changes over time as it is frequency dependent. As such the triangular approximation here corresponds to a simplification of the actual wave distribution that generated the data.

To account for potential misalignment or spreading of waves, two triangular distributions

were estimated: one in the fore-aft direction and one in the side-side direction. As a result, the input vector for the second approach, $\boldsymbol{p}_2$, included two additional inputs:

$$\boldsymbol{p}_2 = \begin{bmatrix} \mathtt{F0}_{\mathtt{FA}} & \mathtt{M0}_{\mathtt{FA}} & \mathtt{F0}_{\mathtt{SS}} & \mathtt{M0}_{\mathtt{SS}} & \mathtt{F175}_{\mathtt{FA}} & \mathtt{M175}_{\mathtt{FA}} & \mathtt{F175}_{\mathtt{SS}} & \mathtt{M175}_{\mathtt{SS}} & \mathtt{Fw}_{\mathtt{FA}} & \mathtt{Fw}_{\mathtt{SS}} \end{bmatrix}^T, \quad (5.36)$$

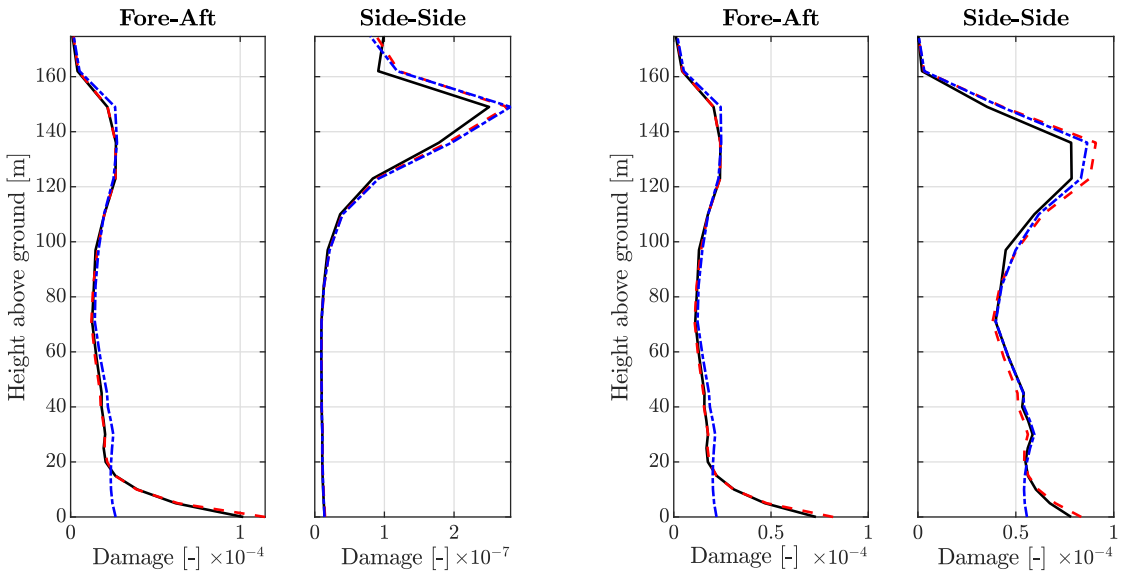
where $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$ are the triangular approximation's forces at the MWL in the fore-aft and side-side directions, respectively. It should be noted that the values of $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$ are assumed to be independent for identification purposes, even though in the simulation these were the projections of a wave excitation applied at a yaw angle. In other words, the identification model is not making use of this additional information. In fact, the yaw angle can be identified from the estimated wave inputs, as demonstrated in Section 5.4.5 (ii).

(i) Case Study 2 results

Regarding the influence of the three sensor setups, similar observations and conclusions to those discussed in Case Study 1 were also found here. Therefore, results are only presented for the 'Sparse: θ-a' sensor setup. First a discussion is made on the damage distributions, before presenting the results from input estimation. A comparison is drawn between the estimated results obtained from using the two input vectors $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$. Since the damage distributions provide insight into the performance of state estimation for the entire substructure, the response at individual DOFs is not presented here.

Figure 5.15 shows the damage distributions caused by the two excitations studied in Case Study 2, corresponding to wave yaw angles of 0° and 30°.

These figures demonstrate that, overall, the fatigue damage has been accurately estimated. However, a drop in performance was observed at nodes closer to the ground-level interface when the wave inputs were not estimated and were instead assumed to be zero (i.e., when using $\boldsymbol{p}_1$), whereas the triangular approximation achieved a level of accuracy comparable to that of Case Study 1 (see Figure 5.7 vs Figure 5.15). It is interesting to note that when the waves are aligned in the fore-aft direction, the estimated damage distribution in the side-side direction is similar for both input vectors $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$ (see Figure 5.15a). This was expected due

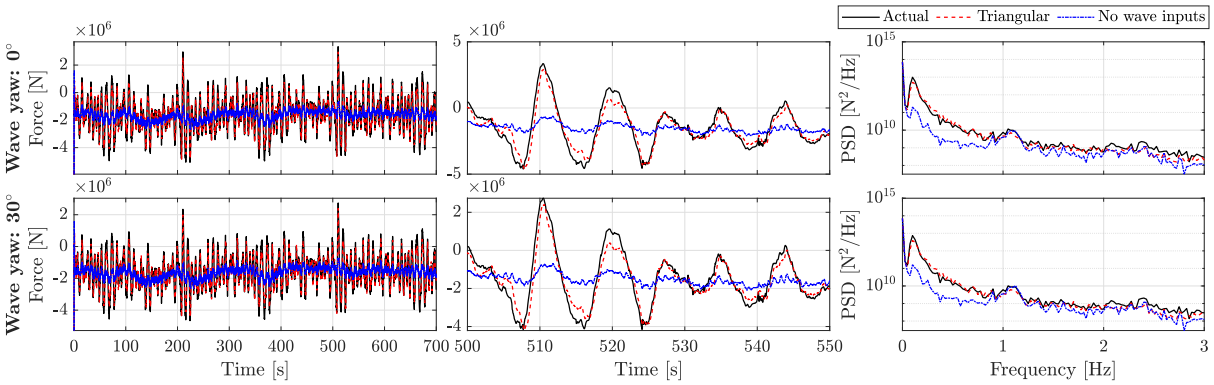


(a) Wave yaw angle: 0° **(b)** Wave yaw angle: 30°

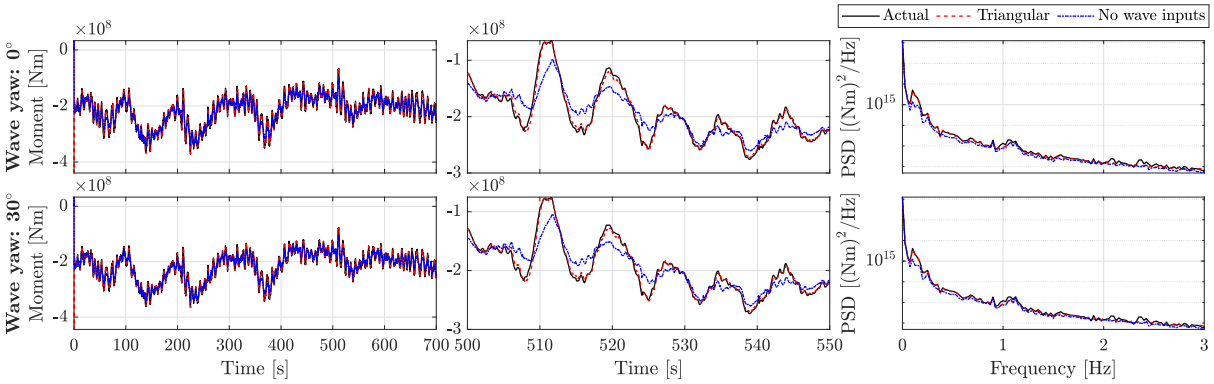
Figure 5.15: Case Study 2 damage distributions. (solid black line: actual damage, blue dash-dotted line: damage when no wave inputs were considered, red dashed line: damage when triangular approximation was used).

to the symmetry of the substructure, resulting in a reduced-order model that is decoupled in the fore-aft and side-side directions. Therefore, when waves are aligned in the fore-aft direction, there is minimal impact on the estimated response in the side-side direction, rendering $\mathtt{Fw}_{\mathtt{SS}}$ unnecessary. However, when a yaw angle is present, a drop in accuracy was observed when wave inputs were excluded (see Figure 5.15b), and the inclusion of $\mathtt{Fw}_{\mathtt{SS}}$ in the estimated quantities was required to restore previous levels of accuracy.

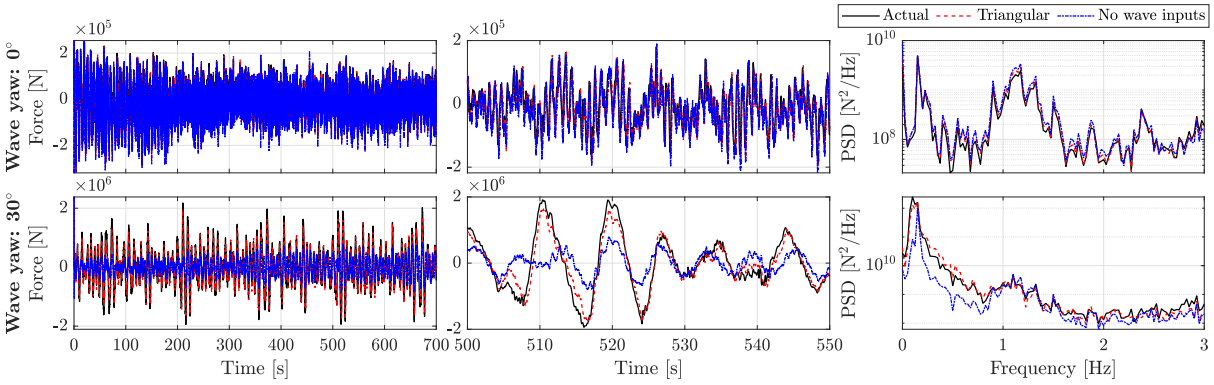
Figure 5.16 presents the results for input estimation at the ground-level interface. Here it can be seen that when wave inputs are not included, the inputs are underestimated. Examining the power spectral densities reveal that although the peaks corresponding to the first and second tower modes are estimated at the correct frequencies, the magnitude has been underestimated around the first mode, contributing to the errors seen in the time domain. One possible explanation for this is that the interface forces and moments have to compensate for the additional loading, which is present due to the waves but not estimated. Therefore, it can be concluded that for accurate estimation of the interface inputs, it is important to account for the presence of wave loading in the input vector. Moreover, it can be seen that using a simple triangular approximation for this purpose yields highly accurate estimated inputs.



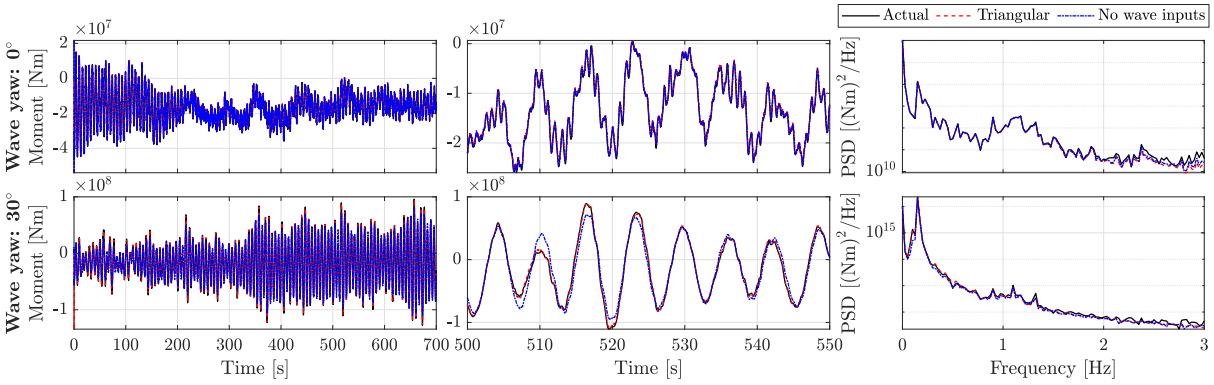
(a) Fore-aft shear force ($\mathtt{FO}_{\mathtt{FA}}$)



(b) Fore-aft bending moment ($\mathtt{MO}_{\mathtt{FA}}$)



(c) Side-side shear force ($\mathtt{FO}_{\mathtt{SS}}$)



(d) Side-side bending moment ($\mathtt{MO}_{\mathtt{SS}}$)

Figure 5.16: Ground level interface input estimation for Case Study 2 (wind and wave loading). For each sub-figure, left to right shows: entire response duration, 10 s magnified excerpt, power spectral density. Top row shows results for wave yaw = 0° and the bottom row shows wave yaw = 30°.

It is worth mentioning that the presence of waves had less impact on the accuracy of the estimated inputs at the tower-top interface, with results comparable to those of Case Study 1. As a result, they are not included here.

(ii) Case Study 2: estimating wave misalignment from $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$

The angle of misalignment between wind and waves can have significant impacts on fatigue, installation, and operational efficiency [190–192]. Its effect on fatigue is evident in Figure 5.15, where the side-side damage distribution from the $30°$ simulation shows an increase in accumulated damage of approximately three orders of magnitude compared to the $0°$ simulation. Therefore, the ability to estimate the wave misalignment angle is of clear engineering interest.

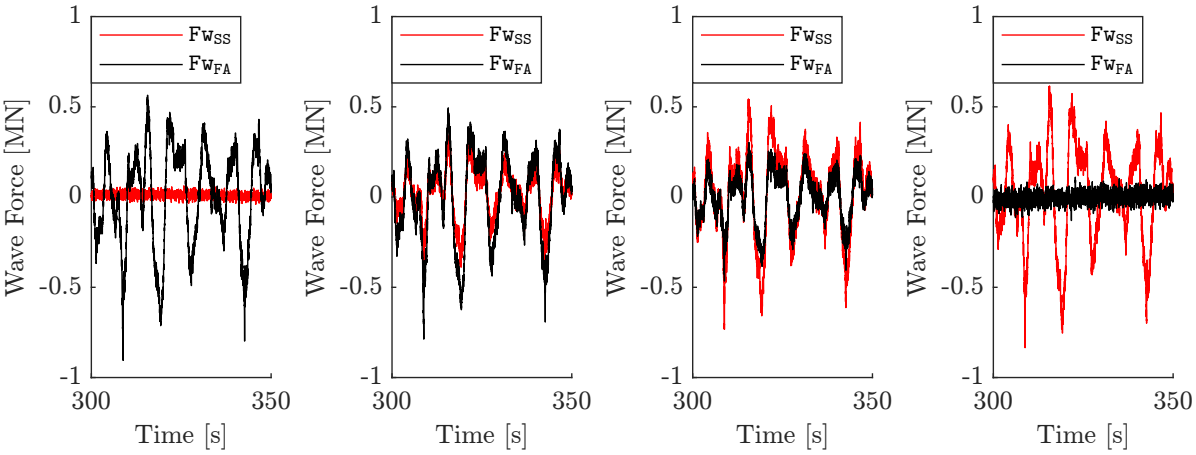
This section demonstrates how the dominant wave direction can be inferred from the estimated triangular wave approximation inputs: $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$. Alongside the previously considered simulations with wave yaw misalignment of $0°$ and $30°$, two further cases with misalignments of $60°$ and $90°$ were included to provide additional data. For each dataset, the 'Sparse: θ-a' was used to estimate $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$. In these simulations, the wave misalignment angle is assumed to be constant over the ten-minute duration. While in reality it may vary, this assumption is reasonable because the direction of the waves typically evolves slowly compared to the time periods considered here.

Figure 5.17 shows 50 second excerpts of these estimated wave inputs for each simulation. A visual inspection reveals that, as the wind-wave misalignment increases, the overall magnitude of the $\mathtt{Fw}_{\mathtt{FA}}$ decreases, while that of $\mathtt{Fw}_{\mathtt{SS}}$ increases. This trend suggests that the estimated wave inputs are capturing the effect of the changing yaw angle.

Based on the estimated values of $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$, the resultant wave force can be resolved as shown in Figure 5.18. The direction of this vector defines the wave misalignment angle, α_w, which is given by:

$$\tan \alpha_w = \frac{\mathtt{Fw}_{\mathtt{SS}}}{\mathtt{Fw}_{\mathtt{FA}}}. \tag{5.37}$$

While Equation 5.37 can, in principle, be used to compute α_w at each time step, directly



(a) 0° simulation **(b)** 30° simulation **(c)** 60° simulation **(d)** 90° simulation

Figure 5.17: Time series excerpts of estimated triangular wave inputs ($\mathrm{Fw_{FA}}$ and $\mathrm{Fw_{SS}}$) from simulations with a range of wind-wave misalignment angles.

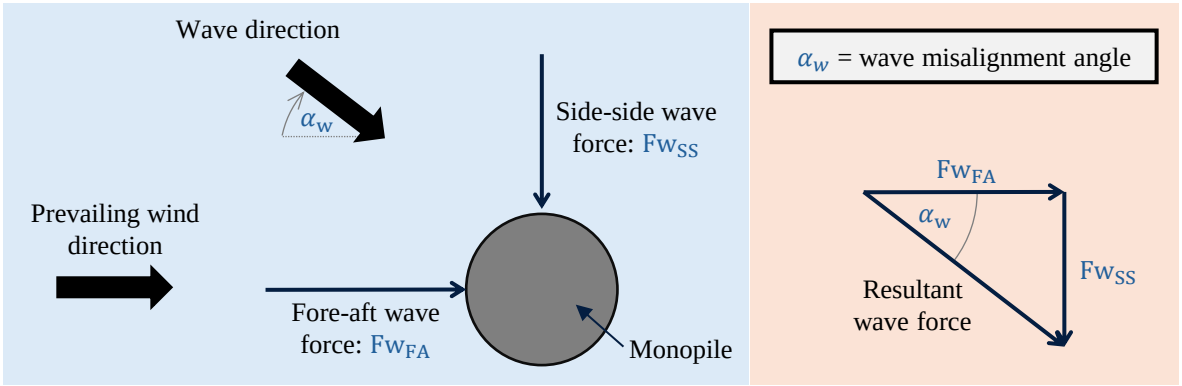


Figure 5.18: Illustration showing the wind and wave directions, and how the resultant of the estimated wave inputs $\mathrm{Fw_{FA}}$ and $\mathrm{Fw_{SS}}$ relates to the wave misalignment angle.

applying this equation resulted in a noisy or spiky signal. This was attributed to estimation errors in $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$, as well as numerical instability arising from division by small values. Therefore, two alternative and more sophisticated methods were implemented to estimate α_w. The first considers a time-varying Kalman filtering approach similar to that used for estimating the compliance parameters in Section 5.4.4 (iv), while the second is based on the Arias Intensity, a measure traditionally applied in earthquake engineering [193]. Both methods are described in more detail below.

Approach 1: time-varying Kalman filtering

Assuming constant α_w and defining a new state variable as $x = \tan \alpha_w$, Equation 5.37 can be rearranged into a discrete-time state-space model:

$$\begin{aligned} x_k &= x_{k-1} \\ \mathtt{Fw}_{\mathtt{FA,k}} &= \mathtt{Fw}_{\mathtt{SS,k}} \cdot x_k, \end{aligned} \tag{5.38}$$

where the subscript k denotes the current time step. In this formulation, $\mathtt{Fw}_{\mathtt{FA,k}}$ is considered as an observation, and $\mathtt{Fw}_{\mathtt{SS,k}}$ maps the state variable, x_k, to the output.

After estimating x_k using the time-varying Kalman filter algorithm, the wave misalignment angle can be recovered as:

$$\alpha_{w,\,k} = \arctan x_k. \tag{5.39}$$

Figure 5.19a presents the time histories of α_w for each simulated dataset. The results demonstrate that this method provides highly accurate estimates of α_w based on the estimated wave inputs.

Approach 2: Arias Intensity

The Arias Intensity Tensor [193], $\mathbf{I}_A$, provides a means to determine the dominant directions of energy content in two orthogonal signals, such as $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$. It can be defined based on

the estimated wave inputs as:

$$\mathbf{I}_A = \frac{\pi}{2g} \int_0^T \begin{bmatrix} \mathtt{Fw}_{\mathtt{FA}}^2 & \mathtt{Fw}_{\mathtt{FA}} \cdot \mathtt{Fw}_{\mathtt{SS}} \\ \mathtt{Fw}_{\mathtt{SS}} \cdot \mathtt{Fw}_{\mathtt{FA}} & \mathtt{Fw}_{\mathtt{SS}}^2 \end{bmatrix} dt, \tag{5.40}$$

where T is the duration of the signals and g is the acceleration due to gravity.

An eigendecomposition of $\mathbf{I}_A$ results in two orthogonal eigenvectors that represent the principal axes of energy distribution. The wave misalignment angle α_w can then be defined as the angle between the major axis (associated with the largest eigenvalue) and the fore-aft direction.

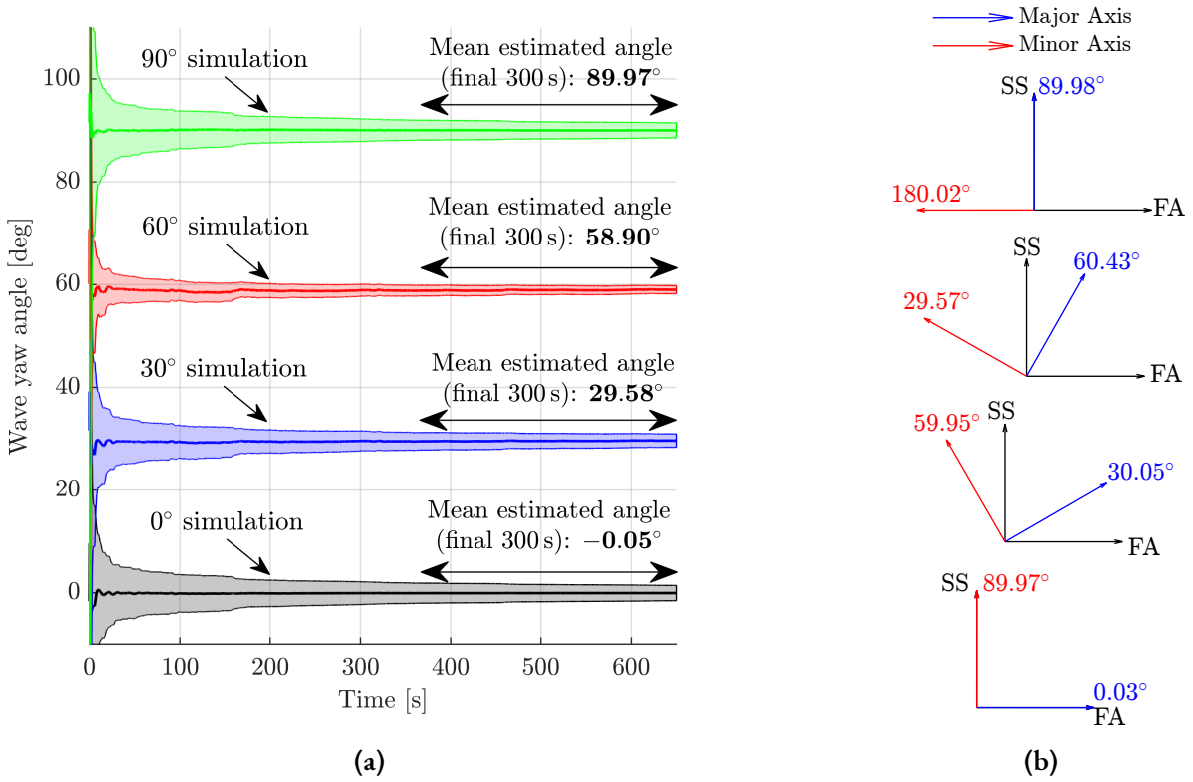
Note that the constant $\pi/(2g)$ in Equation 5.40 does not affect the orientation of the eigenvectors, but it is retained here for consistency with the original definition of the Arias Intensity in earthquake engineering.

Figure 5.19b shows the orientations of the major and minor axes obtained from the eigenvectors of $\mathbf{I}_A$ for each simulation. The figures show that angles between the major axis and the fore-aft direction provide accurate estimates of the wave yaw misalignment angles used to generate the simulated data.

—

Comparing the two methods, the recursive structure of the Kalman filtering approach makes it advantageous for online implementation and better suited to conditions where the wave direction varies over time within a given dataset. On the other hand, the Arias Intensity approach is attractive for its simplicity, as it requires no tuning or hyperparameters, unlike the Kalman filter.

Nonetheless, Figure 5.19 shows that both methods perform comparably well, providing precise and reliable estimates of α_w from the estimated wave inputs $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$. Furthermore, although $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$ cannot be directly validated against a known solution in this numerical case study, the results provide strong evidence that $\mathtt{Fw}_{\mathtt{FA}}$ and $\mathtt{Fw}_{\mathtt{SS}}$ are physically meaningful.



(a) (b)

Figure 5.19: Estimated wave misalignment angles from simulations with varying wind-wave misalignment, using (a) the Kalman filtering approach and (b) the Arias Intensity tensor approach. In (a), the solid lines indicate mean values for α_w, and the shaded area represents the 95 % confidence interval. In (b), the direction of the major axis aligns with the wave misalignment angle and the subplots from top to bottom correspond to simulations with true misalignment angles of 90°, 60°, 30°, and 0°, respectively.

5.5 Summary

In this chapter, a substructuring approach has been applied for input-state estimation of the tower-monopile component of an offshore wind turbine. By applying a cut at the top of the tower and at the ground level, the resulting substructure only requires modelling details which relate to the tower and monopile. This is advantageous, as it significantly reduces the model uncertainty.

The accuracy of the proposed approach was evaluated through numerical case studies conducted on a 15 MW offshore wind turbine. The Simulink Model was used to generate data, while a reduced-order linear finite-element model was used for identification purposes. These case studies demonstrated the capability of the suggested framework to achieve high performance for state and input estimation. Regarding the inputs, an offset was observed in the estimated

interface forces, which was attributed to the linear finite-element model not accounting for the non-linear p-delta effects due to gravity forces. This was confirmed by simulating a less realistic dataset, where forces due to gravity were not considered; for this case, the offset was negligible. Despite this offset, the overall trend in the estimated inputs was still found to be captured well, with high accuracy in the frequency domain. Similar results for the inputs were obtained irrespective of whether displacement or inclinometer measurements were employed. Moreover, the sparse sensor setup was shown to produce accurate estimates for the interface inputs. However, due to having fewer accelerometers, the response was estimated with a time lag, and was attributed to the reduced contribution from terms in the direct feedthrough matrix $\mathbf{J}$.

The states were also found to be estimated very accurately and were not affected by the aforementioned p-delta effects. However, when the displacement sensors were swapped for inclinometers in the measurement set, a drift was observed in the estimated displacements. This was attributed to the fact that the considered substructure is unconstrained, so that inclinometers and accelerometers can't provide any information on the absolute position of the wind turbine. As a consequence of observability of this model and the specific sensor setup, the Augmented Kalman Filter is not able to estimate the rigid body response of the wind turbine. This conclusion was reinforced by examining the estimated modal displacements, where it was shown that all modes apart from the rigid body modes were accurately estimated. Note that if one displacement sensor is available to supplement the inclinometers and accelerometers, then the rigid body motion can also be estimated. This conclusion could potentially advise future measurement campaigns to consider installing at least one GPS sensor, for example in the nacelle, to provide this information. However, since the rigid body motion does not contribute to the deformation of the substructure, the predicted accumulated damage was found to be estimated with equal accuracy regardless of the sensor type.

In a second case study, the influence of both wind and wave loading was investigated. Here it was shown that the performance declines if inputs are not considered to account for the wave loading. Consequently, such excitations should be included in the inputs to be estimated. It was then demonstrated that a triangular distribution can be used to approximate the wave loading.

This straightforward approximation only required two additional inputs and was shown to yield results with high accuracy in both the estimated damage distributions and the interface inputs. Furthermore, two methods were presented for inferring the wind-wave misalignment angle from these estimated wave inputs, both of which demonstrated strong performance.

Due to the significantly reduced model uncertainty, the proposed approach is considered to be a more robust method for estimating the states and inputs of the tower substructure, compared with trying to model the entire system. In addition, the method is computationally efficient and can be implemented in real-time, making it suitable for online applications.

A key advantage of the framework lies in its minimal assumptions, meaning that very few would need to be removed for real-world applications. However, it is important to highlight its main limitations. In this study, the substructure model properties were assumed to be known, whereas in real-world applications, there may exist uncertainties related to the definition of certain parameters, such as water level, hydrodynamic added mass coefficient, and modelling of the transition piece. If such parameters are thought to introduce significant modelling error, then the framework could be developed to include parameter estimation. In addition, accelerometers alone were found to be insufficient to accurately estimate inputs, resulting in a decrease in performance for state and fatigue estimation. Consequently, the number of inclinometers required to supplement the accelerometers used in this work may be greater than that which is typically available in practice. While collocation at the interface degrees of freedom was assumed in this work, this is not a strict requirement for successful input-state estimation; however, such an arrangement may still improve performance. An optimal sensor placement strategy was not considered in this work, and further studies are needed to explore this aspect.

An example was presented showing how the macro-element compliance parameters can be identified from the input-state estimation results. However, the identifiability of all three parameters relied on the observability of rigid-body modes, and therefore required displacement sensors. While including a foundation model could allow for observability of the rigid displacements, the suggested approach allows for removing assumptions about the soil-structure interaction model without consequences for the estimated fatigue life of the structure.

To build on the successful performance of this substructuring approach for input-state estimation, the next chapter further develops the framework by focusing on the identification of foundation parameters using a sensor setup representative of real-world applications, without relying on displacement sensors.

CHAPTER 6

A Substructuring Approach for Parameter Estimation

6.1 Introduction

Chapter 5 demonstrated the benefits of a substructuring approach for input-state estimation of an unconstrained tower-monopile component, where both the rotor nacelle assembly (RNA) and the soil-structure interaction were excluded from the model. In this chapter, the foundation is reintroduced, and the substructuring framework is extended to allow for identification of the foundation parameters.

As discussed in Chapter 1, the joint estimation of parameters, inputs and states been explored in previous studies [132, 143, 145, 146], where variations of the Kalman filter are used to recursively estimate these quantities simultaneously. While these methods have demonstrated successful identification of model parameters, their performance can be sensitive to the choice of process and measurement noise covariances, potentially leading to biased or unstable results. Nabiyan *et al.* [113] implemented a sequential method for updating the inputs and the foundation stiffness of a 2 MW offshore wind turbine, based on the approach proposed by Ebrahimian *et al.* [149]. Their method is based on Bayesian inference to iteratively update model parameters and inputs by minimising the difference between measurements and model-predictions in the time-domain. The authors show the successful identification of the rotational foundation stiffness using a mechanics-based model, where the RNA is represented as a lumped mass.

An alternative framework for parameter estimation is proposed in this chapter, that leverages a time-domain likelihood function derived from input-state estimation results obtained using an Augmented Kalman Filter [118]. While this is similar in concept to the frameworks in

[113, 149], the proposed method avoids having to model the rotor nacelle assembly by using a substructuring approach.

The mathematical formulation of the new framework is described in Section 6.2. A numerical case study is then presented in Section 6.3 to evaluate its performance for identifying the macro-element foundation parameters of an offshore wind turbine. This case study is designed to complement the real-world application considered in the following chapter. The influence of how the underlying foundation stiffness affects parameter identifiability is explored. Additionally, it is shown that the results are largely insensitive to errors in the hydrodynamic added mass parameters.

6.2 Time-Domain Bayesian Inference with Substructuring

To formulate a likelihood function based on time-domain data, knowledge of the inputs driving the dynamic system are required. To obtain these, input-state estimation methods [194] can be used to estimate the unknown inputs as shown in Chapter 5, allowing a direct comparison between model predictions and measurements in the time domain for a given realisation of the uncertain foundation parameters.

Starting from the stochastic augmented state-space model introduced in Chapter 5 (Equations 5.18–5.19), the dependence on uncertain foundation parameters $\boldsymbol{\theta} \in \mathbb{R}^{n_\theta}$, where n_θ denotes the number of parameters, can be included as:

$$\boldsymbol{x}_k^a = \mathbf{A}_d^a(\boldsymbol{\theta})\boldsymbol{x}_{k-1}^a + \boldsymbol{w}_{k-1} \tag{6.1}$$

$$\boldsymbol{y}_k = \mathbf{G}^a(\boldsymbol{\theta})\boldsymbol{x}_k^a + \boldsymbol{v}_k, \tag{6.2}$$

Here, the augmented system matrix $\mathbf{A}_d \in \mathbb{R}^{(n_s+n_p)\times(n_s+n_p)}$ and output matrix $\mathbf{G}^a \in \mathbb{R}^{n_y\times(n_s+n_p)}$ are parameterised by $\boldsymbol{\theta}$, and can be defined by the structural model (described later in Section 6.3.2). n_s, n_p, and n_y are the number of states, inputs, and measurements, respectively. As in Chapter 5, the process noise $\mathbf{w}_k$ and measurement noise $\mathbf{v}_k$ are assumed to be independent zero-mean Gaussian random variables, with covariance matrices $\mathbf{Q}^a \in \mathbb{R}^{(n_s+n_p)\times(n_s+n_p)}$ and $\mathbf{R} \in \mathbb{R}^{n_y\times n_y}$, as defined in Equations 5.20–5.21.

The aim is to derive a probabilistic model for the prediction error $\boldsymbol{\eta}(\boldsymbol{\theta}) \in \mathbb{R}^{n_y}$, which evaluates the discrepancy between the model predictions $\boldsymbol{y}(\boldsymbol{\theta}) \in \mathbb{R}^{n_y}$ and measurements $\hat{\boldsymbol{d}} \in \mathbb{R}^{n_y}$. This is achieved by applying a similar concept to Rao-Blackwellisation [195], whereby $\boldsymbol{\theta}$ is marginalised out of the state-space model (Equations 6.1 and 6.2). Given a realisation of $\boldsymbol{\theta}$, the Augmented Kalman Filter (AKF) [118] can then be used for input-state estimation.

An overview of the AKF algorithm was provided in the previous chapter (Algorithm 2). Once $\boldsymbol{\theta}$, $\mathbf{R}$ and $\mathbf{Q}^a$ are specified, the AKF uses measurements at each time step k, to compute the posterior mean $\bar{\boldsymbol{x}}_k^a$ and covariance $\boldsymbol{\Sigma}_k^{x^a}$ of the augmented state vector. These estimates (or a subset thereof) are then passed through the output equation to evaluate the model-predicted output vector $\boldsymbol{y}_k(\boldsymbol{\theta})$ and its corresponding covariance matrix $\boldsymbol{\Sigma}_k^y$:

$$\boldsymbol{y}_k(\boldsymbol{\theta}) = \mathbf{G}^a(\boldsymbol{\theta})\bar{\boldsymbol{x}}_k^a, \tag{6.3}$$

$$\boldsymbol{\Sigma}_k^y = \mathbf{A}_d^a \boldsymbol{\Sigma}_k^{x^a} \mathbf{A}_d^{aT} + \mathbf{R}. \tag{6.4}$$

The prediction error at time step k is computed as:

$$\boldsymbol{\eta}_k(\boldsymbol{\theta}) = \boldsymbol{y}_k(\boldsymbol{\theta}) - \hat{\boldsymbol{d}}_k. \tag{6.5}$$

The likelihood function can then be defined by modelling $\boldsymbol{\eta}_k(\boldsymbol{\theta})$ as a zero-mean multivariate Gaussian [89]:

$$p(\hat{\boldsymbol{d}}_k|\boldsymbol{\theta}) = (2\pi)^{-n_\theta/2} \det\left(\boldsymbol{\Sigma}_k^y\right)^{-1/2} \exp\left(-\frac{1}{2}\boldsymbol{\eta}_k^T \left(\boldsymbol{\Sigma}_k^y\right)^{-1} \boldsymbol{\eta}_k\right). \tag{6.6}$$

For a given batch of data, the overall likelihood function is defined as the average log-likelihood across all time steps:

$$p(\hat{\boldsymbol{d}}_{N_{\text{burn}}:N_{\text{batch}}}|\boldsymbol{\theta}) = \frac{1}{N_{\text{batch}} - N_{\text{burn}}} \sum_{k=N_{\text{burn}}}^{N_{\text{batch}}} \ln p(\hat{\boldsymbol{d}}_k|\boldsymbol{\theta}) \tag{6.7}$$

where N_{batch} is the total number of time steps in the batch of data and N_{burn} is the number of initial measurements discarded to allow the AKF estimates to converge.

It should be noted that although the derivation is presented here using the AKF for input-state estimation, the framework can be generally applied using any variant of the Kalman filter that performs the same task [194]. The AKF was selected here due to the linear substructure model (see Section 6.3.2), its relatively straightforward implementation and its low computational cost.

6.3 Numerical Example

6.3.1 Wind Turbine Model for Data Generation

Figure 6.1a illustrates the finite-element model used for data generation, where the monopile, transition piece and tower were represented using 1D Timoshenko beam elements. The soil-structure interaction was modelled using the macro-element foundation and the rotor nacelle assembly was represented as a lumped mass with associated inertia. To complement the real-world case study presented in the following chapter, the specific model was developed to replicate the large-scale turbine considered there.

For the tower, transition piece, and monopile, detailed diameter and wall thickness distributions were used to define the stiffness and mass matrices for each element. Material properties were assumed constant for all components, with Young's modulus of 210 GPa, density of 7850 kg/m^3, Poisson's ratio of 0.3 and shear coefficient of 0.5. For submerged elements, the mass density was increased to account for the water levels inside and outside the monopile, which were assumed equal to eachother and the mean water level. The internal and external added masses, m_i^{add} and m_e^{add}, were defined as:

$$m_i^{\text{add}} = \rho_w \frac{\pi D_i^2}{4} \tag{6.8}$$

$$m_e^{\text{add}} = \rho_w \frac{\pi D_o^2}{4} \times C_a, \tag{6.9}$$

where $\rho_w = 1025$ kg/m^3 is the water density, and D_i and D_o are the monopile's inner and outer diameter, respectively. C_a is the hydrodynamic added mass coefficient, which was initially set equal to 1, corresponding to the theoretical value for flow around a cylinder [196]. Therefore, the added mass for submerged elements was defined by three parameters: the internal water

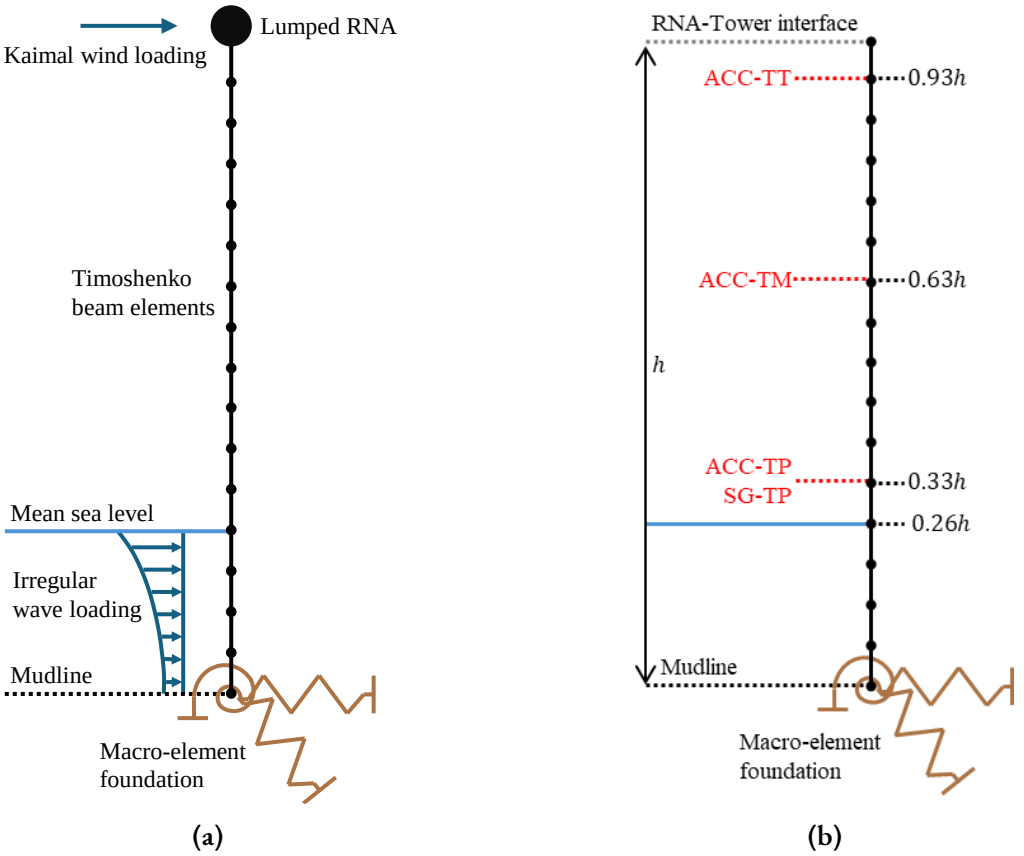


Figure 6.1: (a) Finite element (FE) model used for data generation. (b) FE model for the tower-foundation substructure and locations of sensors. The FE model is 3D, constructed with 1D beam elements, but the figure provides a 2D schematic representation in the fore–aft plane. The side–side direction is perpendicular to the figure (into the page).

level ζ_i, external water level ζ_e and C_a. Although assumptions must be made about these parameters, Section 6.3.3 (ii) shows that they have minimal influence on the estimated foundation parameters for the numerical case study considered in this chapter.

Secondary structural components were represented as lumped masses and included in the global mass matrix.

A baseline macro-element foundation, $\mathbf{K}^0_{\text{macro}}$ (or $\mathbf{C}^0_{\text{macro}}$), was used for data generation and validation purposes:

$$\mathbf{K}^0_{\text{macro}} = \begin{bmatrix} k^0_{11} & k^0_{12} \\ k^0_{12} & k^0_{22} \end{bmatrix} = \begin{bmatrix} 0.0277 & -0.1730 \\ -0.1730 & 2.2126 \end{bmatrix} \times 10^{11} \tag{6.10}$$

$$\mathbf{C}^0_{\text{macro}} = \left(\mathbf{K}^0_{\text{macro}}\right)^{-1} = \begin{bmatrix} c^0_{11} & c^0_{12} \\ c^0_{12} & c^0_{22} \end{bmatrix} = \begin{bmatrix} 0.7035 & 0.0550 \\ 0.0550 & 0.0088 \end{bmatrix} \times 10^{-9}, \tag{6.11}$$

where the units of k_{11}^0, k_{22}^0 and k_{12}^0 are N/m, Nm/rad and Nm/m, respectively, with corresponding compliance units of m/N, rad/Nm and m/Nm. These values were based on the initial foundation design derived from geotechnical site investigation data for the real-world wind turbine. This is discussed further in the following chapter (Section 7.3.1).

A 600 second forward simulation was run to generate synthetic data for the numerical case study, with a fixed sampling period of 0.04 seconds. The dynamics were driven by Kaimal wind and JONSWAP wave loading. The Kaimal wind load was defined by a mean wind speed of $10.59\,\text{m/s}$ and a turbulence intensity of $15\,\%$. Figure 6.2a shows the time history of the tower-top wind speed. From this wind speed, denoted U, the wind thrust force T and the aerodynamic torque Q were computed as:

$$T = \frac{1}{2} C_T \rho A U^2, \tag{6.12}$$

$$Q = \frac{1}{2} C_Q \rho A R U^2, \tag{6.13}$$

where C_T and C_Q are the thrust and torque coefficients, respectively. The rotor swept area and radius are denoted by A and R, and the density of air is taken as $\rho = 1\,\text{kg/m}^3$. For this numerical example, the thrust and torque coefficients were assumed to be $C_T = 0.804$ and $C_Q = 0.055$. The thrust force T was applied as a point force at the tower top node in the fore-aft direction, and the torque Q was applied about the fore-aft axis.

The wave kinematics were derived from a JONSWAP spectrum [165], defined by a significant wave height of $H_s = 3\,\text{m}$ and a peak wave period of $T_p = 10\,\text{s}$. The resulting wave surface elevation time history is shown in Figure 6.2b. The wave kinematics were used to compute the distributed forces along the submerged nodes according to the Morison Equation [164].

The response was taken from accelerometers and strain gauges distributed along the tower, as shown in Figure 6.1b, to be used as measurements for identifying the foundation parameters. These sensors were chosen to match the set-up of the real-world case study presented in Chapter 7, in which the turbine was instrumented with three biaxial accelerometers located approximately 33 %, 63 % and 93 % along the height of the substructure: a transition-piece accelerometer (ACC-TP), a tower-middle accelerometer (ACC-TM), and a tower-top accelerome-

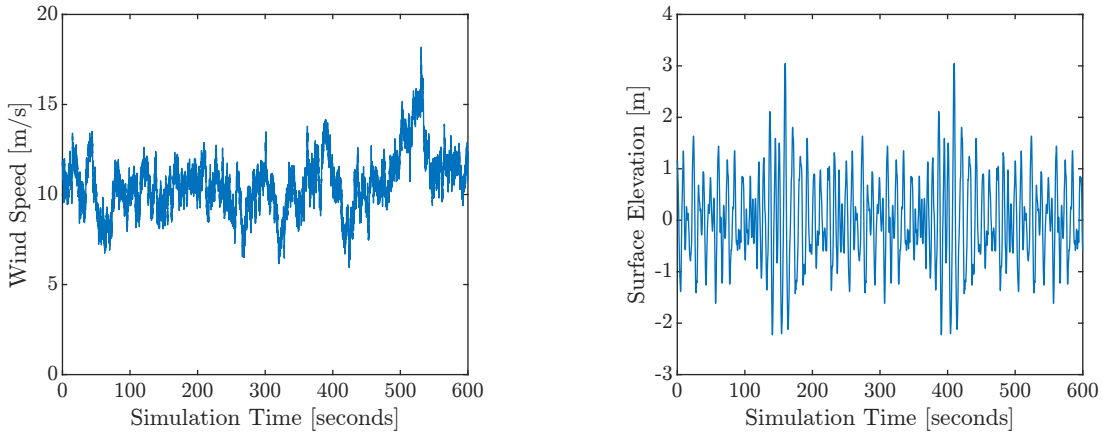


Figure 6.2: (a) Tower top wind speed defined by a Kaimal spectrum. (b) Wave surface elevation defined by a JONSWAP spectrum.

ter (ACC-TT). In this chapter, the biaxial accelerometers were aligned to measure the response in the fore-aft and side-side directions. In addition, six strain gauges were installed at the transition piece, distributed around the cross section as shown in Figure 6.3, and measuring strain along the vertical direction. For each signal, 1 % zero-mean white noise was added to introduce measurement error.

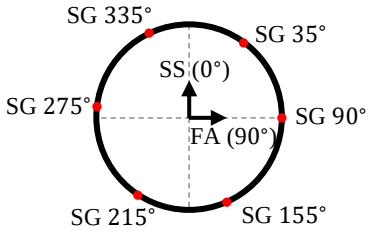


Figure 6.3: Distribution of strain gauges (SGs) at the transition piece. Abbreviations: FA = fore-aft, SS = side-side.

The lateral response was dominant, so the axial and torsional degrees of freedom (DOFs) were neglected. Therefore, the response here only captures the axial strain due to bending about the fore-aft and side-side axes.

6.3.2 Tower-Foundation Substructure Model

For identification of the foundation parameters, the substructure shown in Figure 6.1b was used. Only the foundation and the tower are modelled; in other words, it is identical to the model used to generate the data, with the added mass from the lumped rotor nacelle assembly

removed.

The dynamic system for the remaining substructure was then modelled using the second order equation of motion for a linear time-invariant system:

$$\mathbf{M}\ddot{\boldsymbol{u}}(t) + \mathbf{C}\dot{\boldsymbol{u}}(t) + \mathbf{K}(\boldsymbol{\theta})\boldsymbol{u}(t) = \mathbf{S}_p\boldsymbol{p}(t), \tag{6.14}$$

where $\mathbf{M}$, $\mathbf{C}$, and $\mathbf{K}(\boldsymbol{\theta})$ are the mass, damping and stiffness matrices, respectively and $\boldsymbol{u}(t)$ is the displacement vector at time t. $\mathbf{K}(\boldsymbol{\theta})$ depends on the unknown macro-element foundation parameters $\boldsymbol{\theta}$. $\boldsymbol{p}(t)$ is the vector of inputs which drive the system and $\mathbf{S}_p$ is a selection matrix which maps these inputs to the degrees-of-freedom at which they act. In this work, $\boldsymbol{p}(t)$ comprised two perpendicular, lateral point forces applied at the RNA-tower interface node.

Equation 6.14 was converted to a reduced-order representation to improve performance for input-state estimation and lower computational cost, by transforming from physical to modal coordinates using Equation 5.2. The number of modes in the reduction basis and degrees of freedom for the considered substructure are denoted as n_m and n_{dof}, respectively. The reduction basis $\boldsymbol{\Phi} \in \mathbb{R}^{n_{\text{dof}} \times n_m}$ is derived using a similar substructuring approach to that presented in Chapter 5, where it was shown to perform well for input-state estimation of an unconstrained substructure. In Chapter 5, the tower-monopile substructure was isolated from both the RNA and the soil-structure interaction. In contrast, the present chapter models the foundation explicitly, meaning the substructure is constrained. As a result, rigid body modes are no longer present, and only the tower-top degrees of freedom correspond to the interface with unmodelled components (i.e., the RNA). The resulting reduction basis then consists of only the retained free-interface normal modes, $\boldsymbol{\phi}_n$ and the attachment modes, $\boldsymbol{\phi}_a$:

$$\boldsymbol{\Phi} = \begin{bmatrix} \boldsymbol{\phi}_n & \boldsymbol{\phi}_a \end{bmatrix}. \tag{6.15}$$

The normal modes, $\boldsymbol{\phi}_n$, are derived following an eigendecomposition of the substructure's stiffness and mass matrices for a given realisation of $\boldsymbol{\theta}$ and a free-interface at the tower-top

interface:

$$\mathbf{K}\boldsymbol{\Psi} = \mathbf{M}\boldsymbol{\Psi}\boldsymbol{\Sigma}, \tag{6.16}$$

where $\boldsymbol{\Psi}$ and $\boldsymbol{\Sigma}$ are diagonal matrices containing the eigenvectors and associated eigenvalues, respectively. $\boldsymbol{\phi}_n$ is then defined as the retained subset of eigenvectors.

It should be noted that, since $\mathbf{K}$ depends on $\boldsymbol{\theta}$, so do $\boldsymbol{\Psi}$ and $\boldsymbol{\Sigma}$. Therefore, for each realisation of $\boldsymbol{\theta}$, this eigendecomposition must be repeated.

Computing the attachment modes $\boldsymbol{\phi}_a$ is simpler here compared to Chapter 5, since the substructure is constrained, meaning that the stiffness matrix $\mathbf{K}$ is not singular and can be inverted. Therefore, $\boldsymbol{\phi}_a$ is computed by solving the static equilibrium equation:

$$\boldsymbol{\phi}_a = \mathbf{K}^{-1} \begin{bmatrix} \mathbf{0} \\ \mathbf{I}_\mathrm{s} \end{bmatrix}, \tag{6.17}$$

where $\mathbf{I}_\mathrm{s} \in \mathbb{R}^{n_a \times n_a}$ is an identity matrix corresponding to the n_a interface degrees of freedom, and the block zero matrix has dimensions $(n_\mathrm{dof} - n_a) \times n_a$ to ensure compatibility.

The specific reduction basis $\boldsymbol{\Phi}$ used here comprised six retained normal modes, and the four attachment modes associated with the RNA-tower interface degrees of freedom. The six retained normal modes correspond to the first three pairs of bending modes. They are plotted in Figure 6.4 for the case when the baseline macro-element foundation, $\mathbf{K}^0_\mathrm{macro}$, was assumed.

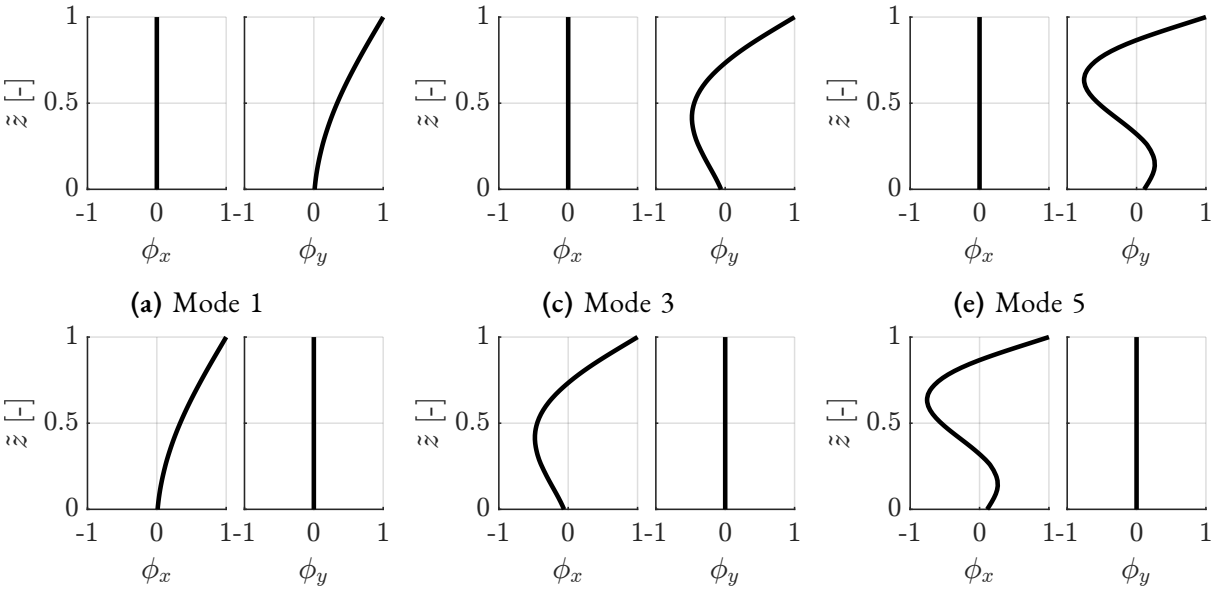
Equation 5.2 can be substituted into Equation 6.14 to derive the reduced-order equation of motion:

$$\tilde{\mathbf{M}}\ddot{\boldsymbol{z}}(t) + \tilde{\mathbf{C}}\dot{\boldsymbol{z}}(t) + \tilde{\mathbf{K}}\boldsymbol{z}(t) = \boldsymbol{\Phi}^T\mathbf{S}_p\boldsymbol{p}(t), \tag{6.18}$$

where $\boldsymbol{z}(t) \in \mathbb{R}^{n_m \times 1}$ is the displacement vector in modal coordinates and, $\tilde{\mathbf{M}} = \boldsymbol{\Phi}^T\mathbf{M}\boldsymbol{\Phi}$, $\tilde{\mathbf{C}} = \boldsymbol{\Phi}^T\mathbf{C}\boldsymbol{\Phi}$ and $\tilde{\mathbf{K}} = \boldsymbol{\Phi}^T\mathbf{K}\boldsymbol{\Phi}$.

The state vector $\boldsymbol{x}(t) = \begin{bmatrix} \boldsymbol{z}(t) & \dot{\boldsymbol{z}}(t) \end{bmatrix}^T$ is then introduced to transform the second order equation of motion to a continuous state space representation:

$$\dot{\boldsymbol{x}}(t) = \mathbf{A}_c\boldsymbol{x}(t) + \mathbf{B}_c\boldsymbol{p}(t) \tag{6.19}$$



(a) Mode 1 **(c)** Mode 3 **(e)** Mode 5

(b) Mode 2 **(d)** Mode 4 **(f)** Mode 6

Figure 6.4: Example of retained normal modes computed using $\mathbf{K}^0_{\text{macro}}$. $\tilde{z}$ denotes the normalised height, where $\tilde{z} = 0$ corresponds to the mudline and $\tilde{z} = 1$ to the tower top. x and y are aligned with the fore-aft and side-side directions, respectively.

$$\boldsymbol{y}(t) = \mathbf{G}\mathbf{x}(t) + \mathbf{J}\boldsymbol{p}(t), \tag{6.20}$$

where:

$$\mathbf{A}_c = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\tilde{\mathbf{M}}^{-1}\tilde{\mathbf{K}} & -\tilde{\mathbf{M}}^{-1}\tilde{\mathbf{C}} \end{bmatrix}, \quad \mathbf{B}_c = \begin{bmatrix} \mathbf{0} \\ \tilde{\mathbf{M}}^{-1}\boldsymbol{\Phi}^T\mathbf{S}_p \end{bmatrix}. \tag{6.21}$$

The measurement vector $\boldsymbol{y}(t)$ corresponds to the response from the installed accelerometers and strain gauges. Therefore, the matrices $\mathbf{G}$ and $\mathbf{J}$ take the following form:

$$\mathbf{G} = \begin{bmatrix} -\mathbf{S}_a\boldsymbol{\Phi}\tilde{\mathbf{M}}^{-1}\tilde{\mathbf{K}} & -\mathbf{S}_a\boldsymbol{\Phi}\tilde{\mathbf{M}}^{-1}\tilde{\mathbf{C}} \\ \mathbf{T}_\varepsilon\mathbf{S}_\varepsilon\boldsymbol{\Phi} & \mathbf{0} \end{bmatrix}, \quad \mathbf{J} = \begin{bmatrix} \mathbf{S}_a\boldsymbol{\Phi}\tilde{\mathbf{M}}^{-1}\boldsymbol{\Phi}^T\mathbf{S}_p \\ \mathbf{0} \end{bmatrix}, \tag{6.22}$$

where $\mathbf{S}_a$ is the selection matrix to extract accelerations corresponding to the measured accelerometers' DOFs, and $\mathbf{S}_\varepsilon$ is a selection matrix that extracts the displacements associated with the element where strain is measured. The matrix $\mathbf{T}_\varepsilon$ transforms these displacements into strains using the element-level stiffness matrix associated with the measured strain DOFs. Therefore, $\mathbf{T}_\varepsilon$ implicitly captures the procedure of computing internal forces (e.g., bending moments) from a stiffness matrix and subsequently converting them into strain at specific points in the cross

section using beam theory.

It should be noted that both the accelerations and strain measurements were used for input-state estimation, as implied by $\mathbf{G}$ and $\mathbf{J}$ in Equation 6.22. However, for defining the likelihood function (Equations 6.3–6.7), only the strain measurements were used in the identification of the foundation parameters—i.e., only the bottom block rows of $\mathbf{G}$ and $\mathbf{J}$ were considered.

Finally, a zero-order hold discretisation is applied to the continuous system matrices $\mathbf{A}_c$ and $\mathbf{B}_c$, with sampling frequency f_s:

$$\mathbf{A}_d = \exp\left(\mathbf{A}_c / f_s\right) \quad \text{and} \quad \mathbf{B}_d = \mathbf{A}_c^{-1} \left[\mathbf{A}_d - \mathbf{I}\right] \mathbf{B}. \tag{6.23}$$

The discrete state-space matrices $\mathbf{A}_d$, $\mathbf{B}_d$, $\mathbf{G}$ and $\mathbf{J}$ are then assembled into the augmented state and output matrices, $\mathbf{A}_d^a$ and $\mathbf{G}^a$ (see Equation 5.22), which can be used to define the likelihood function derived in Section 6.2.

The process and measurement noise matrices, $\mathbf{Q}^a$ and $\mathbf{R}$, were kept constant for all realisations of $\boldsymbol{\theta}$. They were both assumed to be diagonal matrices, where the elements in $\mathbf{R}$ were set as as 0.01 times the variance of each measurement signal. The process noise matrix $\mathbf{Q}^a$ was defined using three manually-tuned covariances associated with $\boldsymbol{z}$, $\dot{\boldsymbol{z}}$ and $\boldsymbol{p}$. Through trial and error, values of 10^{-5} for the states $\boldsymbol{z}$ and $\dot{\boldsymbol{z}}$, and 10^{10} for $\boldsymbol{p}$, were found to perform well as the diagonal entries of $\mathbf{Q}^a$.

6.3.3 Identification of the foundation parameters

Two parameterisation of the macro-element foundation were considered to be identified from the measured data. The first involves identifying a separate value for each compliance parameter of the macro-element defined in Equation 3.23: $\boldsymbol{\theta}_1 = \{c_{11},\, c_{22},\, c_{12}\}$.

The second parameterisation uses the initial design macro-element, $\mathbf{C}^0_{\text{macro}}$, to constrain the relationship between the springs. Therefore, this parameterisation is more robust when the soil parameters are correlated [197, 198], and allows for a single parameter, γ, to scale the entire

macro-element as:

$$\mathbf{C}^0_{\text{macro}} = \gamma \left(\mathbf{K}^0_{\text{macro}}\right)^{-1} = \gamma \begin{bmatrix} c^0_{11} & c^0_{12} \\ c^0_{12} & c^0_{22} \end{bmatrix}. \tag{6.24}$$

Uniform priors were assumed for all parameters, defined by the bounds listed in Table 6.1. The burn in time and batch length required to formulate the likelihood function (Equation 6.7) were 200 seconds and 600 seconds, corresponding to $N_{\text{burn}} = 5000$ and $N_{\text{batch}} = 15000$, respectively. To obtain the posterior distributions, a brute-force approach was employed, where random samples were drawn from the prior distribution using Latin hypercube sampling and evaluated using Equation 6.7.

Table 6.1: Lower and upper bounds used to define the uniform prior distributions for all parameter. The initial design compliance values (c^0_{11}, c^0_{22}, and c^0_{12}) are given in Equation 6.11.

Parameter	Lower Bound	Upper Bound
c_{11}	$0.1 \times c^0_{11}$	$100 \times c^0_{11}$
c_{11}	$0.1 \times c^0_{22}$	$5 \times c^0_{22}$
c_{11}	$0.1 \times c^0_{12}$	$5 \times c^0_{12}$
γ	0.1	5

Figure 6.5 presents the posterior distribution for the compliance parameterisation, where the parameters have been normalised by their true values. The corresponding errors in the *maximum a posteriori* (MAP) estimates were: $-10.18\,\%$, $-6.84\,\%$, and $31.01\,\%$ for c_{11}, c_{22}, and c_{12}, respectively. These results show that only the rotational compliance c_{22} can be reliably identified from these measurements. For the other parameters (c_{11} and c_{12}), the posterior distributions contains flat regions, indicating that these parameters are practically unidentifiable. Moreover, due to the presence of measurement error, c_{11} and c_{12} compensate for one another, resulting in larger errors in the MAP estimates. The lack of identifiability observed here is likely due to the limited sensitivity of the predicted measurements to these parameters. However, for very high values of c_{11}, the posterior probability starts to decrease. This suggests that while the measurements indicate the lateral spring is stiff, the data cannot resolve exactly how stiff it is. To a point this behaviour can be expected for foundations with very large lateral stiffness: whether they are modelled with the exact value of a very high stiffness, or another

similarly large value, or as perfectly rigid constraints, the differences would be very small to tell apart. Therefore, using this framework, the data is only informative enough to rule out the possibility of the foundation being soft. Nonetheless, unlike the results from Chapter 4, which showed cases where the posterior distributions were confidently incorrect, the framework considered here avoids overconfidence and instead highlights the practical unidentifiability of certain parameters—a result that can still be valuable to the user.

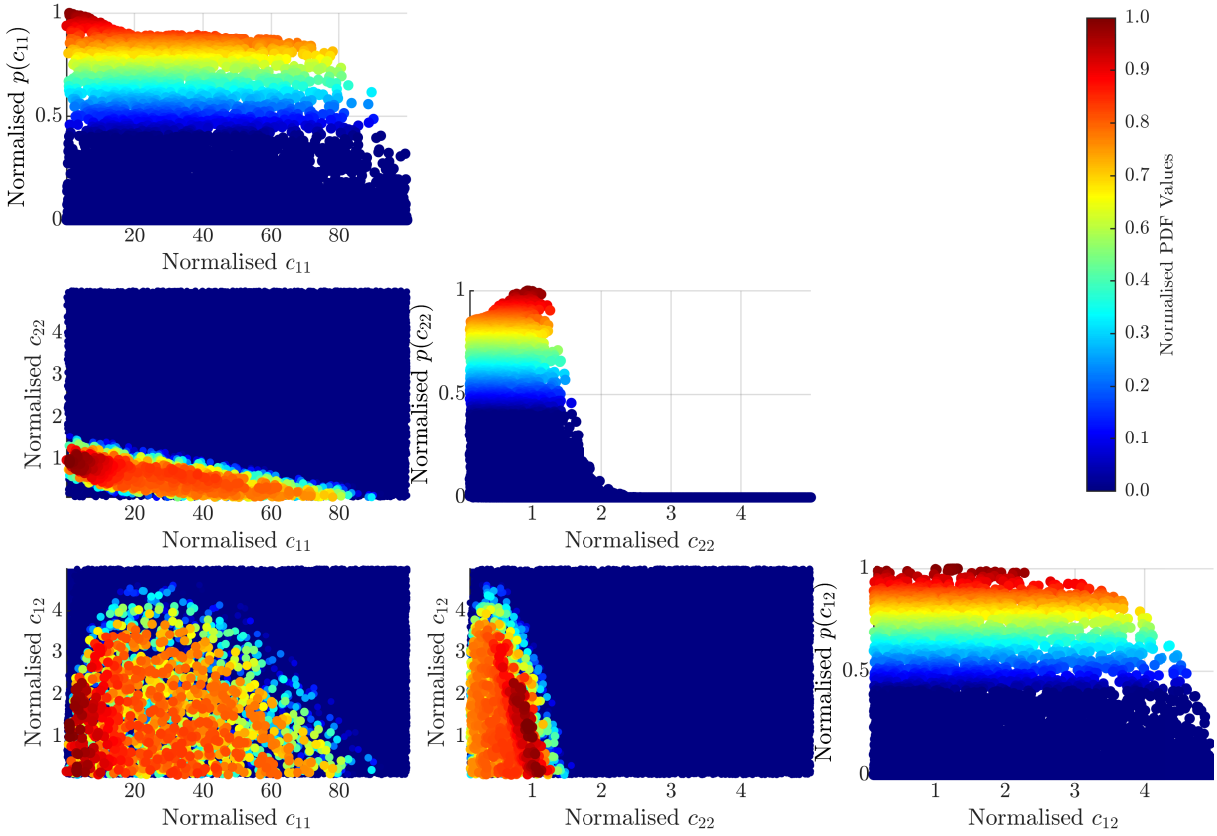


Figure 6.5: Posterior distribution for compliance parameterisation. Parameters are normalised by their true values. Probability density functions are normalised by the maximum value.

To allow for all three macro-element foundation parameters to be identified, additional information must be introduced to constrain the relationship between c_{22}, c_{11} and c_{12}, thereby reducing the number of unknowns. This can be provided by using the second parameterisation in terms of γ. Figure 6.6 shows the posterior distribution for γ, with a MAP estimate error of $-1.89\,\%$. The well-defined peak and low estimation error indicate that γ is identifiability.

Figure 6.6 was obtained by assuming a baseline macro-element foundation in which the relative differences in each compliance were known exactly. To investigate the impact of using

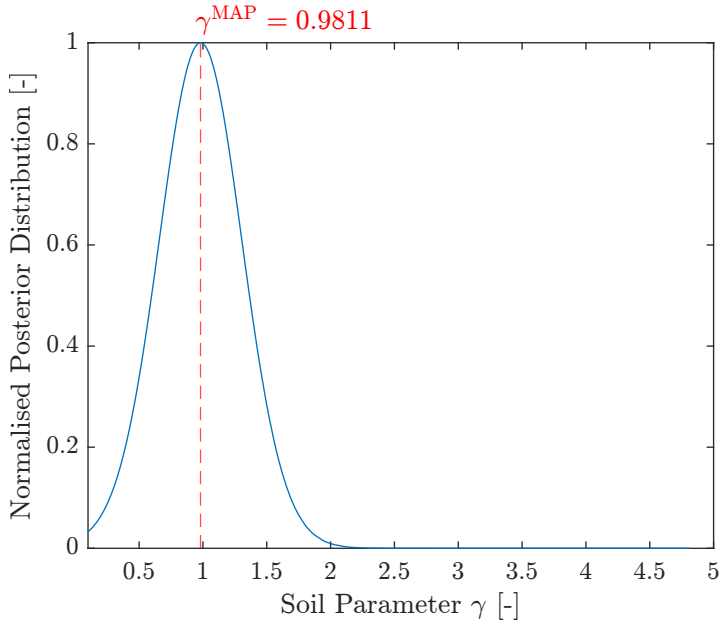


Figure 6.6: Posterior distribution for γ. The probability density function is normalised by its maximum value.

an incorrect baseline foundation, either a 5 % or 10 % error was introduced to each compliance parameter individually, and γ was identified as before. Table 6.2 shows the MAP estimates of γ for each case. The results show that γ^{MAP} is relatively insensitive to errors in c^0_{11} and c^0_{12}. However, when an error is present for c^0_{22}, the framework compensates by adjusting γ^{MAP} to match c^0_{22}, at the expense of c_{11} and c_{12}. This indicates that, for this case study, the results are dominated by the rotational spring c_{22}, which is consistent with the results obtained when identifying each parameter individually (Figure 6.5).

Table 6.2: *Maxima a Posteriori* (MAP) estimates for γ for different assumed errors in the baseline compliance matrix $\mathbf{C}^0_{\text{macro}}$.

Error in $\mathbf{C}^0_{\text{macro}}$			**MAP Estimate γ^{MAP} [-]**
c^0_{11} error	c^0_{22} error	c^0_{12} error	
0 %	0 %	0 %	0.9811
5 %	0 %	0 %	0.9805
10 %	0 %	0 %	0.9801
0 %	5 %	0 %	0.9408
0 %	10 %	0 %	0.9037
0 %	0 %	5 %	0.9747
0 %	0 %	10 %	0.9683

(i) Influence of underlying foundation stiffness

The previous results were obtained using data from a forward simulation where the model was defined by the baseline macro-element foundation $\mathbf{K}^0_{\text{macro}}$. However, the practical identifiability of γ may depend on the relative stiffness of the underlying foundation. To investigate this, additional forward simulations were carried out to obtain measurement data for foundations with stiffnesses equal to 0.5, 2, 4, and 10 times $\mathbf{K}^0_{\text{macro}}$, and γ was identified for each case.

Figure 6.7 presents the estimated posterior probability distributions for γ using data obtained from simulations with the different foundations. The results show that as the underlying foundation stiffness increases, the posterior distributions widen, reflecting reduced identifiability of γ. This observation occurs because the measurements become less sensitive to changes in γ as the foundation increases in stiffness. This supports the discussion from the previous section: when the foundation springs become very stiff it becomes difficult to differentiate their effect from a rigid constraint.

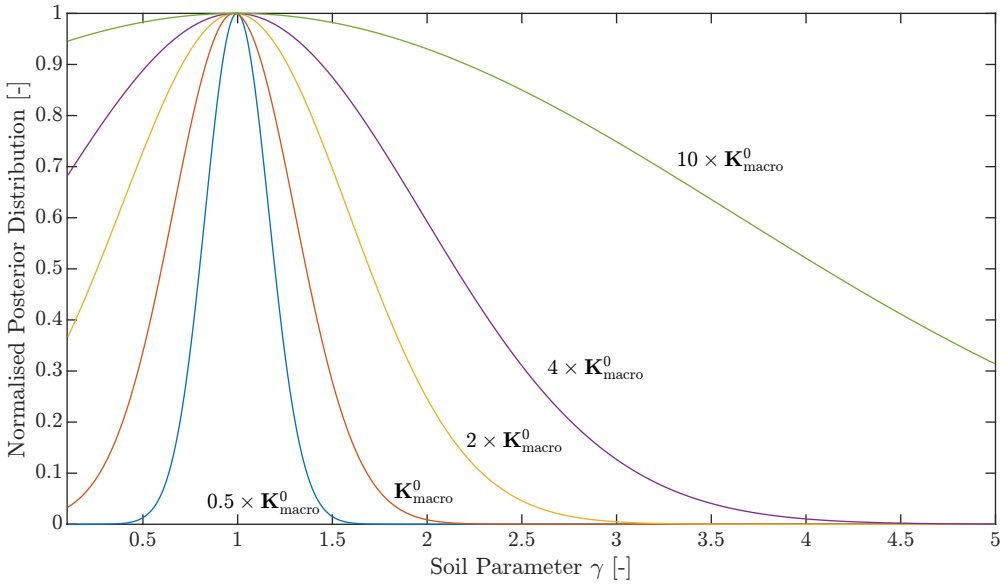


Figure 6.7: Influence of the underlying foundation stiffness on the identifiability of γ.

The corresponding errors in the MAP estimates are provided in Table 6.3, which shows that the most probable estimates for γ were accurately estimated for all cases, with no clear trend linked to underlying foundation stiffness. However, in practical applications—where modelling error may also be present—this cannot be guaranteed, and the reduced identifiability of γ in stiffer foundations may lead to larger errors in the MAP estimates.

Table 6.3: Influence of the underlying foundation stiffness on the *Maximum a Posteriori* (MAP) estimates for γ.

Underlying Foundation Stiffness	Error in MAP Estimate for γ
$0.5 \times \mathbf{K}^0_{\text{macro}}$	$-0.40\,\%$
$1 \times \mathbf{K}^0_{\text{macro}}$	$-1.99\,\%$
$2 \times \mathbf{K}^0_{\text{macro}}$	$-1.40\,\%$
$4 \times \mathbf{K}^0_{\text{macro}}$	$-1.19\,\%$
$10 \times \mathbf{K}^0_{\text{macro}}$	$-0.42\,\%$

(ii) Influence of model error from water levels and added mass coefficient

The results from the numerical examples shown so far have demonstrated that γ can be identified using the proposed substructuring approach which avoids modelling the rotor nacelle assembly, thereby bypassing significant modelling errors associated with this component. However, the remaining substructure model was assumed to be exact. In practice, uncertainties may exist in the parameters governing the added mass of submerged elements—the internal water level ζ_i, external water level ζ_e, and the hydrodynamic added mass coefficient C_a (see Equations 6.8 and 6.9). It is therefore important to consider how errors in these parameters could affect the identifiability of γ. To investigate this, forwards simulations were run under different parameterisations ζ_i, ζ_e, and C_a, while retaining the same substructure model as previously introduced for identification.

For data generation in this section, five values of C_a were considered, varying from 0.5 to 2. The water levels were differed by ± 4 metres relative to the mean water level to reflect possible tidal variations. Two cases were considered for the water levels. In the first, only the external water level was varied and the internal water level kept fixed at the mean water level. In the second, both the external and internal water levels were varied. This resulted in 90 combinations arising from variations in ζ_i, ζ_e and C_a. Note also that data was generated by applying identical loading and using the baseline macro-element foundation $\mathbf{K}^0_{\text{macro}}$ for all cases. The substructure model used for identification was defined by using the 'incorrect' values of $C_a = 1$, with ζ_i and ζ_e set to the mean water level.

Figure 6.8 presents the identified MAP values for γ and the corresponding 95 % confidence

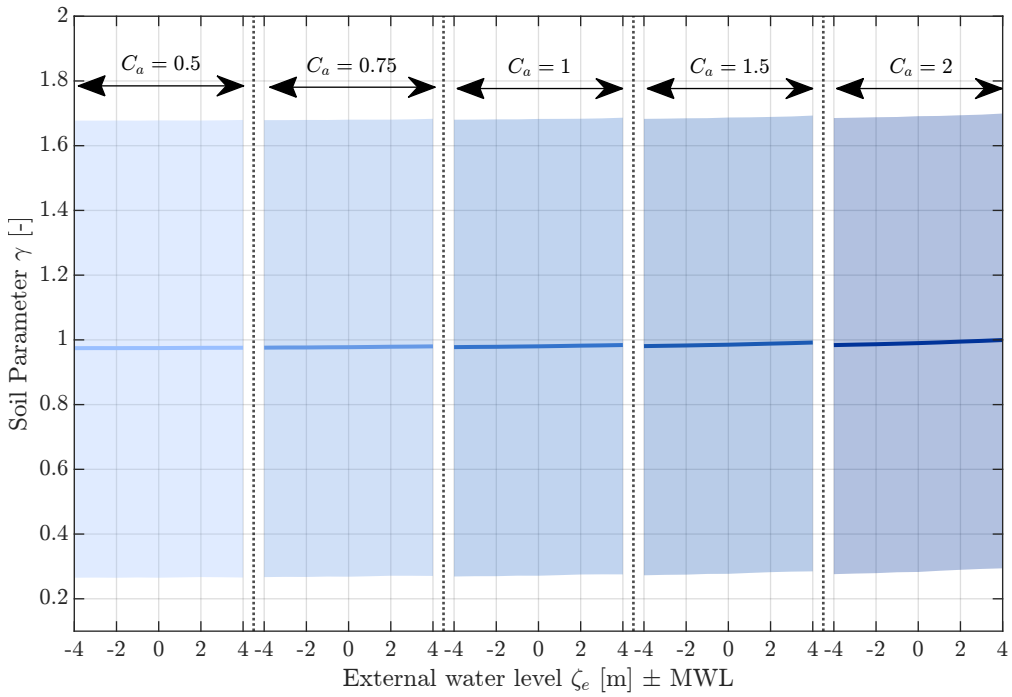
interval, across all parameterisations. The results show that the estimated posterior distributions are very consistent, indicating that γ is insensitive to variations in ζ_i, ζ_e or C_a. A minor trend can be observed: as any of these parameters increase, the MAP estimates for γ also increase and the distributions shift towards higher values. This result was expected, since larger values of ζ_i, ζ_e or C_a increase the added mass of submerged elements, thereby lowering natural frequencies present in the data. In response, higher values of γ are identified to compensate for the corresponding reduction in dynamic stiffness.

In addition, the results indicate that as the added mass of the submerged elements increase—i.e., ζ_i, ζ_e or C_a increase—their influence on the identified distributions for γ also increases. For example, the left hand side of Figure 6.8a exhibits a flatter gradient compared to the right hand side of Figure 6.8b. Therefore, as ζ_i, ζ_e or C_a increase, so does their influence on the measured response, leading to identified distributions for γ that are more sensitive to errors in these parameters. Nevertheless, over the range considered here, ζ_i, ζ_e or C_a do not significantly affect the identifiability of γ to a degree that would necessitate corrective measures. However, their influence may be greater in other settings—such as shorter wind turbines or deeper waters—and should therefore be assessed for the specific structure under consideration. If this sensitivity is found to be significant, then these parameters could be considered as uncertain and identified alongside γ. However, this was considered outside the scope of the current study.

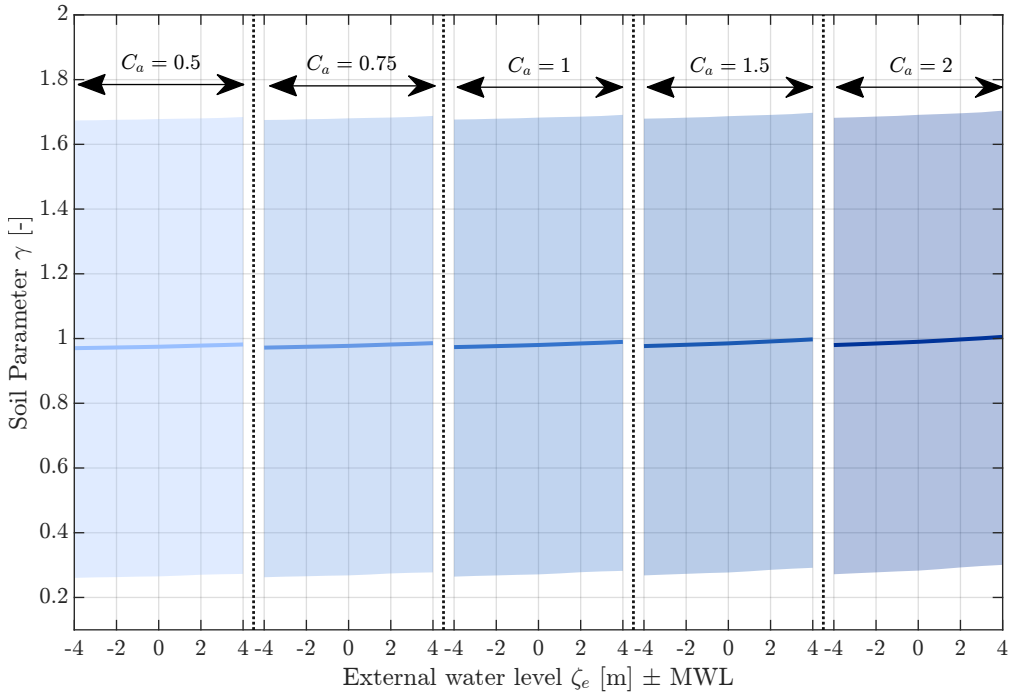
6.4 Summary

This chapter has presented a numerical case study to evaluate the potential of a substructuring-based framework for identifying the foundation parameters of an offshore wind turbine. The results indicate that the turbine's dynamic behaviour is sensitive only to the rotational spring. For the relative soil-structure stiffness considered here, the available measurements show minimal sensitivity to the lateral spring, except at very high compliance values (normalised $c_{11} > 80$) where it begins to influence the response.

The identifiability of the foundation parameters is dependent on the underlying foundation stiffness: softer foundations lead to more identifiable parameters, likely because the system dynamics are more sensitive to the foundation behaviour. Additionally, the identifiability of



(a)



(b)

Figure 6.8: Posterior estimates for γ obtained from data with different parameterisations of the internal water level ζ_i, external water level ζ_e and hydrodynamic added mass coefficient C_a. The solid line indicates the *Maximum a Posteriori* estimates and the shaded region represents the 95 % confidence interval. (a) Results when ζ_e is varied and ζ_i is fixed equal to the mean water level. (b) Results when $\zeta_i = \zeta_e$.

γ was found to be robust to uncertainties in the hydrodynamic added mass coefficient and variations in the internal and external water levels.

Overall, the results demonstrate that the proposed framework can reliably identify foundation parameters while avoiding the modelling challenges associated with the rotor nacelle assembly highlighted in Chapter 4. Furthermore, this numerical case study shows that the method is applicable in realistic conditions using limited accelerometers and strain gauges, and provides a strong basis for the real-world case study presented in the following chapter.

CHAPTER 7

Real-World Case Study

7.1 Introduction

The primary aim of this chapter is to apply the structural health monitoring methods developed in earlier chapters to identify the foundation parameters of a real-world offshore wind turbine. In doing so, this chapter evaluates their practical applicability and compares the results with insights gained from the numerical case studies presented in Chapters 3–6. This comparison helps to reinforce the conclusions drawn previously and to increase confidence in the results obtained for this real-world case study. Two Bayesian inference methods are applied for parameter estimation and are summarised below. Their mathematical formulations are not repeated here as they have already been introduced.

Method 1: Bayesian Model Updating using Modal Data

Method 1 applies the vector-based Bayesian model updating framework introduced in Section 3.3.2. This method relies on modal data identified from operational modal analysis and is widely used in the literature. The method requires modelling the entire structure, including the rotor nacelle assembly (RNA). However, accurately modelling the flexibility of the blades requires detailed design information, which were not available here due to commercial confidentiality. As a result, a flexible RNA model was implemented by scaling blade properties from the IEA-10 MW reference turbine [9], based on the known length and lumped mass of the actual blades. However, the true distributions of mass, stiffness, twist profiles, and aerofoil properties are likely to differ significantly, introducing model errors that may affect the accuracy of the identified foundation parameters. Therefore, a second independent method is also considered to provide cross validation of the results.

Method 2: Substructuring-based approach for parameter estimation

Method 2 applies the substructuring-based parameter estimation framework introduced in Chapter 6. In this approach, the rotor nacelle assembly is excluded from the model, and input-state estimation is performed using an augmented Kalman filter applied to the reduced tower-foundation system. The results are subsequently employed to formulate a probabilistic model to estimate the foundation parameters. This framework is shown to be better suited to scenarios where there is significant RNA model error. Moreover, the input-state estimation results from the updated substructure model can be leveraged for virtual sensing applications, such as fatigue estimation.

An additional contribution of this chapter lies in the processing of a real-world dataset that has not previously been analysed. Therefore, Section 7.2 is assigned to describing the wind turbine setup and the post-processing of the field data. Operational modal analysis results are presented in Section 7.2.1, where an automated process for selecting suitable modes across multiple datasets is described. This automated process follows an approach similar to that proposed by Reynders *et al.* [199].

Following this, Section 7.3 provides an overview of the numerical models used for each method, and Section 7.4 presents the results from applying the two frameworks to the real-world case study.

7.2 Field Data

The performance of the two methods for parameter identification is evaluated using field data from one of RWE's large-scale offshore wind turbines, located in the North Sea. The wind turbine is supported by a monopile foundation embedded in a medium-dense sand.

Figure 7.1 provides a schematic of the wind turbine, showing its main dimensions and the layout of the sensors. The structure is instrumented with three biaxial accelerometers measuring horizontal accelerations which are positioned at the transition piece (TP), tower-middle (TM), and tower-top (TT), as well as six strain gauges installed on the transition piece and measuring strains along the vertical direction. The TP, TM, and TT sensors are located at approximately 33 %, 63 %, and 93 % of the distance from ground level to hub height, respectively. Both the

accelerometers and strain gauges sample data at a frequency of 25 Hz.

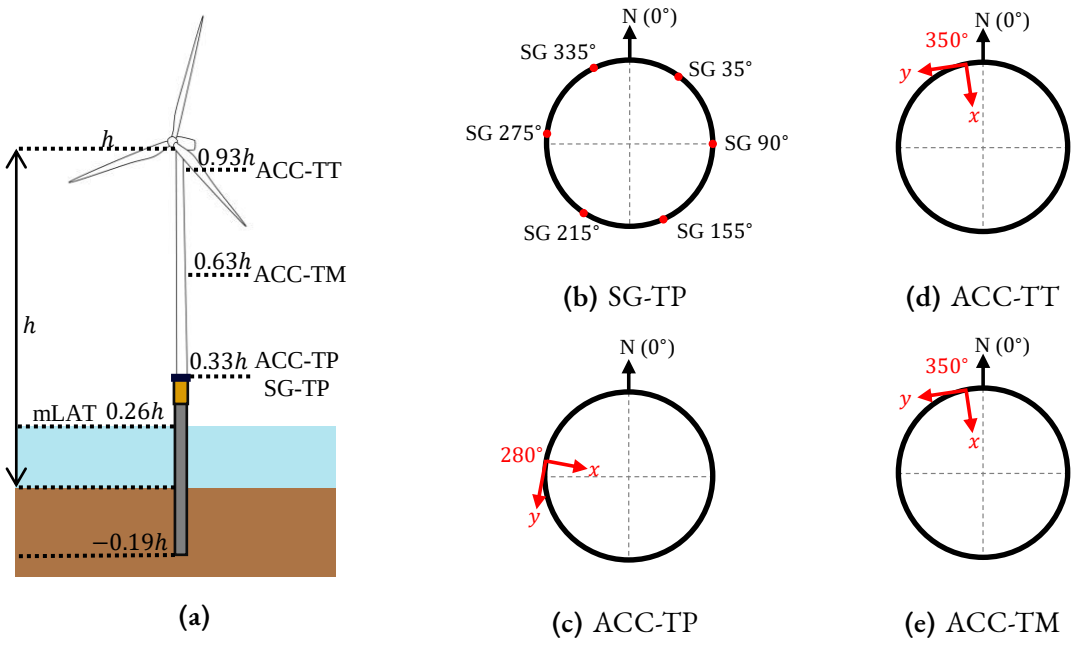


Figure 7.1: (a) Wind turbine geometry and tower sensor positions. (b) Transition piece strain gauges: SG-TP (c) Transition piece accelerometer: ACC-TP (d) Tower top accelerometer: ACC-TT (e) Tower-middle accelerometer: ACC-TM

The accelerometer data were high-pass filtered above 0.05 Hz to remove DC offsets. Strain gauge data were only required for Method 2, which uses a model that neglects axial degrees of freedom (see Section 6.3.2). Therefore, the strain component due to the axial force of the cross-section needed to be removed from the measured signals. To achieve this, it was assumed that the stresses are described by beam theory and hence, the distribution of strain around the inner diameter of the cross section varies as:

$$\varepsilon(\alpha) = X + Y\sin(\alpha) + Z\cos(\alpha), \tag{7.1}$$

where X is a constant, independent of the angular position α (measured clockwise from 0°, in accordance with Figure 7.1), which represents the component of strain due to the axial force. The constants Y and Z correspond to the components of strain due to the bending moments. Using the six strain gauge signals, linear regression was applied to estimate X, Y and Z, after which the axial component X was subtracted from each measured strain signal.

In addition to the vibration data, the supervisory control and data acquisition (SCADA) sys-

tem records operational data such as wind speed, nacelle yaw and blade pitch. These parameters are logged as ten-minute averages, along with their corresponding standard deviations.

In this work, SCADA data was only required for Method 1. The nacelle yaw was used to transform acceleration signals from their local x-y coordinate system to a fore-aft (FA) and side-side (SS) coordinate system. Consequently, the yaw angle can influence the identified modal parameters, which is discussed further in Section 7.2.1. In addition, the Simulink model—used for Method 1 only—must be linearised about a steady-state operating point to compute modal parameters. Therefore, the blade pitch and wind speed were used as parameters for a steady-state simulation to define this operating point.

By contrast, Method 2 does not model the RNA or require a steady-state simulation. Therefore, it does not depend on data from the SCADA system. For Method 2, the acceleration signals are instead transformed to a consistent North–East coordinate system based on their sensor installation orientation (see Figure 7.1).

A full day of data was available for analysis. The data was stored in successive segments of ten minutes each. During this specific day, the wind turbine was found to be continuously idling for a 7-hour and 30-minute period. As the methods and modelling assumptions used in this chapter (e.g., linear-elastic springs in the macro-element foundation) are better suited to idling conditions, only this idling period was considered for parameter identification, resulting in a total of 45 ten-minute datasets.

7.2.1 Operational Modal Analysis

Since Method 1 requires modal data, the accelerometer data were used to estimate the structure's modal parameters for each ten-minute dataset during the idling period. Figure 7.2 presents the nacelle yaw angle for the continuously idling window used in this work. The nacelle yaw exhibits some variation between datasets but generally stays within a relatively narrow range. However, the standard deviation associated with the nacelle yaw—represented by the shaded area in Figure 7.2—indicates that the nacelle yaw was not stationary for most of the idling datasets. This is important to note because the yaw angle is assumed constant over the corresponding ten-minute segment when transforming raw data from the accelerometers'

local x-y coordinate system to the FA–SS directions; any variability in the yaw angle will therefore introduce measurement error. While splitting each ten-minute segment into shorter time windows could be investigated to reduce this influence on the identified modal parameters, this is left future work.

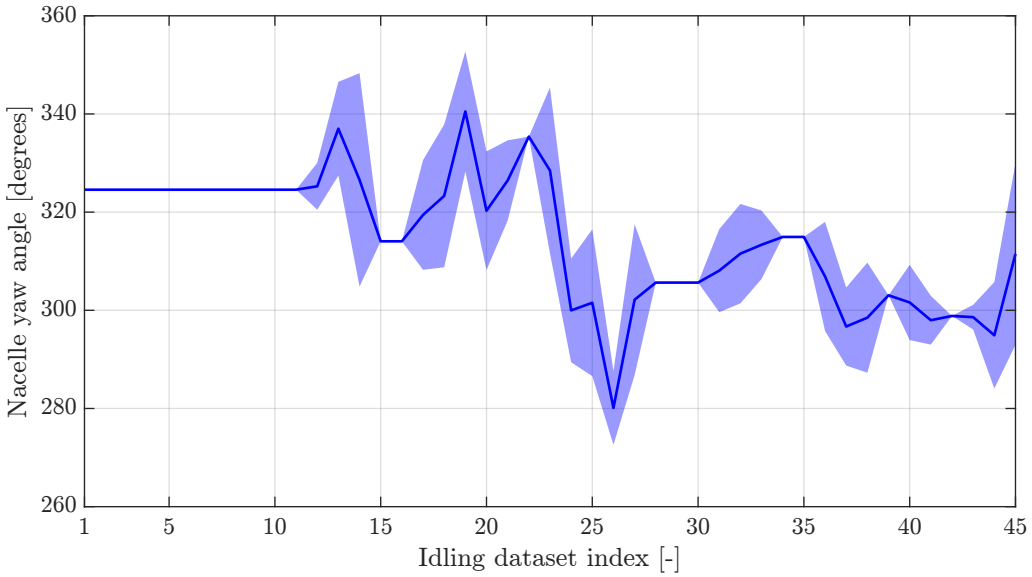


Figure 7.2: Ten-minute averages and corresponding ± 2 standard deviations of nacelle yaw angles from the SCADA system.

Once transformed to the FA–SS coordinate system, the six acceleration signals were used to estimate the structure's natural frequencies and mode shapes by applying the stochastic subspace algorithm SSI-cov [200, 201]. This is the same algorithm used in Chapter 3 (see Section 3.2), which was also implemented using the MACEC toolbox [171] in this chapter. For a given model order, n_{mo}, the algorithm estimates the system's eigenvalues $\hat{\boldsymbol{\lambda}} \in \mathbb{R}^{n_{\text{mo}}}$, and mode shapes in the sensor basis $\hat{\boldsymbol{\Phi}} \in \mathbb{R}^{n_y \times n_{\text{mo}}}$, where n_y is number of measurements. In this chapter, $n_y = 6$, corresponding to the number of acceleration signals.

The ith identified eigenvalue in $\hat{\boldsymbol{\lambda}}$ is defined as:

$$\begin{aligned} \hat{\lambda}_i &= -2\pi \hat{f}_i \hat{\zeta}_i \pm 2\pi \hat{f}_i \sqrt{1-\hat{\zeta}_i^2}, \\ \hat{f}_i &= \frac{|\hat{\lambda}_i|}{2\pi}, \end{aligned} \tag{7.2}$$

where $\hat{f}_i$ and $\hat{\zeta}_i$ are the ith natural frequency in Hz and damping ratio respectively.

For each ten-minute dataset, SSI-cov was applied for n_{mo} ranging from 50 to 220. This enabled the implementation of a semi-automated procedure to discard spurious modes and select a subset of physical modes most suitable for model updating. The procedure is described next, where Stages 1 & 2 are largely based on the fully-automated operational modal analysis approach proposed by Reynders *et al.* [199], with some modifications.

Stage 1: Initial Clearing

The aim of this stage is to clear out any modes that are thought to be certainly spurious and separate them from those that are possibly physical. This is achieved using the following single-mode validation criteria.

These include distance measures for eigenvalues, eigenfrequencies, and damping ratios, defined as follows:

$$d\lambda = \frac{|\lambda_1 - \lambda_2|}{\max(|\lambda_1|, |\lambda_2|)}; \quad df = \frac{|f_1 - f_2|}{\max(f_1, f_2)}; \quad d\zeta = \frac{|\zeta_1 - \zeta_2|}{\max(|\zeta_1|, |\zeta_2|)}. \tag{7.3}$$

which take values of zero in the ideal case of comparing two of the same physical mode.

To evaluate the similarity of mode shapes, the modal assurance criterion (MAC) is used:

$$\text{MAC} = \frac{|\boldsymbol{\phi}_1 \boldsymbol{\phi}_2^*|^2}{\boldsymbol{\phi}_1 \boldsymbol{\phi}_1^* \boldsymbol{\phi}_2 \boldsymbol{\phi}_2^*}, \tag{7.4}$$

where $(\cdot)^*$ indicates the (complex conjugate) transpose of a vector. In the ideal case, the MAC equals 1 when comparing two identical physical modes or takes a low value when comparing unrelated modes.

For each mode at a given model order, df, $d\zeta$ and the MAC are computed with respect to the mode at the nearest lower model order that minimises $d\lambda$. These criteria (df, $d\zeta$ and MAC) are then used as input features for a k-means clustering algorithm [202], with $k = 2$. The result is two clusters: one containing modes classified as possibly physical and the other containing modes classified as certainly spurious.

Hard criteria were then applied to the subset of possibly physical modes to eliminate

some of the spurious modes within this cluster. In this work, modes with damping ratios outside the range of 0% and 20% were discarded. Additionally, the estimated variance of the eigenfrequencies, as computed from the SSI-cov method [60, 61], was used to define the hard criterion that all possibly physical modes must have a coefficient of variation less than or equal to 0.1.

Stage 2: Hierarchical Clustering of Similar Modes

In this stage, modes classified as possibly physical in Stage 1 are clustered based on their similarity using an automated approach. Initially, each possibly physical mode is treated as its own individual cluster. The similarity between clusters is then quantified through a distance metric, which defines the distance between two clusters k and l as:

$$d(k,l) = \frac{|f_k - f_l|}{\max(f_k, f_l)} + 1 - \text{MAC}(\boldsymbol{\phi}_k.\boldsymbol{\phi}_l). \tag{7.5}$$

The minimum distance between any two clusters is compared to a predefined threshold d_{th}, given by:

$$d_{th} = \mu_p + \sigma_p \tag{7.6}$$

where μ_p and σ_p are the mean and standard deviation of the combined feature $p = df + (1 - \text{MAC})$, computed using the df and MAC values from Stage 1 across all possibly physical modes. If the minimum distance between any two clusters is less than d_{th}, then those clusters are merged and the distances between all remaining clusters are updated by computing the average pairwise distance between their constituent elements. This process is repeated iteratively until the minimum distance between any two clusters exceeds d_{th}.

Stage 3: Cleaning of Clusters

The distance threshold d_{th} defined by Equation 7.6 was was found to be problematic: it sometimes caused distinct physical modes to be incorrectly merged into a single cluster, but also prevented the merging of clusters that actually correspond to the same physical mode. In other words, the threshold was both too high and too low. To address these issues, a cleaning stage was introduced.

To ensure that clusters contain distinct physical modes, it was assumed that no cluster should include more than one mode from the same model order. If a cluster violated this condition, it was split using k-means clustering, where k was set to the maximum number of occurrences of any single model order within that cluster.

For each resulting cluster, a reference mode was selected as the mode corresponding to the median damping ratio for that cluster. The clusters were then ordered by the number of modes they contained, from largest to smallest. Starting from the largest cluster, the lowest model order present in that cluster was identified. For all modes at that model order—including those previously discarded as spurious—the mode that minimised the following objective function J was assigned to a new set of clusters:

$$J = \left(\frac{|f_n^{\text{mo}} - f_{\text{ref}}|}{f_{\text{ref}}} \right)^2 + (1 - \text{MAC}(\boldsymbol{\phi}_n^{\text{mo}}, \boldsymbol{\phi}_{\text{ref}})) \tag{7.7}$$

where f_n^{mo} and $\boldsymbol{\phi}_n^{\text{mo}}$ denote the nth identified frequency and mode shape at model order mo, and f_{ref} and $\boldsymbol{\phi}_{\text{ref}}$ are the frequency and mode shape of the reference mode.

This process was repeated for increasing model orders before moving on to the next largest cluster. If any modes assigned to a new cluster had already been assigned to a previous one, they were excluded to prevent duplication across clusters. Finally, any cleaned cluster with less than five modes was considered spurious or not sufficiently stabilised, and was therefore discarded.

To demonstrate the outcome of Stages 1 to 3, Figure 7.3a presents an example stabilisation diagram from a representative dataset, indicating both the cleared modes and the cleaned clusters.

Stage 4: Semi-Automated Mode Set Selection

Up to this point, the process has been fully automated. The result of Stage 3 is set of clusters which are more likely to correspond to individual physical modes. However, Figure 7.3a shows that a large number of clusters were identified, some of which may still contain spurious modes. The aim of this stage is to select a subset of modes suitable for model updating.

This was carried out using a semi-automated approach. For a single ten-minute dataset, six

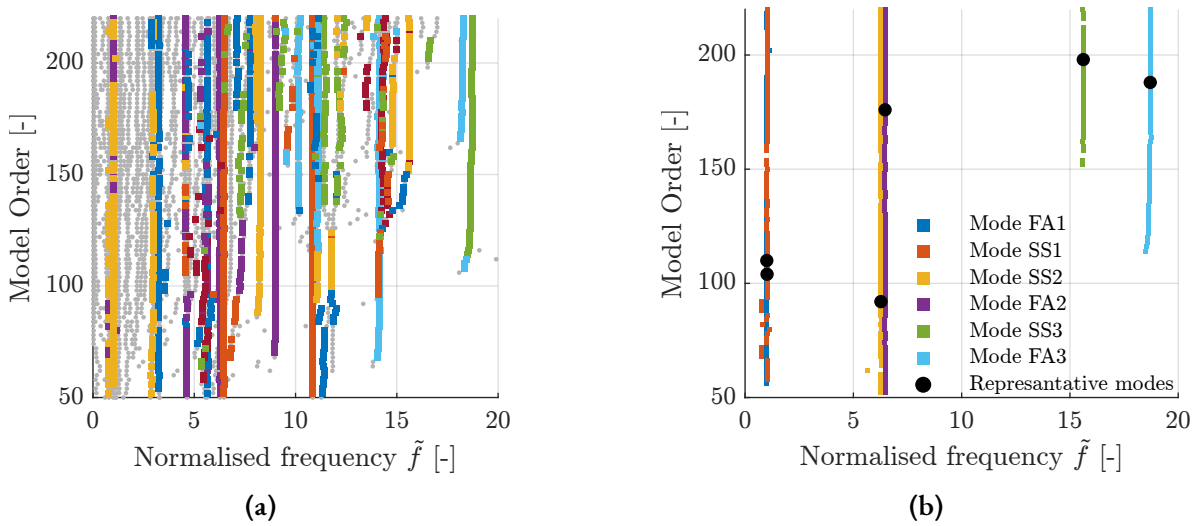


Figure 7.3: (a) Example set of cleaned clusters from a representative dataset—grey dots show rejected modes and colours indicate common clusters. (b) Example set of selected clusters and representative modes used for model updating.

clusters were manually selected. For each of these clusters, the mode corresponding to the median frequency was chosen as representative. The selection of these six clusters was based on the author's judgement, as those which were considered to correspond to tower-dominated mode shapes and therefore expected to be more sensitive to the foundation parameters. This manual selection was only performed for the first dataset within the idling period. For subsequent datasets, each cluster was assigned a representative mode based on the median frequency, and clusters were then selected by identifying those with the smallest absolute frequency difference with respect to the manually selected modes.

Figure 7.3b provides an example of the six selected clusters and their representative modes. The modes are labelled as 'FA' or 'SS' based on the direction of their dominant mode shape component, and numbered as 1 to 3 to indicate whether they correspond to a 1st, 2nd or 3rd tower mode.

Figure 7.4 shows the six identified frequencies over the idling period. Overall, the frequencies were estimated with relatively good consistency. Figure 7.5a presents the averaged normalised frequencies across the idling datasets, along with their corresponding coefficients of variation. All modes display very low variability, with mode SS2 showing the highest standard deviation, approximately 2% of its mean.

Figures 7.5b - 7.5g show the identified mode shapes in the sensor basis, averaged over

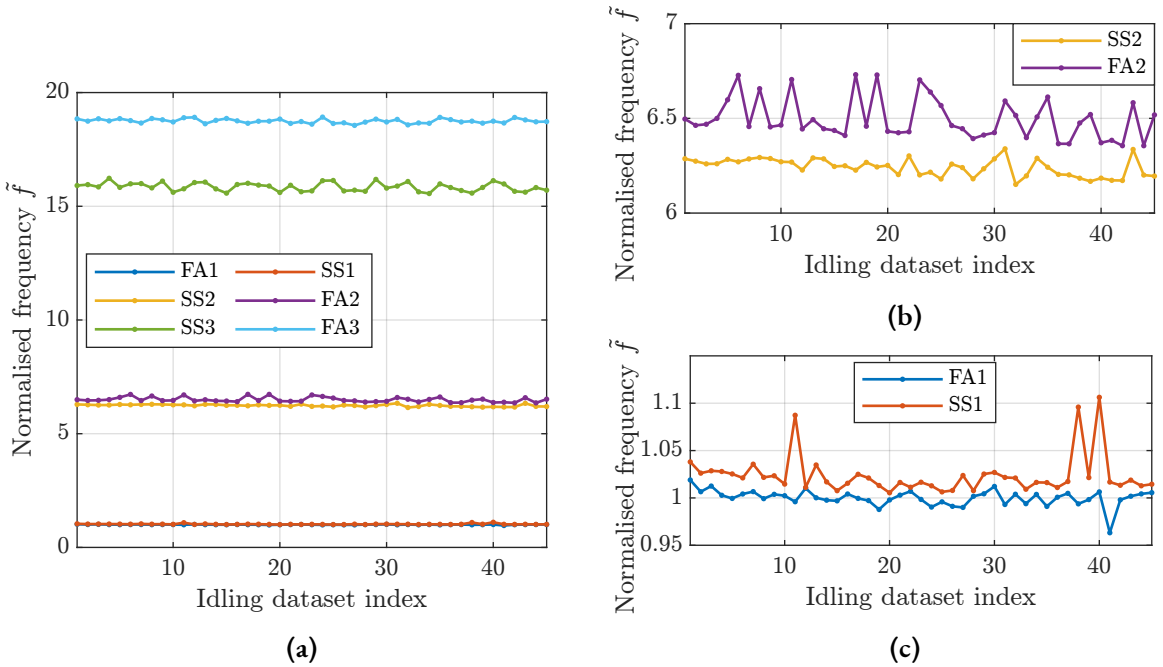


Figure 7.4: Natural frequencies selected stable modes over idling period. (a) All selected modes. (b) Modes SS2 and FA2. (c) Modes FA1 and SS1.

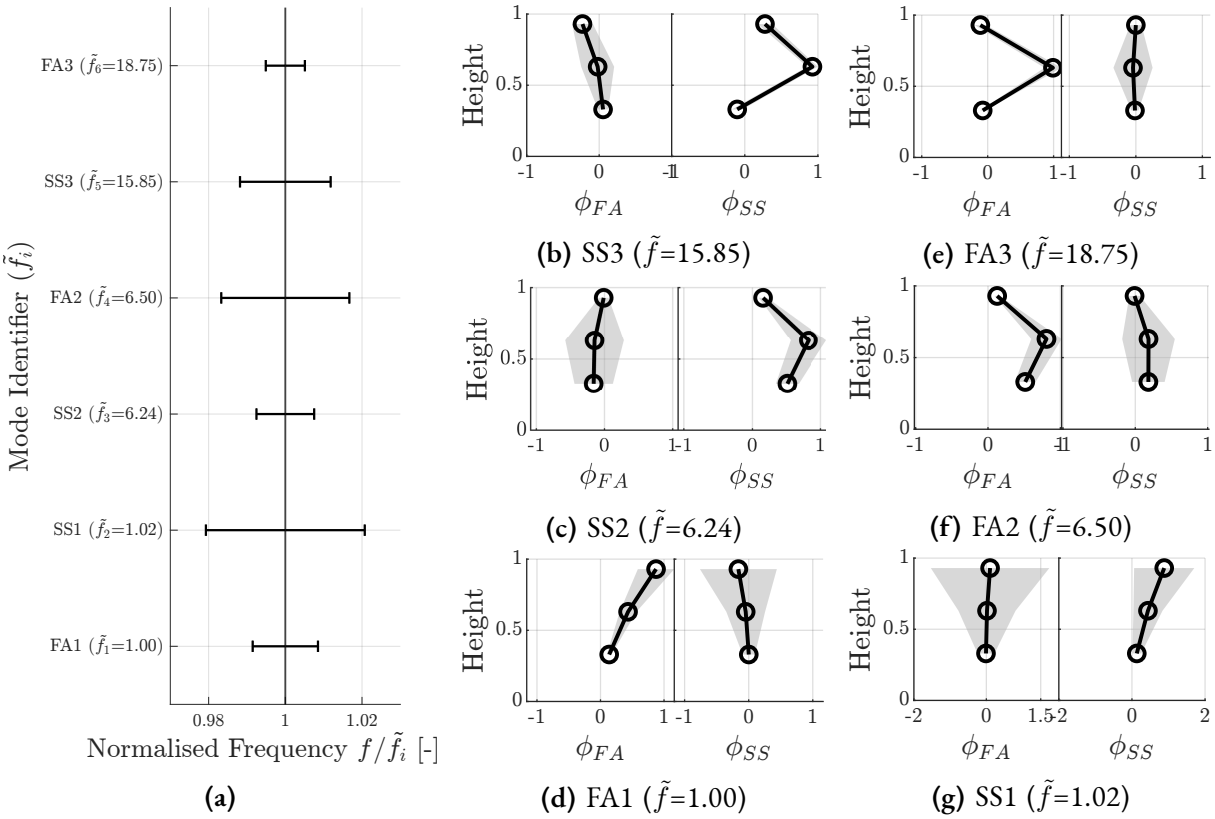


Figure 7.5: Averaged identified modal parameters over the idling period ± one coefficient of variation. Height is the normalised distance from mudline to tower top (mudline = 0, tower top = 1). (a) Frequencies normalised by average value for each mode. (b)-(g) Mode shapes projected in the fore-aft (ϕ_{FA}) and side-side (ϕ_{SS}) directions.

the idling period. The shaded regions indicate ± one coefficient of variation, reflecting the variability of each identified mode shape component across datasets. The mode shapes have been complex normalised using Equation 1.7. For each mode, the components aligned in the dominant direction exhibit lower variability. The first pair of mode shapes (FA1 & SS1) were identified with the greatest variability, most likely due to their close spacing in terms of frequency. As discussed in Chapter 1, this close spacing makes it challenging for the SSI-cov algorithm to differentiate between the two modes, often resulting in an identified mode shape that falls within their shared subspace. Similar behaviour has been reported in both numerical [21] and real-world [19] case studies.

Despite the variability in some of the identified modal parameters, the averaged values over the idling period are expected to provide more reliable estimates. In Section 7.4, both the individual ten-minute datasets and the averaged idling dataset are used to identify the foundation parameters. Moreover, the variability in each modal parameter is highly valuable for model updating, as it can be used to define the measurement uncertainty covariance matrix, $\boldsymbol{\Sigma}_D$, in the likelihood function for Method 1 (Equation 3.11). This allows greater weight to be assigned to modal parameters identified with lower variability, which was shown to reduce bias in the identified foundation parameters in the numerical case studies presented in Chapter 3.

7.3 Numerical Models

In this section, the numerical models used for identifying the foundation parameters are described. Two models were developed, one for each method.

Method 1 uses the modal properties of the full system (including the nacelle and blades) as data, and thus when parameters are updated, the relevant model must include the entire structural system. Consequently, the Simulink Model described in Chapter 2 was used to capture the complete dynamics of the offshore wind turbine structure. The specific model setup is described in Section 7.3.2, including details on the blade scaling method applied to implement the flexible rotor nacelle assembly. It will be referred to as the *Simulink Model* in this chapter.

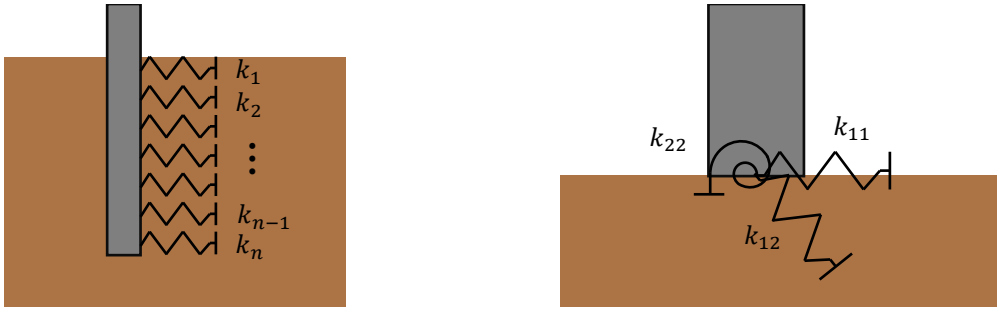
Method 2, in contrast, utilises time-domain measurements and substructuring techniques

to avoid modelling the full system, as shown in Chapter 6. The model used here is identical to that introduced in Chapter 6, where the structure was 'cut' at the top of the tower. For further details, the reader is referred back to Section 6.3.2. This model will be referred to as the *Substructure Model* in this chapter.

7.3.1 Foundation Model

Both models implement the same macro-element foundation representation. As before, the foundation is assumed to behave elastically, which is justified by the small monopile displacements expected under the idling conditions considered here.

Initially, the soil-structure interaction was modelled using a set of lateral elastic springs distributed along the depth of the embedded monopile, as illustrated in Figure 7.6a. Preliminary stiffness values for each spring were based on a modified API p-y method [32], in which the API formulation was adjusted during the design phase to account for the influence of the pile diameter and increased stiffness at small strains.



(a) Depth-wise springs **(b)** Macro-element springs

Figure 7.6: Foundation models

If a finite-element model is used to model the monopile with a depth-wise foundation, an equivalent macro-element stiffness matrix can be derived by using static condensation to retain only the degrees of freedom at the mudline [19]. This was applied to compute the 'initial design' macro-element stiffness matrix $\mathbf{K}^0_{\text{macro}}$ based on the preliminary depth-wise lateral springs:

$$\mathbf{K}^0_{\text{macro}} = \begin{bmatrix} k^0_{11} & k^0_{12} \\ k^0_{12} & k^0_{22} \end{bmatrix} = \begin{bmatrix} 0.0277 & -0.1730 \\ -0.1730 & 2.2126 \end{bmatrix} \times 10^{11}, \tag{7.8}$$

where the units of k_{11}^0, k_{22}^0 and k_{12}^0 are N/m, Nm/rad and Nm/m, respectively.

As with previous chapter, the soil conditions were assumed to be axisymmetric, and the axial and torsional responses were neglected.

7.3.2 Model for Method 1: Simulink Model

The main features of the Simulink Model used for Method 1 were provided in Chapter 2. Here, the inputs that are specific to the wind turbine of this case study are described, based on the available design information.

(i) Tower, monopile and transition piece

The tower, transition piece and monopile were modelled using Simscape's Reduced Order Flexible Solid block [203]. As discussed in Chapter 2, this block requires reduced-order stiffness and mass matrices, which were derived from a finite element model. The underlying finite element model and its parameters (geometry, material properties, and distribution of stiffness and mass) are identical to those described in the previous chapter (Section 6.3.2), and are not repeated here.

Following assembly of the full-order finite element matrices for the tower-monopile substructure (excluding the macro-element foundation), a Craig-Bampton reduction [160] was applied, where 11 external interface nodes were defined and 20 normal modes were retained. Proportional damping was applied to the reduced-order element, using a damping ratio of 1%.

(ii) Flexible Rotor Nacelle Assembly

Design information relating to this part of the structure was limited due to its commercial sensitivity. In this work, the known quantities included the RNA's total lumped mass, its associated inertia and centre of mass, as well as the overall blade length and mass.

However, to model the blades as flexible components, details of the cross-sectional properties along the blade length are required as inputs to the Simulink Model. These include geometry, mass, inertia and stiffness distributions. Due to the unavailability of such detailed design information, the known blade properties from the IEA 10 MW reference turbine [9] were

scaled to the turbine considered here. Following the *Blade Mass* scaling approach introduced in Chapter 4, two scale factors were defined: a length scale factor λ_L and a mass scale factor λ_M.

λ_L was defined as the ratio of the blade lengths of the target and reference turbines. Similarly, λ_M was defined using the ratio of their blade masses. These scale factors were then applied to each of the IEA 10 MW blade property distributions, in accordance with their units. For a reminder of the blade property distributions, their units, and the combined scale factors applied, the reader is referred to Table 4.1 in Chapter 4.

It should be noted that the blade twist distribution was not altered in the scaling process. Additionally, the dimensionless aerofoil lift and drag coefficients—which influence the aerodynamic forces acting on the blades—were also retained from the reference turbine.

Once the blades were defined, the hub and nacelle were represented by a single lumped mass with associated inertia. Its properties (including centre of mass, total mass and moment of inertia) were adjusted to match the provided overall RNA lumped mass properties.

(iii) Steady-State Operating Point for Linearisation

As outlined in Chapter 2, the full Simulink Model of the offshore wind turbine is geometrically non-linear. Consequently, it must be linearised about a steady-state operating point in order to obtain model predictions for the modal parameters.

The wind speed, blade pitch angle, and hydrodynamic forces were defined as system inputs, while the tower accelerations at the installed sensor locations were assigned as outputs. This choice of input-output structure aligns with typical offshore wind turbine monitoring setups, where environmental loads (e.g., wind and wave excitations) and control actions (e.g., blade pitch) act as external drivers, and the structural response is captured through tower-mounted accelerometers. The model was linearised about a steady-state operating point corresponding to idling conditions in still water, with constant wind speed and blade pitch values derived from averaged SCADA data recorded during the idling period.

Once the resulting linearised system matrices ($\mathbf{A}|_{op}$ and $\mathbf{G}|_{op}$) were obtained, the model-

predicted modal parameters were computed from an eigendecomposition, as shown in Section 2.3.7 (i).

7.4 Application of Methods to Field Data

In the section, the performance of each framework for identifying the foundation parameters is evaluated using the field data described in Section 7.2. The same two parameterisations of the macro-element foundation used in Chapter 6 were also considered here: (i) the compliance parameterisation, where the three compliance terms from Equation 3.23 (c_{11}, c_{22}, and c_{12}) are identified individually, and (ii) the γ parameterisation, where a single parameter γ scales the initial design macro-element compliance matrix, $\mathbf{C}^0_{\text{macro}}$ (Equation 6.11).

For both parameterisations and identification frameworks, the updating parameters were assigned uninformative uniform prior distributions. Their respective upper and lower bounds are summarised in Table 7.1.

Table 7.1: Lower and upper bounds of the uniform prior distributions applied to macro-element foundation parameters and the model error parameters within likelihood function (Equation 3.11). The initial design compliance values, c^0_{11}, c^0_{22} and c^0_{12}, are provided in Equation 6.11.

Parameter Type	Parameter	Lower Bound	Upper Bound
Macro-element	c_{11}	$0.05c^0_{11}$	$2c^0_{11}$
	c_{22}	$0.05c^0_{22}$	$2c^0_{22}$
	c_{12}	$0.05c^0_{11}$	$2c^0_{12}$
	γ	0.05	2
Likehood	σ_f	0	0.2
	σ_ϕ	0	0.5

To compute the posterior distribution, a brute-force sampling approach was initially employed, in which a large number of samples were evaluated using the prior and likelihood functions. While straightforward, this method becomes increasingly computationally expensive as the dimensionality of the parameter space grows. In this chapter, this limitation was significant only for Method 1. To overcome this challenge and improve sampling efficiency, the improved Transitional Markov Chain Monte Carlo (iTMCMC) method [175] was adopted. The iTMCMC algorithm generates samples sequentially through a series of intermediate distri-

butions, which smoothly transition from the prior to the posterior. This allows the sampling process to concentrate more densely in regions of higher posterior probability, thereby improving convergence and reducing computational waste. Importantly, all posterior evaluations from each stage of the iTMCMC process were retained and aggregated to construct the final approximation of the posterior distribution. This strategy was found to yield more consistent and stable estimates compared to using only the samples from the final stage of the algorithm.

After sampling, the *maxmimum a posteriori* (MAP) estimate was obtained by minimising the negative log-posterior using MATLAB's `fminsearch` algorithm [172].

7.4.1 Method 1 Results

For Method 1, the six modes identified in Section 7.2.1 were used to define the likelihood function (Equation 3.11). Results were obtained using the estimated modal parameters from each ten-minute dataset, as well as from a dataset where the modal parameters across the idling period were averaged. When updating the individual macro-element compliance parameterisation, 15000 samples were obtained from each iTMCMC stage. However, due to the reduced dimensionality when updating the γ parameterisation, only 4000 samples per stage were generated to improve the computational cost.

(i) Estimation of individual compliance parameters (Method 1)

Figure 7.7 presents the posterior distribution obtained by updating each compliance parameter individually, using the averaged idling dataset. The wide posterior distributions for each macro-element compliance indicate significant uncertainty associated with each parameter and their practical unidentifiability. Moreover, the marginal distribution for the lateral spring c_{11} is concentrated near its lower bound. These observations likely arise from a combination of the high sensitivity of the foundation parameters to measurement error and their ability to compensate for one another to produce similar modal parameters.

The posterior distributions for the likelihood model error parameters, σ_f and σ_ϕ, provide insight into the model accuracy. In this case, Figure 7.7 shows large values for σ_f suggesting significant modelling errors in the predicted frequencies. σ_ϕ exhibits lower values, indicating

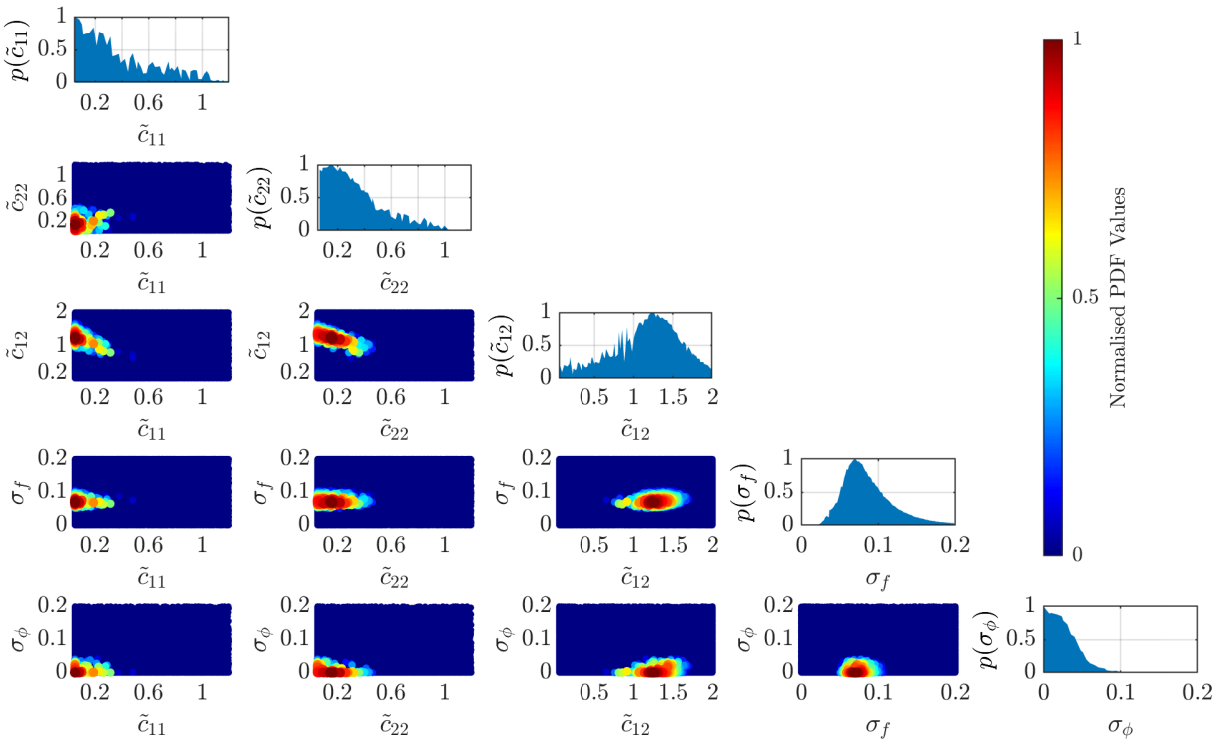


Figure 7.7: Posterior distributions for macro-element compliances and model error parameters from updating with the **averaged idling dataset**. All units are dimensionless: compliances are normalised by their initial design values and probability densities are normalised by the *maximum a posteriori* estimate.

comparatively smaller modelling error in the mode shapes. However, due to the sparsity of sensors, it is likely that the mode shape components can still be matched reasonably well, even in the presence of modelling inaccuracies.

Figure 7.8 shows the results from updating with each individual ten-minute dataset. The figure plots the MAP estimates and associated 95 % confidence interval for each parameter across the idling period. The results reinforce the observations made above and also reveal significant variability in the MAP estimates between datasets, highlighting the sensitivity of the parameters to measurement error. Furthermore, the uncertainty associated with σ_ϕ was observed to be consistently smaller than σ_f, indicating that the framework is finding a parameterisation which prioritises minimising the error in the mode shape components over the frequencies.

Overall, these results suggest that the Simulink Model used here introduces significant error in predicting the modal parameters, which in turn leads to unreliable estimates for the foundation parameters. Nonetheless, the ability to quantify uncertainty shows the benefit of the Bayesian framework to support the user in making informed engineering decisions. For

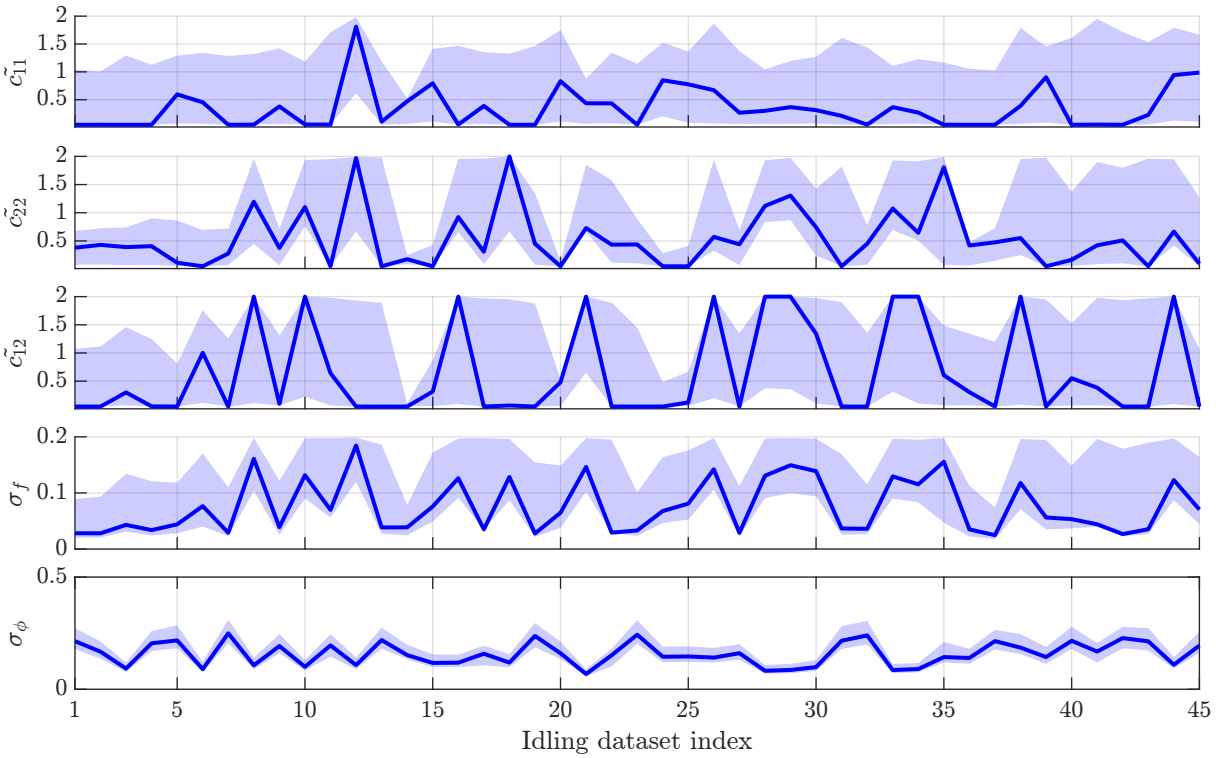


Figure 7.8: MAP estimates and 95% confidence intervals for macro-element compliances and model error parameters when updating with **individual ten-minute datasets** across the idling period. All units are dimensionless: compliances are normalised by their initial design values.

example, a more informative prior can be imposed on the foundation parameters to place less weight on the model and data. This was effectively implemented in this chapter by considering the alternative foundation parameterisation involving γ, which is discussed in the following section.

(ii) Estimation of γ parameterisation (Method 1)

Figure 7.9 shows the posterior distribution obtained by updating γ, using the averaged idling dataset. The relatively narrow marginal distribution for γ indicates improved identifiability of the macro-element compared to the case where each spring is updated independently. However, the distribution shows two peaks, which correspond to parameterisations that favour different modal parameters. This multimodal behaviour highlights the complexity of the posterior distributions and reinforces the need for global optimisation methods to avoid convergence at local minima.

Figure 7.10 presents the *maximum a posteriori* (MAP) estimates of γ and the corresponding

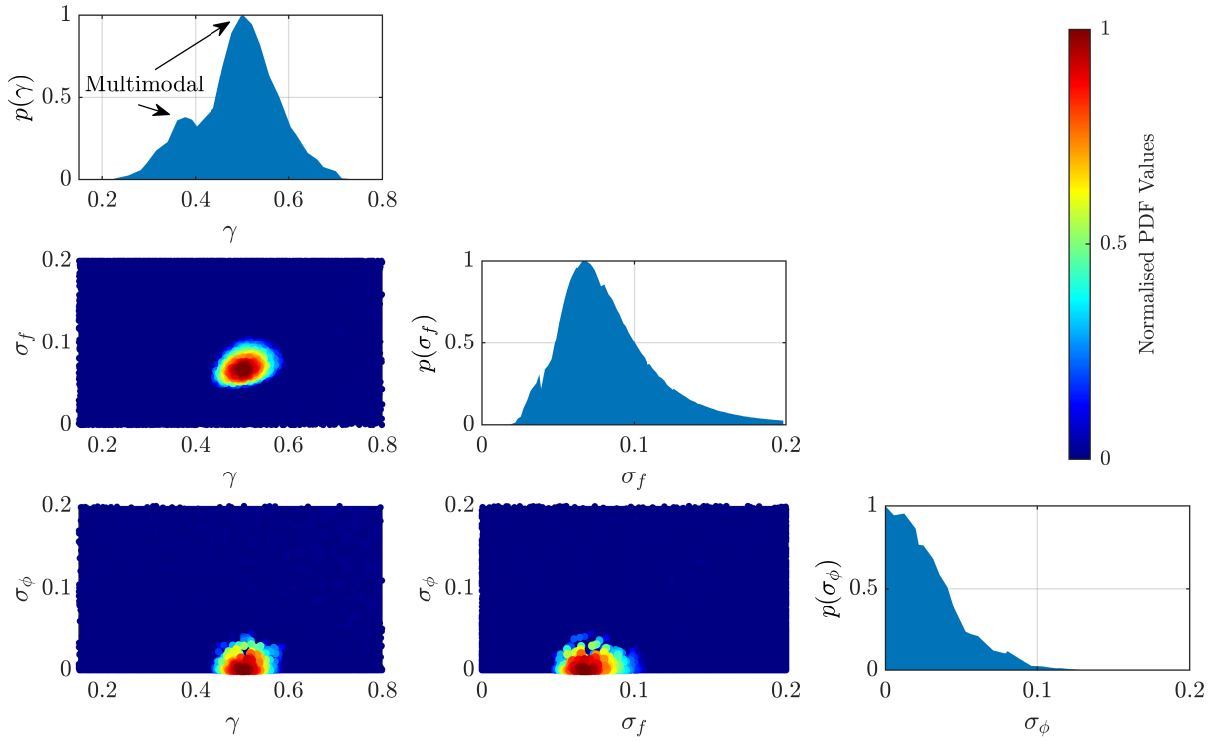


Figure 7.9: Posterior distributions for macro-element γ scaling factor and model error parameters when updating with the **averaged idling dataset**. All units are dimensionless: γ is relative to the initial design and probability densities are normalised by the *maximum a posteriori* estimate.

95 % confidence intervals obtained from the individual ten-minute datasets. The narrow 95 % confidence intervals indicate that γ was generally identified with relatively low uncertainty. However, the notable variability in the MAP estimates across datasets suggests that the optimal values of this macro-element parameterisation remains sensitive to measurement noise. Consequently, the results from the averaged idling dataset, which is expected to exhibit lower measurement noise, were considered more reliable than those derived from the individual ten-minute segments.

The posterior distributions for σ_f and σ_ϕ were similar to those obtained using the individual compliance parameterisation. This further supports the conclusion that the model used here introduces significant prediction error and should be revised before subsequent use. In this case, the primary source of the model error likely arises due to inaccurate scaling of the unknown blade properties. Future modelling efforts should focus on refining this aspect to improve the identifiability of the foundation parameters. Nonetheless, σ_f and σ_ϕ allow the model error to be quantified, improving the interpretability of the inference results whilst recognising limitations

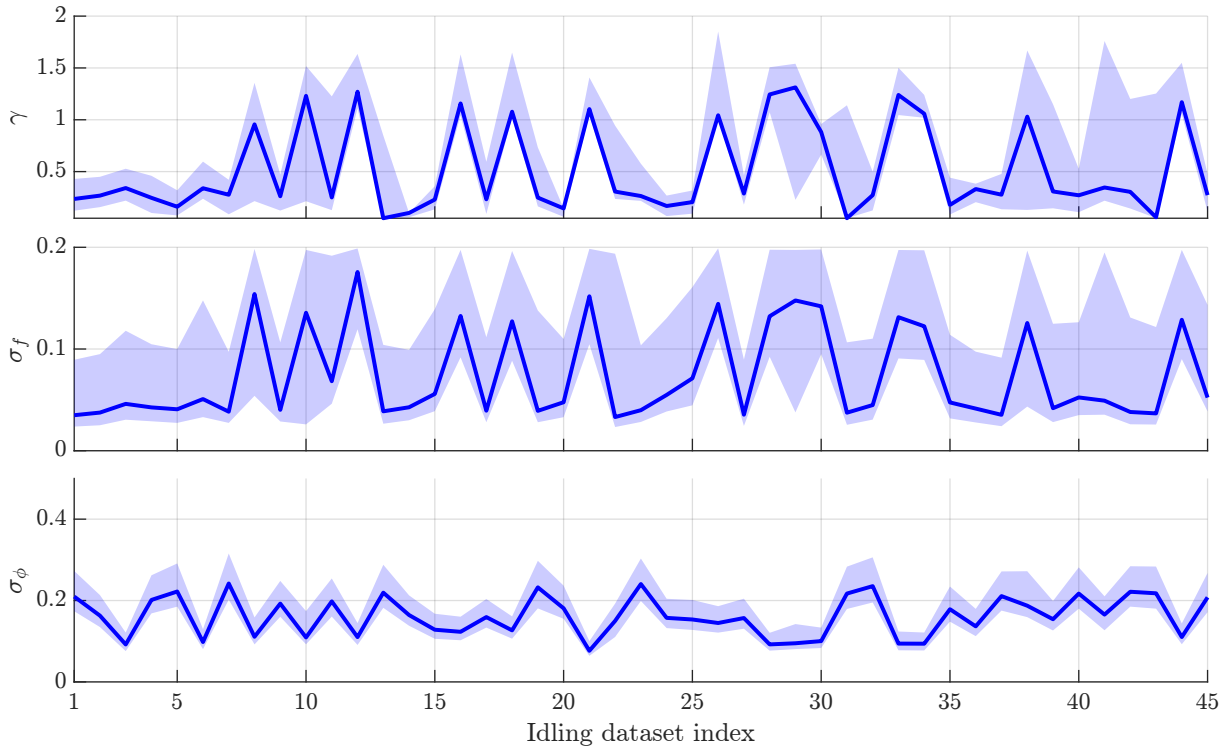


Figure 7.10: MAP estimates and 95% confidence intervals for macro-element γ scaling factor and model error parameters when updating with **individual ten-minute datasets** across the idling period. All units are dimensionless: γ is relative to the initial design.

in the structural model.

(iii) Model Predictions (Method 1)

Model-predicted modal parameters were obtained using three realisations of the macro-element foundation: the initial design $\mathbf{K}^0_{\text{macro}}$, the MAP estimate from updating individual compliance parameters (c_{11}, c_{22}, c_{12}), and the MAP estimates from updating γ. Table 7.2 compares the frequency error and MAC values between each of the MAP estimates and measured data. Figure 7.11 presents this comparison visually and additionally includes the predictions from the initial design foundation $\mathbf{K}^0_{\text{macro}}$. The measured results correspond to the averaged idling dataset.

Table 7.2 and Figure 7.11 show that, although the MAP estimates differ significantly between the two updated parameterisations, the resulting model predictions are remarkably similar. This supports the conclusion that the individual macro-element parameters are practically unidentifiable, as multiple parameter combinations yield comparable predictive performance. Nevertheless, both sets of MAP estimates result in model predictions that align more closely

Table 7.2: Comparison of model-predicted and measured modal parameters for *maximum a posteriori* (MAP) foundation estimates. Results correspond to those obtained from the averaged idling dataset and consider two parameterisations: individual compliances or the γ scaling factor.

Mode	**Compliance MAP:** $\begin{bmatrix}\tilde{c}_{11}\\ \tilde{c}_{12}\\ \tilde{c}_{22}\end{bmatrix} = \begin{bmatrix}0.05\\ 1.25\\ 0.15\end{bmatrix}$		**γ MAP:** $\begin{bmatrix}\tilde{c}_{11}\\ \tilde{c}_{12}\\ \tilde{c}_{22}\end{bmatrix} = \begin{bmatrix}0.50\\ 0.50\\ 0.50\end{bmatrix}$	
	f diff. [%]	MAC [-]	f diff. [%]	MAC [-]
FA1	1.90	0.9985	−0.75	0.9992
SS1	0.62	0.9984	−2.05	0.9992
SS2	−1.18	0.4951	−1.66	0.4891
FA2	−6.48	0.8063	−6.94	0.8000
SS3	−13.81	0.9775	−12.82	0.9677
FA3	−7.70	0.9371	−7.64	0.9339

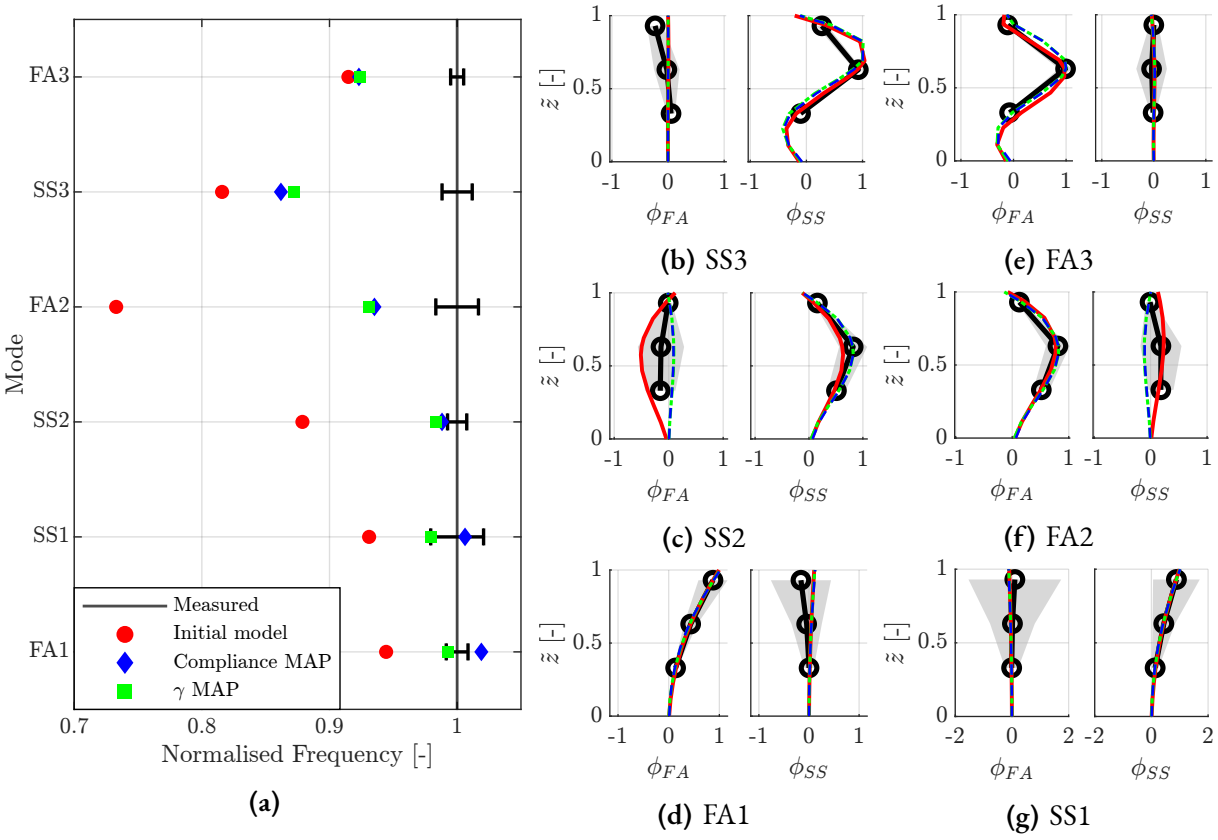


Figure 7.11: Comparison of model predictions for Method 1. Error bars/shaded areas indicate ± one coefficient of variation in the measured values. (a) Frequencies normalised by averaged idling measured frequency for each mode. (b)-(g) Mode shapes. $\tilde{z} = 0$ = mudline; $\tilde{z} = 1$ = tower top. *Solid black: Measured; Solid red: Initial model; Dashed blue: Compliance MAP; Dash-dot green: γ MAP.*

with the measured modal parameters than those based on the initial design values.

The predictions for the first tower modes (FA1 and SS1) show the best agreement with the measured data. These lower-frequency modes are expected to be less influenced by modelling uncertainties in the rotor nacelle assembly (RNA), as their dynamics are dominated by the tower-foundation interaction and occur below the frequency range of the first blade mode. In contrast, the model predictions for higher-order modes exhibit larger discrepancies from the measurements, hinting that the primary source of model error stems from inaccuracies in the RNA representation.

7.4.2 Method 2 Results

Method 2 used data from all installed accelerometers and strain gauges to identify both foundation parameterisations for each ten-minute dataset. The initial augmented state vector $\boldsymbol{x}_0^a$ was assumed to be zero, with an associated covariance matrix $\boldsymbol{\Sigma}_0^{x^a} = 0.01\mathbf{I}$, where $\mathbf{I}$ is the identity matrix. To allow the augmented Kalman filter estimates to converge, the initial 200 seconds were ignored and the likelihood was computed using the remaining data from the ten-minute dataset (i.e., $N_{\text{burn}} = 5000$ and $N_{\text{batch}} = 15000$ in Equation 6.7). As discussed at the beginning of Section 7.4, a brute-force approach was used to evaluate the posterior distribution, where Latin hypercube sampling was used to efficiently draw samples from the uniform priors. For the compliance parameterisation, 1000 samples were used to approximate the posterior, whereas only 750 samples were required for γ.

(i) Estimation of individual compliance parameters (Method 2)

Figure 7.12 shows a representative posterior distribution for the compliance parameters, obtained by updating them independently using an example ten-minute dataset. The figure reveals that the posterior distribution exhibits wide and flat regions, indicating that these individual macro-element parameters are practically unidentifiable. Among the three parameters, only the rotational spring c_{22} could be identified from the measurements. In contrast, the lateral compliance c_{11} was found to consistently tend towards its lower bound, likely due to its sensitivity to measurement error. One potential cause for this could relate to the soil being too stiff for the available measurements to provide information about the lateral response. As a result, this

framework appears to be only capable of classifying the lateral soil conditions as broadly 'soft' or 'stiff'. Therefore, these results motivate the use of the more informative γ parameterisation, which provides the necessary constraint to be able to identify the macro-element foundation.

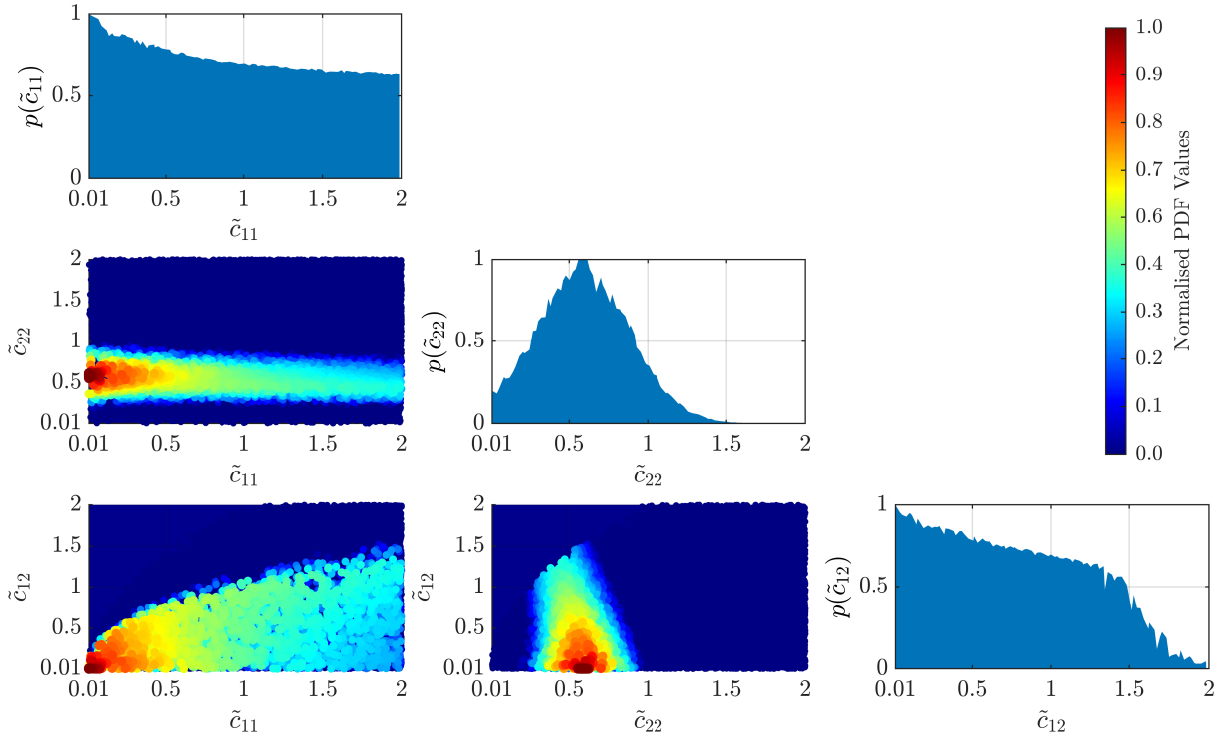
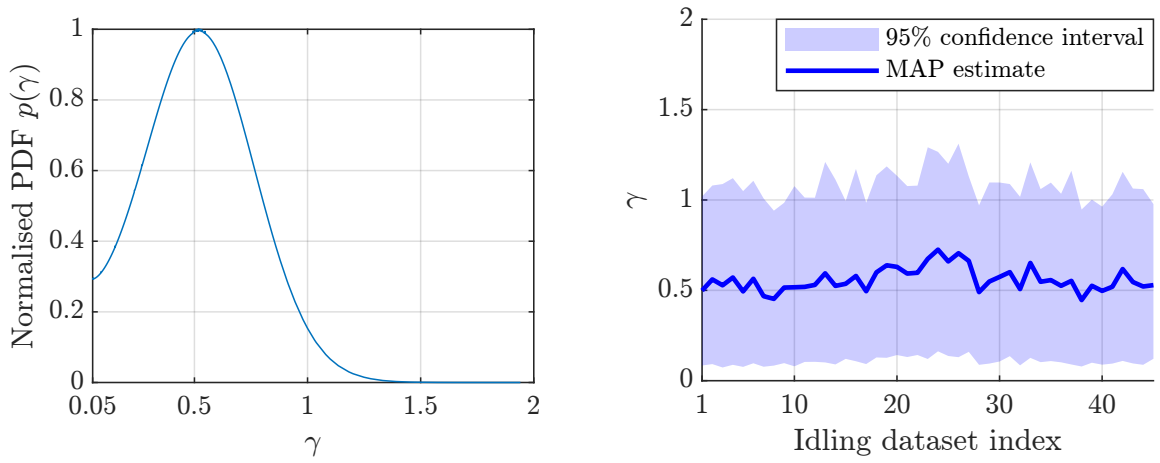


Figure 7.12: Example posterior distribution for the compliance parameters obtained using a representative ten-minute dataset. All units are dimensionless: compliances are normalised by their initial design values and probability densities are normalised by the *maximum a posteriori* estimate

(ii) Estimation of γ parameterisation (Method 2)

Figure 7.13a shows a representative posterior distribution for γ, obtained from an example ten-minute dataset. The distribution exhibits a well-defined peak, indicating the identifiability of this foundation parameter using Method 2. Moreover, Figure 7.13b presents the MAP estimates and corresponding 95% confidence intervals across the entire idling period. These results show the consistency of the posterior distributions across multiple datasets, which increased confidence in both the robustness of the framework and the reliability of the identified values.



(a) (b)

Figure 7.13: Posterior distributions for soil parameter γ, obtained from Method 2. γ is relative to the initial design. (a) Example distribution from a representative ten-minute dataset. (b) MAP estimates and corresponding 95% confidence interval across idling period.

(iii) Input-state estimation results (Method 2)

One of the key advantages of Method 2 is that the inputs and states are obtained as a by-product of the estimation process. This section presents input-state estimation results for three cases in which the macro-element foundation was defined using: the initial design stiffness matrix $\mathbf{K}^0_{\text{macro}}$, the MAP estimates from the γ parameterisation, and the MAP estimates from the compliance parameterisation. The corresponding MAP estimates from the representative ten-minute dataset used in this analysis are listed in Table 7.3.

Table 7.3: MAP estimates for each compliance parameter obtained with Method 2 from a representative ten-minute dataset. These estimates were used to the produce the input-state estimation results presented in Section 7.4.2 (iii). Normalised compliance values are relative to the initial design.

Parameterisation	Normalised Compliance Values (MAP estimates)
Compliance	$[\tilde{c}_{11} \quad \tilde{c}_{12} \quad \tilde{c}_{22}] = [0.05 \quad 0.59 \quad 0.05]$
γ	$[\tilde{c}_{11} \quad \tilde{c}_{12} \quad \tilde{c}_{22}] = [0.52 \quad 0.52 \quad 0.52]$

Figure 7.14 shows time histories for the two estimated inputs: the tower-top lateral forces in two orthogonal directions. The estimated inputs are nearly identical across all three cases, suggesting that the tower-top forces are invariant to the choice of foundation parameters. One possible explanation is that the feedthrough matrix $\mathbf{J}$, which maps the inputs to the outputs,

has negligible dependence on the foundation parameters. As a result, the estimated inputs are primarily determined by the measurements.

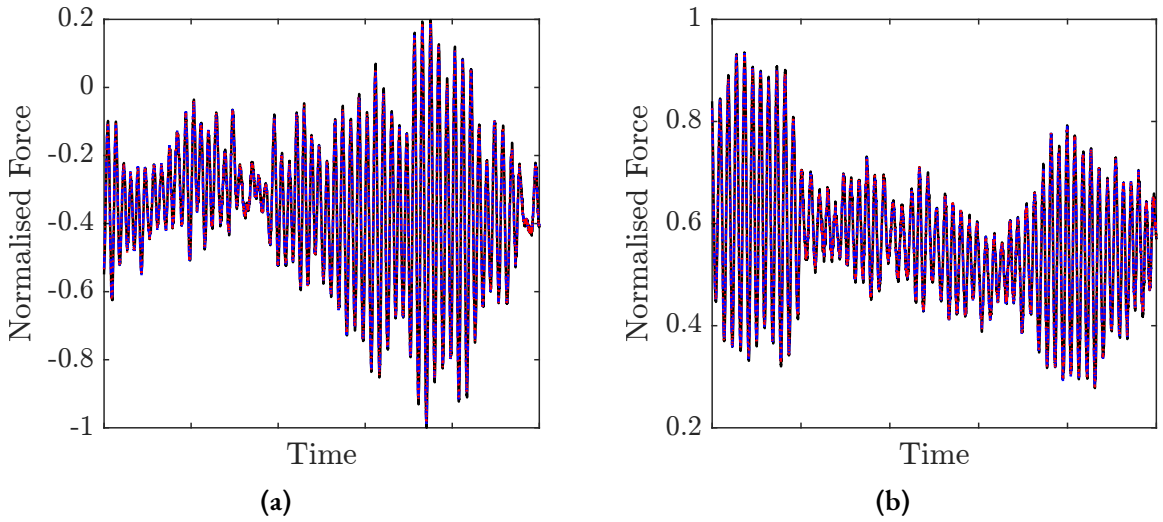
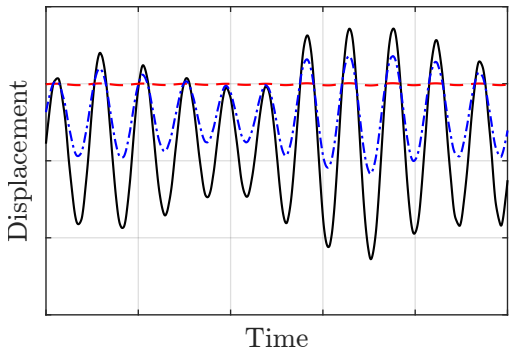


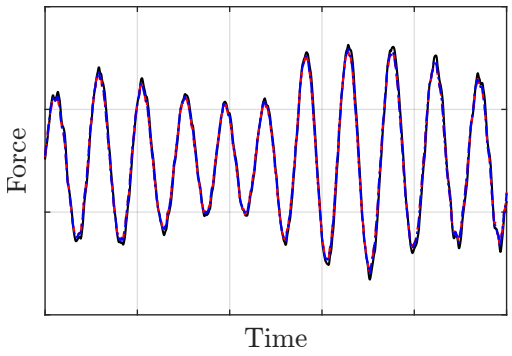
Figure 7.14: Estimated tower-top input forces from three parameterisations of the macro-element foundations: initial design (solid black), compliance MAP estimated (dashed blue), and γ MAP estimated (dotted red). (a) Normalised East–West force; (b) Normalised South–North force.

Figure 7.15 presents time histories for the estimated response at the ground-level. In this case, the estimated displacements and rotations show more sensitivity to the foundation parameters, whereas the corresponding internal shear forces and bending moments are nearly identical. This result may reflect the fact that the internal forces are driven mainly by the external inputs, which were shown to be insensitive to the foundation parameters. However, because the underlying system model changes with each foundation parameterisation, the state trajectories must adjust accordingly to reproduce the same internal forces. It is interesting to note that since the stress response is determined by the internal bending moments at a given location, the resulting fatigue calculations would also be largely insensitive to the foundation parameters for this case study.

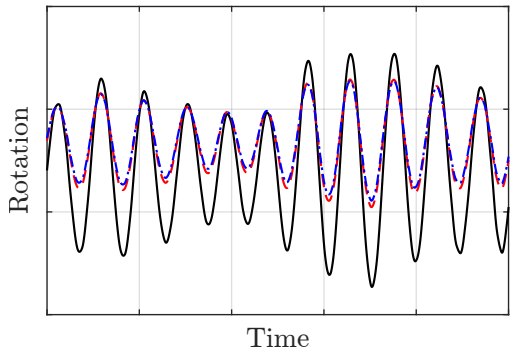
In summary, the estimated inputs—and therefore the internal forces and moments—are near identical across foundation parameterisations, as they are primarily governed by the available measurements in this specific sensor setup. However, the state estimates differ, since changes in foundation stiffness alter the relationship between applied forces and the resulting structural response.



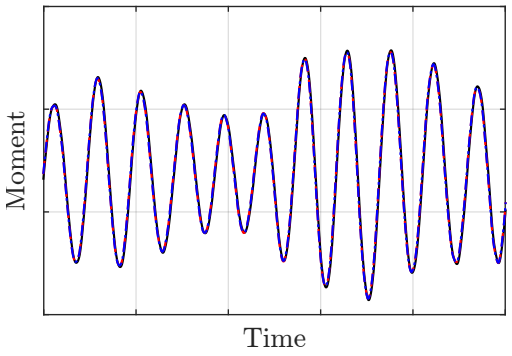
(a) East–West displacement



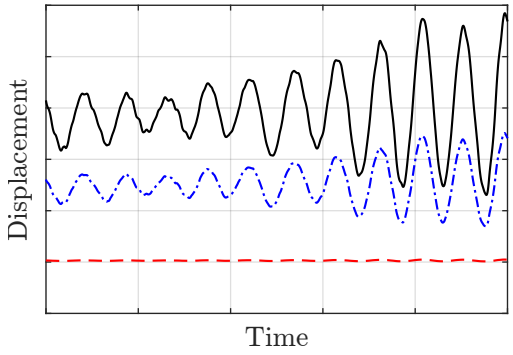
(b) East–West shear force



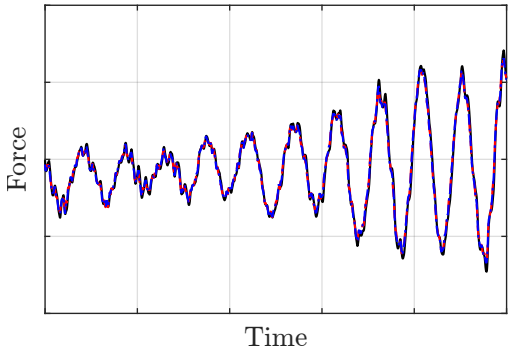
(c) Rotation about South–North axis



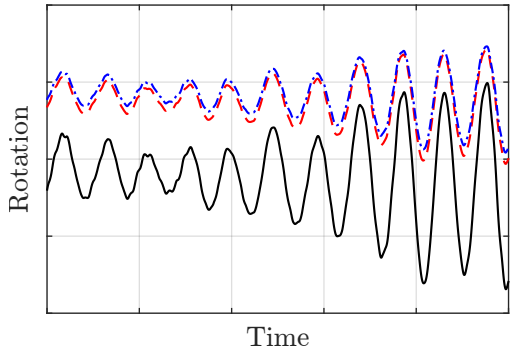
(d) Bending moment about South–North axis



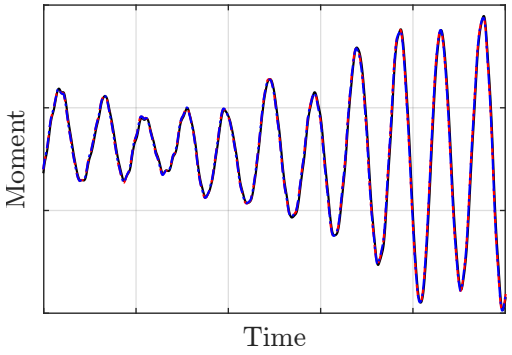
(e) South–North displacement



(f) South–North shear force



(g) Rotation about East–West axis



(h) Bending moment about East–West axis

Figure 7.15: Excerpts of estimated ground-level states, internal shear forces and bending moments. Units have been removed for confidentiality reasons. *Solid black: Initial design; Dashed red: Compliance MAP; Dash-dot blue: γ MAP.*

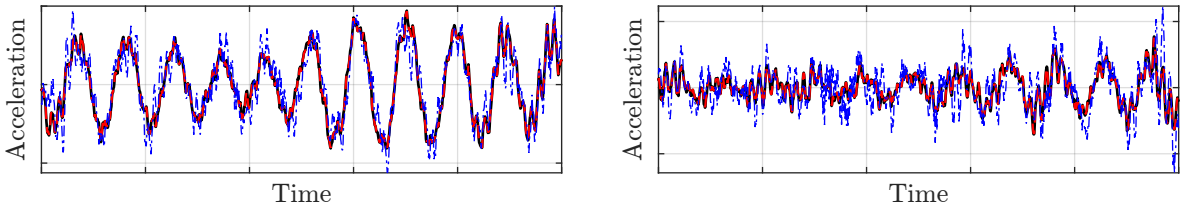
(iv) Model Predictions (Method 2)

To assess the accuracy of the updated model, the measurements were recreated using two approaches. The first uses the augmented posterior state vector from the MAP-estimated foundation parameters to compute the measurements through the output equation (Equation 6.3). This method is referred to as *AKF Reconstructed*.

The second approach uses only the posterior estimates of the inputs, which represent the applied forces at the top of the tower (example time histories are provided in Figure 7.14). A forward simulation was then performed using these inputs, where the substructure dynamics were governed by the full, unreduced mass, damping, and stiffness matrices (including the MAP-estimated foundation parameters), as described by Equation 6.14. The system response corresponding to each measurement was then extracted for comparison. Results from this second approach are referred to as *Forward Simulation*.

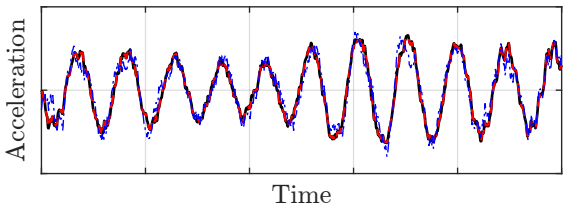
Figure 7.16 presents time history excerpts, comparing the sensor data with the two methods for recreating the measurements, where the foundation was defined by the MAP estimate for γ. The AKF Reconstructed response shows closer agreement with the measurements compared with the Forward Simulation results, which was expected as the AKF incorporates sensor data to correct the states at each time step. Nonetheless, Figure 7.16 also shows strong agreement between the Forward Simulation and the measurements. This is a more revealing test, as it relies heavily on the accuracy of the model. Therefore, the strong performance of the Forward Simulation implies that the inputs are physically meaningful and that the substructure model accurately represents the system dynamics. These results provide greater confidence in the updated foundation parameters and the robustness of the inference framework.

Although Figure 7.16 only presents model predictions using the MAP estimates from the γ parameterisation, Forward Simulation results were also obtained using a model where the substructure was defined by the initial-design foundation and the MAP estimates from the compliance parameterisation. The mean absolute errors (MAEs) for all three foundations are compared in Figure 7.17, where each model-predicted signal is compared against the measured response. The results show that the MAEs are comparable across the two MAP-estimated

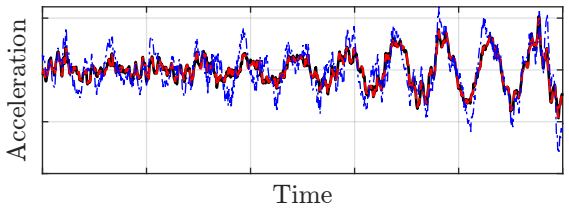


(a) ACC-TPx

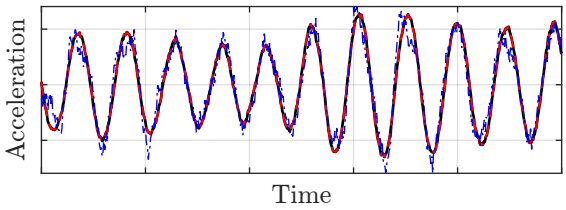
(b) ACC-TPy



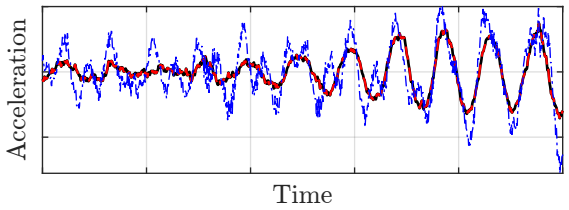
(c) ACC-TMx



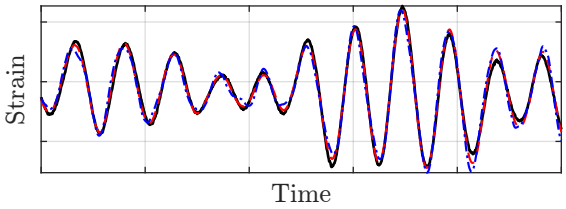
(d) ACC-TMy



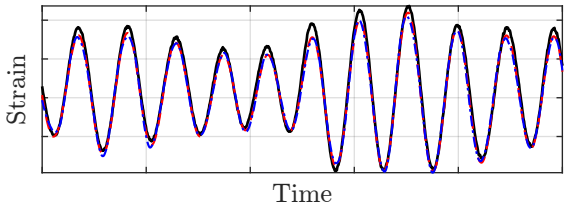
(e) ACC-TTx



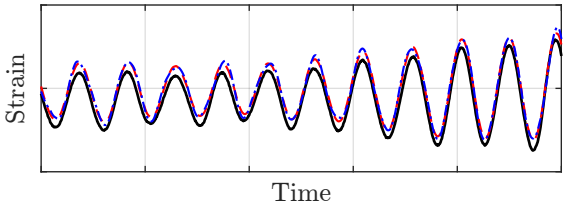
(f) ACC-TTy



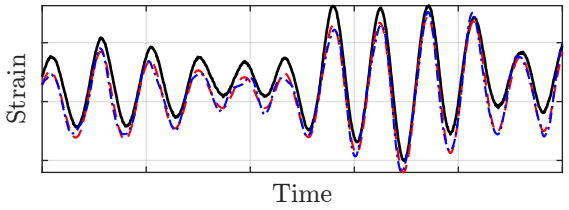
(g) SG 35



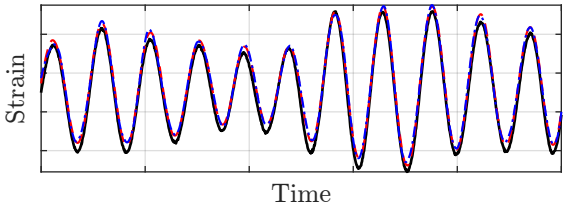
(h) SG 90



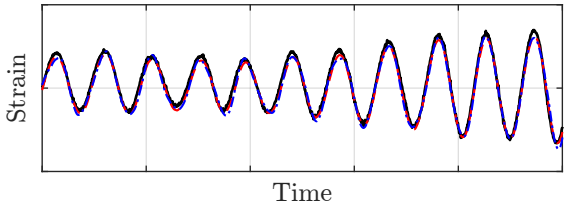
(i) SG 155



(j) SG 215



(k) SG 275



(l) SG 335

Figure 7.16: Excerpts of recreated measurements compared with sensor data. Units are hidden for confidentiality reasons. *Legend: Solid black = Measured; Dashed red = AKF Reconstructed; Dash-dot blue = Forward Simulation.*

foundations. This likely reflects the similar values of the rotational spring c_{22} from both MAP estimates, to which the measurements show significantly greater sensitivity. However, both MAP estimates exhibit reduced MAEs compared to the initial design, indicating that the updated models more accurately predict the response corresponding to the available measurements.

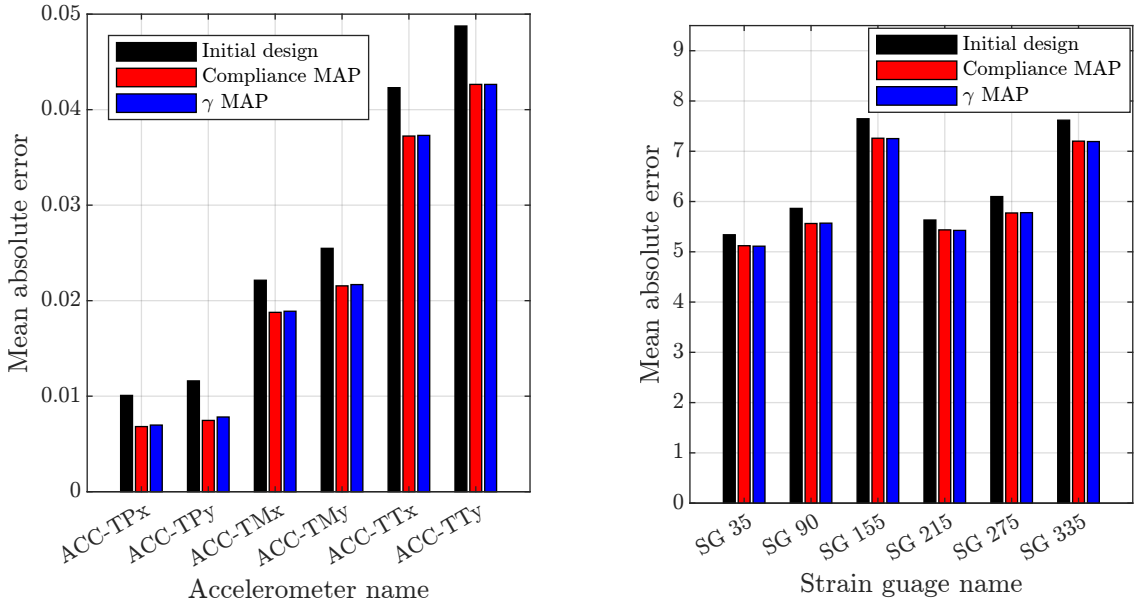


Figure 7.17: Comparison of mean absolute errors between measured sensor data and corresponding responses from Forward Simulations driven by estimated inputs: (a) accelerometer measurements; (b) strain gauge measurements. Acceleration and strain units are anonymised.

7.4.3 Discussion

Having presented the individual results, this section compares the performance between the two methods. In Method 1, the identified model error parameters, σ_f and σ_ϕ, revealed significant prediction error from the Simulink Model, suggesting notable modelling inaccuracies. In contrast, Method 2 produced accurate model predictions, indicating a more reliable structural model. As the Substructure Model of Method 1 does not require explicit modelling of the RNA, this implies that the dominant error in the Simulink Model arises from inaccurate scaling of the flexible blade properties. Consequently, Method 2 provided more consistent estimates of the foundation parameters and showed less sensitivity to measurement error.

In addition to the RNA modelling assumptions, Method 1 relies on a greater number

of user decisions that are likely to affect parameter identifiability, such as the semi-automated mode set selection process or the definition of the covariance matrix $\boldsymbol{\Sigma}_D$. It also assumes a constant yaw angle for each ten-minute dataset to transform the raw sensor data to a FA–SS coordinate system, an assumption which is often only valid for a limited number of datasets in practice. Finally, the OMA method used to extract the modal parameters, SSI-cov, assumes that the excitations to the structure can be reasonably approximated as white noise.

In contrast, fewer modelling assumptions and user choices are required for Method 2, making it a more robust framework for this case study. Moreover, Method 2 performs input-state estimation simultaneously; the results of which can be used for further applications, such as estimating remaining useful life [14, 15].

Despite these differences, both methods shared similar conclusions: the individual compliance parameters were found to be practically unidentifiable, and required reparametrisation of the macro-element in terms of γ. Figure 7.18 compares two distributions of the soil compliance parameter γ: the marginal posterior distribution from the most reliable Method 1 result—based on the averaged idling dataset—and the frequentist distribution of MAP estimates obtained across all idling datasets using Method 2. The close agreement between these distributions, which were obtained from independent methods, provides cross-validation of the frameworks and improves confidence in the identified results.

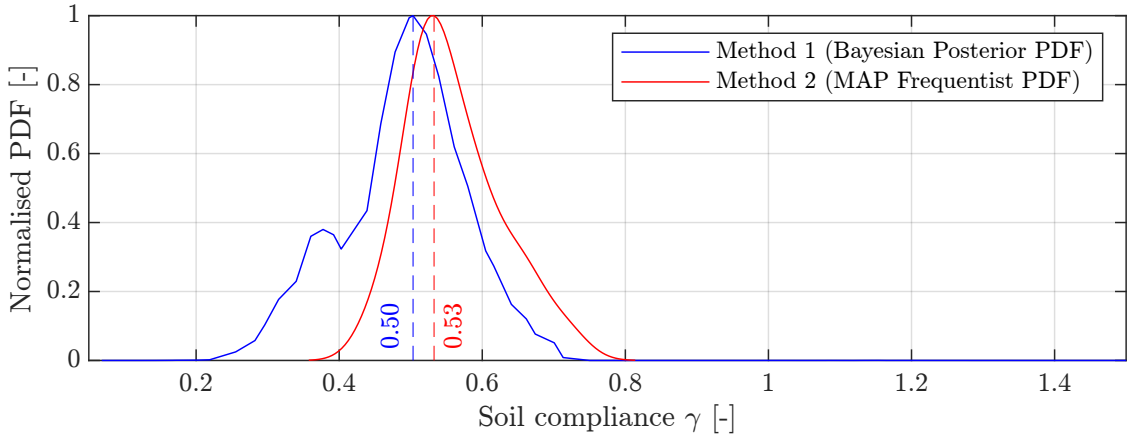


Figure 7.18: Comparison of methods for foundation parameter estimation. Blue: marginal distribution obtained using Method 1 with the averaged idling dataset. Red: frequentist distribution of MAP estimates obtained from Method 2.

For the idling conditions considered here, both frameworks resulted in a most probable

foundation stiffness approximately twice the initial design stiffness. This increase in stiffness was expected given the initial model underestimated the frequencies (see Figure 7.11a). The associated uncertainty for γ was found to be significantly reduced, with 95 % confidence intervals falling within 50 % of the MAP estimate. It should be noted that this level of uncertainty is considered relatively low in the context of soil-structure interaction [21, 204, 205].

The applied frameworks show the successful identification of the macro-element foundation using vibration data from a limited number of sensors. The identified foundation estimates can be used to compare and validate different soil-structure interaction design methods [19]. Alternatively, the updated Substructure Model could be used as a digital twin for ongoing monitoring and asset integrity assessment [13].

7.5 Summary

This chapter has presented a real-world case study where the macro-element foundation parameters of a large-scale offshore wind turbine were estimated. The framework introduced in Chapter 6 and referred to as Method 2 in this chapter, combines input-state estimation with substructuring techniques, to eliminate the requirement of modelling the rotor nacelle assembly. Its performance is compared with a more conventional Bayesian model updating framework based on modal data which was referred to as Method 1.

In both methods, the individual spring parameters of the macro-element foundation were found to be practically unidentifiable, which can be attributed to the low sensitivity of the measurements to the (likely very stiff) lateral spring. To overcome this, the macro-element was re-parameterised using γ to uniformly scale the initial design spring values. Using this parameterisation, both methods identified *Maxima a Posteriori* estimates approximately twice the initial design stiffness, which was expected due to the observed underestimation of frequencies from the initial Simulink Model.

The discrepancy in the updated model predictions indicates that a significant source of error is associated with the RNA. This highlights one of the primary advantages of the proposed substructuring approach (Method 2), in that it does not need to model this component. Furthermore, Method 1 relies on more user-defined assumptions, increasing the risk of introducing

measurement and model error, making it generally less robust in comparison. Nonetheless, it can still provide reasonable estimates of the foundation parameters, and the Bayesian framework allows assessment of both the reliability of the results and the accuracy of the model.

The computational cost of each method is an important further consideration for practical implementation. All results were obtained on a Windows machine equipped with an 11th Gen Intel Core i7-11700K CPU (8 cores, 16 threads, 3.60 GHz) and 64 GB RAM. On this system, a single posterior evaluation required approximately 0.17 seconds for Method 1 and 0.79 seconds for Method 2. While Method 2 has a higher per-sample cost, its computations are fully parallelisable, and it requires fewer samples to approximate the posterior distribution compared to Method 1. As a result, the full posterior for the γ parameterisation was obtained in approximately 3 minutes using 8 CPUs for Method 2, whereas Method 1 required around 1 hour on a single CPU (excluding the time for operational modal analysis and simulating for a steady-state operating point). This demonstrates that Method 2 can deliver results faster than the duration of a ten-minute dataset, making it viable for semi-online estimation. In contrast, Method 1 could not be parallelised at the sample level due to its reliance on the iTMCMC algorithm; however, parallelisation across multiple datasets was employed to reduce the total computational time.

This study provides a basis for more robust and efficient parameter identification in offshore wind turbine systems, but several directions for future research remain. While the current work focused on idling conditions, future investigations should explore the applicability of both methods under operational loading, where aerodynamic and control-related effects introduce greater complexity. Furthermore, alternative foundation modelling approaches, such as the PISA method, could be incorporated and benchmarked against the macro-element model used here to improve the design of monopiles for next-generation turbines. Overall, this chapter has demonstrated that a substructuring approach—when combined with input-state estimation—enables robust and computationally efficient identification of foundation parameters, circumventing the need for detailed modelling of the commercially sensitive rotor nacelle assembly.

CHAPTER 8

Conclusions

Overall this thesis has improved our understanding of structural health monitoring methods for offshore wind turbines, through both numerical and real-world case studies. This conclusions chapter returns to the four work packages (WPs) introduced in Figure 1.5 (Chapter 1), summarising their key contributions and recommendations, before suggesting ideas for future work.

8.1 Key Contributions from Work Packages

WP 1: Bayesian model updating for identifying foundation parameters

Key Recommendation: *Bayesian model updating is recommended over deterministic frameworks for foundation parameter identification, as it quantifies uncertainty and thereby improves interpretability.*

The primary aim of Work Package 1 was to extend previous deterministic work on identifying the macro-element foundation parameters of an offshore wind turbine by using a Bayesian model updating approach for this task. This was addressed in Chapter 3, where a deterministic model updating framework was compared against two Bayesian frameworks: one using a vector-based likelihood function and another using a MAC-based likelihood.

Numerical case studies showed that all three frameworks could identify the compliance parameters of the macro-element foundation. However, the Bayesian frameworks offered clear advantages: through quantification of both parameter and model uncertainty, they provide greater interpretability for the user. Furthermore, the deterministic objective function and Bayesian posterior distributions were often multi-modal, highlighting the importance of global

optimisation methods—another strong motivation for using Bayesian frameworks.

A secondary aim of Work Package 1 was to investigate how uncertainty in estimated modal parameters should be incorporated as measurement error within the updating process. Chapter 3 demonstrated that uncertainties obtained from stochastic subspace identification varied significantly between datasets and did not always improve the results. In contrast, defining measurement uncertainty using a frequentist approach, based on variability across multiple datasets, resulted in more consistent and reliable parameter estimates. However, it should be noted that this numerical investigation generated datasets under the same environmental and operational conditions. In practical applications, frequentist uncertainty would also reflect variability arising from changing environmental and operational states.

Overall, Work Package 1 provided an improved understanding of the benefits and limitations of Bayesian model updating for foundation parameter identification. Its key limitation was the assumption of a perfect structural model, which motivated the investigations in Work Package 2.

WP 2: Impact of RNA scaling when estimating foundation parameters

Key Recommendation: *When model updating with a flexible rotor nacelle assembly, interpretation of the identified foundation parameters should be approached with caution and a Bayesian framework is essential for supporting this task.*

In Chapter 4, the identified compliance parameters were found to be highly sensitive to errors introduced when scaling the flexible blade properties from reference turbines. Neither the geometric scaling nor the blade-mass scaling methods could accurately reproduce the structure's modal parameters. These modelling biases led to significant errors in the estimated compliance parameters, undermining the reliability of the results.

Furthermore, the identified foundation parameters were highly sensitive to the choice of modal parameters used within the updating framework. The first pair of tower modes were least affected by the blade model errors and should always be included. In contrast, the second and third pairs of tower modes were the most susceptible to blade model error and excluding them improved identification results. Higher-order modes were less sensitive to blade scaling, but are

typically more challenging to estimate accurately from ambient vibration data in practice.

In line with expectations, the blade-mass scaling approach proved more robust than geometric scaling, as it preserved lumped mass properties. Furthermore, scaling from a reference turbine of similar size to the target turbine is recommended, as the blades are more likely to share similar structural properties and design philosophy.

Finally, Chapter 4 also proposed the new γ-$\boldsymbol{d}$ parameterisation of the macro-element, which introduces a correlation between the compliance parameters. This parameterisation assigns a more informative prior and was found to reduce the influence of the blade model error on the updating results.

In general, this work package highlights that, without reliable information on rotor and blade properties, the accuracy of model updating for estimating foundation parameters can be significantly compromised. This provides evidence for the need to promote greater openness and accessibility of turbine design data within the industry.

WP 3: Substructuring-based methods for parameter-input-state estimation

Key Recommendation: *If sufficient measurements are available, substructuring methods are recommended to alleviate the modelling challenges highlighted in Work Package 2, and to enable more reliable parameter–input–state estimation.*

This work package demonstrated the potential for substructuring as a means to bypass significant modelling errors and support structural health monitoring. Chapter 5 showed that both the soil-structure interaction and the rotor nacelle assembly could be discarded to enable robust input-state estimation with an unconstrained substructure. In this numerical case study, at least one displacement sensor was required to estimate the rigid-body motion. However, fatigue estimation does not rely on the these rigid-body modes being observable and could still be estimated accurately using more commonly available sensors such as inclinometers and accelerometers.

Chapter 6 extended the substructuring framework for parameter estimation, proposing a new method that constructs a time-domain likelihood function from input-state estimation re-

sults. Numerical case studies showed that the individual compliance parameters were practically unidentifiable due to the measurements lack of sensitivity to the lateral spring. However, a single scale parameter, γ, could be identified accurately using this framework. In addition, the identifiability of γ was shown to depend on the underlying foundation stiffness.

WP 4: Application to field data

Key Recommendation: *The identifiability of foundation parameters depends on the underlying foundation stiffness, which should be assessed beforehand to evaluate the feasibility of parameter estimation.*

This work package demonstrated the practical applicability of the proposed structural health monitoring methods for estimating the macro-element foundation parameters. In general, the observations were consistent with those from the numerical case studies. The Bayesian model updating framework based on modal data, introduced in Chapter 3, was found to depend heavily on a number of user choices. Additionally, the Simulink Model used for this method exhibited significant errors, likely due to inaccuracies introduced by scaling the blade properties. As a result, the substructuring-based approach, introduced in Chapter 6 was thought to be more reliable, resulting in more consistent parameter estimates.

However, this real-world application likely considered a structure with a relatively stiff foundation, meaning the dynamic response was less sensitive to the foundation parameters. As a result, the lateral spring could not be identified accurately from the available measurements. Nonetheless, the γ parameterisation was estimated well, and the updated value (twice the initial design stiffness) led to improved model predictions.

8.2 Future Work

8.2.1 Nonlinear and dynamic soil modelling

A key assumption throughout this thesis was that the soil-structure interaction could be represented using a linear macro-element stiffness matrix. While this assumption may be reasonable under small-displacement conditions, further research is needed to assess the influence of nonlinear and hysteric soil behaviour, which may become significant under extreme loading or over

long-term monitoring campaigns.

Additionally, the potential contributions of soil damping and the mass of the foundation—which were both neglected here—should be investigated to understand their impact on the structure's dynamic behaviour and whether they can be reliably identified from ambient vibration data.

It is worth noting that some of these assumptions were partially addressed (or avoided) in Chapter 5, where an unconstrained substructure model was considered and the ground-level response was estimated. Therefore, this approach could be further developed, to account for more complex soil behaviour.

8.2.2 Experimental validation

Although the practical applicability of some of the developed methods was demonstrated using a real-world case study in Chapter 7, experimental validation under controlled conditions would greatly increase confidence in the results. In particular, the unconstrained substructure model could not be applied to field data due to the limited sensor setup available. Therefore, controlled laboratory experiments would provide a valuable opportunity to test and validate this approach more rigorously.

8.2.3 Optimal sensor placement studies

A recurring conclusion from many of the chapter throughout this thesis was that the performance of the developed methods—and the identifiability of the foundation parameters—depended on the available sensor setup. A valuable extension for future work would be to investigate optimal sensor placement strategies, with the aim of achieving more reliable and robust results for identifying the foundation parameters.

In addition, the structure's relative foundation stiffness was found to influence the observability and identifiability of the system. While this effect was briefly explored in Chapter 6 in the context of substructuring-based parameter estimation, further investigation is needed, particularly in parallel with optimal sensor placement studies.

8.2.4 Damage detection and changing structural behaviour

This thesis focused on parameter estimation under the assumption that the structure remained in a healthy state. An interesting direction for future work could be to explore methods capable of capturing damage progression. In particular, scour is often a concern for monopile-supported offshore wind turbines, and the development of methods to detect this from ambient vibration data could offer significant value to wind farm operators. Numerical case studies could be designed to simulate progressive damage and assess the resulting sensitivity in the dynamic response. Similarly, damage could be introduced elsewhere in the structure to investigate whether it can be reliably identified and localised.

Furthermore, all analyses were carried out for idling or parked conditions. However, the foundation stiffness and dynamic behaviour may vary under different operational or environment conditions. Exploring this variability is essential for developing robust monitoring techniques suitable for continuous, long-term deployment. Hierarchical methods offer promising potential in this context, as they can account for such variability across multiple dataset and allow for more generalisable and adaptive monitoring frameworks.

Bibliography

1. IPCC. in *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (eds Core Writing Team, H. L. & Romero, J.) 35–115 (IPCC, Geneva, Switzerland, 2023).
2. National Oceanic and Atmospheric Administration. *Climate Change Impacts* Accessed: 2025-07-14. `https://www.noaa.gov/education/resource-collections/climate/climate-change-impacts`.
3. L. Joyce & F. Zhao. *Global Wind Energy Report 2022* 2022.
4. R. WIlliams, F. Zhao, N. Melkonyan, M. Hutchinson & J. Cheong. *Global Offshore Wind Report 2025* 2025.
5. G. Costanzo, G. Brindley & P. Tardieu. *Wind energy in Europe: 2024 Statistics and the outlook for 2025-2030* tech. rep. (Wind Europe, 2025).
6. International Renewable Energy Agency (IRENA). *Renewable Power Generation Costs in 2023* Accessed: 2025-07-14. Abu Dhabi, 2024.
7. D. Kallehave, B. W. Byrne, C. LeBlanc Thilsted & K. K. Mikkelsen. Optimization of monopiles for offshore wind turbines. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **373,** 20140100 (2015).
8. J. Jonkman, S. Butterfield, W. Musial & G. Scott. *Definition of a 5-MW reference wind turbine for offshore system development* tech. rep. (National Renewable Energy Lab.(NREL), Golden, CO (United States), 2009).
9. P. Bortolotti *et al. IEA Wind TCP Task 37: Systems engineering in wind energy-WP2. 1 Reference wind turbines* tech. rep. (National Renewable Energy Lab.(NREL), Golden, CO (United States), 2019).
10. E. Gaertner *et al. IEA wind TCP task 37: definition of the IEA 15-megawatt offshore reference wind turbine* tech. rep. (National Renewable Energy Lab.(NREL), Golden, CO (United States), 2020).
11. C. R. Farrar & K. Worden. *Structural health monitoring: a machine learning perspective* (John Wiley & Sons, 2012).
12. L. R. Barroso & R. Rodriguez. Damage detection utilizing the damage index method to a benchmark structure. *Journal of Engineering Mechanics* **130,** 142–151 (2004).
13. A. Rasheed, O. San & T. Kvamsdal. Digital twin: Values, challenges and enablers from a modeling perspective. *IEEE access* **8,** 21980–22012 (2020).
14. J. S. Nielsen & J. D. Sørensen. Bayesian estimation of remaining useful life for wind turbine blades. *Energies* **10,** 664 (2017).
15. R. Pacheco-Blazquez, J. Garcia-Espinosa, D. Di Capua & A. Pastor Sanchez. A Digital Twin for Assessing the Remaining Useful Life of Offshore Wind Turbine Structures. *Journal of Marine Science and Engineering* **12,** 573 (2024).
16. K. Sunday & F. Brennan. Influence of soil–structure modelling techniques on offshore wind turbine monopile structural response. *Wind Energy* **25,** 998–1012 (2022).
17. F. S. Wong. Uncertainties in dynamic soil-structure interaction. *Journal of Engineering Mechanics* **110,** 308–324 (1984).
18. F. Nadim. in *Geotechnical safety and risk V* 23–35 (IOS Press, 2015).
19. R. A. McAdam, M. N. Chatzis, M. Kuleli, E. F. Anderson & B. W. Byrne. Monopile Foundation Stiffness Estimation of an Instrumented Offshore Wind Turbine through Model Updating. *Structural Control and Health Monitoring* **2023** (2023).
20. P. Martakis, D. Taeseri, E. Chatzi & J. Laue. A centrifuge-based experimental verification of Soil-Structure Interaction effects. *Soil Dynamics and Earthquake Engineering* **103,** 1–14 (2017).
21. E. F. Anderson. *Identification of Offshore Wind Turbine Foundation Properties from Monitoring Data* DPhil thesis (University of Oxford, 2022). Unpublished.
22. L. Ramirez, D. Fraile & G. Brindley. *Offshore Wind in Europe: Key trends and statistics 2020* windeurope.org (2021).
23. Y. Shkara, R. Schelenz & G. Jacobs. *The effect of blade-tower interaction on the structure loading of multi megawatt horizontal axis wind turbine* in *Journal of Physics: Conference Series* **1037** (2018), 072033.
24. NREL. *OpenFAST*
25. DTU. *Horizontal Axis Wind turbine simulation Code 2nd generation (HAWC2*

26. DNV. *Wind turbine design software - Bladed*
27. J. Jonkman & W. Musial. *Offshore code comparison collaboration (OC3) for IEA Wind Task 23 offshore wind technology and deployment* tech. rep. (National Renewable Energy Lab.(NREL), Golden, CO (United States), 2010).
28. J. Meng. *Dynamic Modelling of Floating Wind Turbines* DPhil thesis (University of Oxford, 2025). Unpublished.
29. J. Meng, R. A. McAdam & M. N. Chatzis. *A Simulink Model for the Dynamic Analysis of Floating Wind Turbines* in *"Dynamics of Civil Structures, Volume 2": Conference Proceedings of the Society for Experimental Mechanics Series.* (2024), 229–239.
30. G. Murphy, D. Igoe, P. Doherty & K. Gavin. 3D FEM approach for laterally loaded monopile design. *Computers and Geotechnics* **100,** 76–83 (2018).
31. K. Kaltekis *et al.* A review of constitutive models for an FEA-based monopile design. *SUT Offshore Site Investigation and Geotechnics*, SUT–OSIG (2023).
32. API. *Recommended Practice for Planning, Designing & Constructing Fixed Offshore Platforms* tech. rep. API-RP.2A-WSD (American Petroleum Institute, 2010).
33. Det Norske Veritas (DNV). *DNVGL-ST-0126 - support structure for wind turbines.* Oslo, Norway: Det Norske Veritas. (2016).
34. L. C. Reese & H. Matlock. *Non-dimensional Solutions for Laterally-loaded Piles with Soil Modulus Assumed Proportional to Depth* (Association of Drilled Shaft Contractors, 1956).
35. H. Matlock. *Correlation for Design of Laterally Loaded Piles in Soft Clay.* Paper presented at the Offshore Technology Conference, Houston, Texas. 1970.
36. W. R. Cox, L. C. Reese & G. B. R. *Field Testing of Laterally Loaded Piles In Sand.* Paper presented at the Offshore Technology Conference, Houston, Texas. 1974.
37. L. C. Reese, W. R. Cox & F. D. Koop. *Analysis of Laterally Loaded Piles in Sand.* Paper presented at the Offshore Technology Conference, Houston, Texas. 1974.
38. P. Doherty & K. Gavin. Laterally loaded monopile design for offshore wind farms. *Proceedings of the Institution of Civil Engineers - Energy* **165,** 7–17 (2012).
39. B. W. Byrne *et al.* New design methods for large diameter piles under lateral loading for offshore wind applications. *Frontiers in offshore geotechnics III* **1,** 705–710 (2015).
40. B. W. Byrne *et al.* PISA design model for monopiles for offshore wind turbines: application to a stiff glacial clay till. *Géotechnique* **70,** 1030–1047 (2020).
41. H. J. Burd *et al.* PISA design model for monopiles for offshore wind turbines: application to a marine sand. *Géotechnique* **70,** 1048–1066 (2020).
42. D. M. Taborda *et al.* Finite-element modelling of laterally loaded piles in a dense marine sand at Dunkirk. *Géotechnique* **70,** 1014–1029 (2020).
43. L. Zdravković *et al.* Finite-element modelling of laterally loaded piles in a stiff glacial clay till at Cowden. *Géotechnique* **70,** 999–1013 (2020).
44. H. J. Burd *et al.* New data analysis methods for instrumented medium-scale monopile field tests. *Géotechnique* **70,** 961–969 (2020).
45. B. W. Byrne *et al.* Monotonic laterally loaded pile testing in a stiff glacial clay till at Cowden. *Géotechnique* **70,** 970–985 (2020).
46. R. A. McAdam *et al.* Monotonic laterally loaded pile testing in a dense marine sand at Dunkirk. *Géotechnique* **70,** 986–998 (2020).
47. L. Zdravković *et al.* Ground characterisation for PISA pile testing and analysis. *Géotechnique* **70,** 945–960 (2020).
48. B. W. Byrne *et al. Cyclic laterally loaded medium scale field pile testing for the PISA project* in *4th International Symposium on Frontiers in Offshore Geotechnics (postponed)* (2020), 1323–1332.
49. H. J. Burd *et al.* Application of the PISA design model to monopiles embedded in layered soils. *Géotechnique* **70,** 1067–1082 (2020).
50. S. Martin *et al. Medium-scale field testing of laterally loaded monopiles in glacial till at Cowden* in *SUT Offshore Site Investigation and Geotechnics* (2023), SUT–OSIG.
51. I. B. Løken & A. M. Kaynia. Effect of foundation type and modelling on dynamic response and fatigue of offshore wind turbines. *Wind Energy* **22,** 1667–1683 (2019).
52. K. S. Skau *et al. REDWIN–REDucing cost in offshore WINd by integrated structural and geotechnical design* in *Journal of Physics: Conference Series* **1104** (2018), 012029.

53. A. M. Page, K. S. Skau, H. P. Jostad & G. R. Eiksund. A new foundation model for integrated analyses of monopile-based offshore wind turbines. *Energy Procedia* **137,** 100–107 (2017).
54. R. Bergua *et al.* OC6 Phase II: Integration and verification of a new soil–structure interaction model for offshore wind design. *Wind Energy* **25,** 793–810 (2022).
55. E. Anderson, R. McAdam & M. N. Chatzis. The Effects of Non-Proportional Damping on the Identification of Offshore Wind Turbine Foundation Properties. *Structural Control and Health Monitoring (accepted)* (2024).
56. P. Van Overschee & B. De Moor. *Subspace identification for linear systems* 1996.
57. E. Reynders. System identification methods for (operational) modal analysis: review and comparison. *Archives of Computational Methods in Engineering* **19,** 51–124 (2012).
58. B. J. Schwarz & M. H. Richardson. Experimental modal analysis. *CSI Reliability week* **35,** 1–12 (1999).
59. K. F. Abdulraheem & G. Al-Kindi. A Simplified wind turbine blade crack identification using Experimental Modal Analysis (EMA). *International Journal of Renewable Energy Research* **7,** 715–722 (2017).
60. E. Reynders, R. Pintelon & G. De Roeck. Uncertainty bounds on modal parameters obtained from stochastic subspace identification. *Mechanical systems and signal processing* **22,** 948–969 (2008).
61. E. Reynders, K. Maes, G. Lombaert & G. De Roeck. Uncertainty quantification in operational modal analysis with stochastic subspace identification: Validation and applications. *Mechanical Systems and Signal Processing* **66,** 13–30 (2016).
62. E. P. Reynders. Uncertainty quantification in data-driven stochastic subspace identification. *Mechanical Systems and Signal Processing* **151,** 107338 (2021).
63. K. Motte, W. Weijtjens, C. Devriendt & P. Guillaume. *Operational modal analysis in the presence of harmonic excitations: a review* in *Dynamics of Civil Structures, Volume 2: Proceedings of the 33rd IMAC, A Conference and Exposition on Structural Dynamics, 2015* (2015), 379–395.
64. L. D. Avendano-Valencia, E. N. Chatzi & D. Tcherniak. Gaussian process models for mitigation of operational variability in the structural health monitoring of wind turbines. *Mechanical Systems and Signal Processing* **142,** 106686 (2020).
65. Y. Ou, E. N. Chatzi, V. K. Dertimanis & M. D. Spiridonakos. Vibration-based experimental damage detection of a small-scale wind turbine blade. *Structural Health Monitoring* **16,** 79–96 (2017).
66. S. Vettori, G. Gomes, E. Di Lorenzo, B. Peeters & E. Chatzi. *Influence of the input model for virtual sensing of wind turbine blades* in *30th International Conference on Noise and Vibration Engineering (ISMA 2022) in conjunction with the 9th International Conference on Uncertainty in Structural Dynamics (USD 2022), Leuven, Belgium, September 12-14, 2022* (2023).
67. G. Oliveira, F. Magalhães, Á. Cunha & E. Caetano. Continuous dynamic monitoring of an onshore wind turbine. *Engineering Structures* **164,** 22–39 (2018).
68. D. Tcherniak & G. C. Larsen. *Application of OMA to an Operating Wind Turbine: now including Vibration Data from the Blades* in *5th International Operational Modal Analysis Conference* (2013).
69. C. Devriendt, M. Elkafafy, G. De Sitter, P. J. Jordaens & P. Guillaume. Continuous dynamic monitoring of an offshore wind turbine on a monopile foundation. *ISMA2012* (2012).
70. P. Zhang, Z. He, C. Cui, L. Ren & R. Yao. Operational Modal Analysis of Offshore Wind Turbine Tower under Ambient Excitation. *Journal of Marine Science and Engineering* **10,** 1963 (2022).
71. H.-C. Kim, M.-H. Kim & D.-E. Choe. Structural health monitoring of towers and blades for floating offshore wind turbines using operational modal analysis and modal properties with numerical-sensor signals. *Ocean Engineering* **188,** 106226 (2019).
72. X. Chen *et al.* Method for monitoring scour depth of pile foundations based on modal identification. *Smart Materials and Structures* **30,** 075008 (2021).
73. R. Shirzadeh, C. Devriendt, M. A. Bidakhvidi & P. Guillaume. Experimental and computational damping estimation of an offshore wind turbine on a monopile foundation. *Journal of Wind Engineering and Industrial Aerodynamics* **120,** 96–106 (2013).
74. A. Bajric, R. Brincker & S. Thöns. *Evaluation of damping estimates in the presence of closely spaced modes using operational modal analysis techniques* in *Proceedings of the 6th international operational modal analysis conference, Gijòn, Spain* **735** (2015), 449–461.
75. S. Chauhan, A. Phillips & R. Allemang. *Damping estimation using operational modal analysis* in *Proceedings of the 26th International Modal Analysis Conference, Orlando, Florida* (2008).
76. M. I. Friswell & J. Mottershead. *Finite element model updating in structural dynamics* (Kluwer Academic Publishers, 1995).
77. E. Simoen, G. De Roeck & G. Lombaert. Dealing with uncertainty in model updating for damage assessment: A review. *Mechanical Systems and Signal Processing* **56,** 123–149 (2015).

78. D. Augustyn, U. Smolka, U. T. Tygesen, M. D. Ulriksen & J. D. Sørensen. Data-driven model updating of an offshore wind jacket substructure. *Applied Ocean Research* **104,** 102366 (2020).
79. M. Henkel. *Validation of virtual sensing for the fatigue assessment of subsoil and submerged components of offshore wind turbines* PhD thesis (Vrije Universiteit Brussel, 2022).
80. C. Papadimitriou. Optimal sensor placement methodology for parametric identification of structural systems. *Journal of sound and vibration* **278,** 923–947 (2004).
81. C. Papadimitriou & L. S. Katafygiotis. in *Engineering design reliability handbook* 525–544 (CRC Press, 2004).
82. E. Simoen, C. Papadimitriou & G. Lombaert. On prediction error correlation in Bayesian model updating. *Journal of Sound and Vibration* **332,** 4136–4152 (2013).
83. C. Argyris, C. Papadimitriou, P. Panetsos & P. Tsopelas. Bayesian model-updating using features of modal data: Application to the metsovo bridge. *Journal of Sensor and Actuator Networks* **9,** 27 (2020).
84. I. Behmanesh, B. Moaveni, G. Lombaert & C. Papadimitriou. Hierarchical Bayesian model updating for structural identification. *Mechanical Systems and Signal Processing* **64,** 360–376 (2015).
85. E. Simoen, B. Moaveni, J. P. Conte & G. Lombaert. Uncertainty quantification in the assessment of progressive damage in a 7-story full-scale building slice. *Journal of Engineering Mechanics* **139,** 1818–1830 (2013).
86. B. Moynihan *et al.* System identification and finite element model updating of a 6 MW offshore wind turbine using vibrational response measurements. *Renewable Energy* **219,** 119430 (2023).
87. X. Shi & M. Chatzis. An efficient algorithm to test the observability of rational nonlinear systems with unmeasured inputs. *Mechanical Systems and Signal Processing* **165,** 108345 (2022).
88. W. E. Walker *et al.* Defining uncertainty: a conceptual basis for uncertainty management in model-based decision support. *Integrated assessment* **4,** 5–17 (2003).
89. J. Beck & L. Katafygiotis. Updating Models and Their Uncertainties. *I-Baysian Statistical Framework* (1998).
90. L. S. Katafygiotis & J. L. Beck. Updating models and their uncertainties. II: Model identifiability. *Journal of Engineering Mechanics* **124,** 463–467 (1998).
91. K. Tatsis, E. Chatzi & E.-M. Lourens. *Reliability prediction of fatigue damage accumulation on wind turbine support structures* in *Proceedings of the 2nd International Conference on Uncertainty Quantification in Computational Sciences and Engineering* (2017), 76–89.
92. B. Moaveni, J. P. Conte & F. M. Hemez. Uncertainty and sensitivity analysis of damage identification results obtained using finite element model updating. *Computer-Aided Civil and Infrastructure Engineering* **24,** 320–334 (2009).
93. K. Prajapat & S. Ray-Chaudhuri. Prediction error variances in Bayesian model updating employing data sensitivity. *Journal of Engineering Mechanics* **142,** 04016096 (2016).
94. R. J. Allemang. *A correlation coefficient for modal vector analysis* in *Proc. of the 1st IMAC* (1982), 110–116.
95. P. Xu, J. Chen, J. Li, S. Fan & Q. Xu. Using Bayesian updating for monopile offshore wind turbines monitoring. *Ocean Engineering* **280,** 114801 (2023).
96. M. Song, I. Behmanesh, B. Moaveni & C. Papadimitriou. Modeling error estimation and response prediction of a 10-story building model through a hierarchical Bayesian model updating framework. *Frontiers in Built Environment* **5,** 7 (2019).
97. M. Song, I. Behmanesh, B. Moaveni & C. Papadimitriou. Accounting for modeling errors and inherent structural variability through a hierarchical bayesian model updating approach: an overview. *Sensors* **20,** 3874 (2020).
98. M. Song, B. Moaveni, C. Papadimitriou & A. Stavridis. Accounting for amplitude of excitation in model updating through a hierarchical Bayesian approach: Application to a two-story reinforced concrete building. *Mechanical Systems and Signal Processing* **123,** 68–83 (2019).
99. X. Jia, O. Sedehi, C. Papadimitriou, L. S. Katafygiotis & B. Moaveni. Hierarchical Bayesian modeling framework for model updating and robust predictions in structural dynamics using modal features. *Mechanical Systems and Signal Processing* **170,** 108784 (2022).
100. O. Sedehi, C. Papadimitriou & L. S. Katafygiotis. Hierarchical Bayesian uncertainty quantification of Finite Element models using modal statistical information. *Mechanical Systems and Signal Processing* **179,** 109296 (2022).
101. M. Song, B. Moaveni & E. Hines. Hierarchical Bayesian quantification of aerodynamic effects on an offshore wind turbine under varying environmental and operational conditions. *Mechanical Systems and Signal Processing* **224,** 112174 (2025).

102. S.-K. Au, J. M. Brownjohn, B. Li & A. Raby. Understanding and managing identification uncertainty of close modes in operational modal analysis. *Mechanical Systems and Signal Processing* **147,** 107018 (2021).
103. W. D'ambrogio & A. Fregolent. Higher-order mac for the correlation of close and multiple modes. *Mechanical systems and signal processing* **17,** 599–610 (2003).
104. S. Adhikari. *Damping models for structural vibration* PhD thesis (University of Cambridge, 2001).
105. P. Brewick & A. Smyth. An investigation of the effects of traffic induced local dynamics on global damping estimates using operational modal analysis. *Mechanical Systems and Signal Processing* **41,** 433–453 (2013).
106. S. S. Rao. *Vibration of continuous systems* (John Wiley & Sons, 2019).
107. R. Fillod. Contribution on Identification of Linear Mechanical Structures. *Doctor of Sciences Physics Dissertation, University of Besançon, France* (1980).
108. P. Avitabile & P. Pingle. Prediction of full field dynamic strain from limited sets of measured data. *Shock and vibration* **19,** 765–785 (2012).
109. E. Akintunde, S. E. Azam & D. G. Linzell. Probabilistic bridge fatigue evaluation at virtual sensing locations using kernel density estimation. *International Journal of Fatigue* **176,** 107885 (2023).
110. B. Moynihan, E. M. Tronci, M. C. Hughes, B. Moaveni & E. Hines. Virtual sensing via Gaussian Process for bending moment response prediction of an offshore wind turbine using SCADA data. *Renewable Energy* **227,** 120466 (2024).
111. A. Iliopoulos *et al.* A modal decomposition and expansion approach for prediction of dynamic responses on a monopile offshore wind turbine using a limited number of vibration sensors. *Mechanical Systems and Signal Processing* **68,** 84–104 (2016).
112. M. Song, S. Christensen, B. Moaveni, A. Brandt & E. Hines. Joint parameter-input estimation for virtual sensing on an offshore platform using output-only measurements. *Mechanical Systems and Signal Processing* **170,** 108814 (2022).
113. M.-S. Nabiyan, F. Khoshnoudian, B. Moaveni & H. Ebrahimian. Mechanics-based model updating for identification and virtual sensing of an offshore wind turbine using sparse measurements. *Structural Control and Health Monitoring* **28,** e2647 (2021).
114. R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* **82,** 35–45 (1960).
115. C. Papadimitriou, C.-P. Fritzen, P. Kraemer & E. Ntotsios. Fatigue predictions in entire body of metallic structures from a limited number of vibration sensors using Kalman filtering. *Structural Control and Health Monitoring* **18,** 554–573 (2011).
116. S. Gillijns & B. De Moor. Unbiased minimum-variance input and state estimation for linear discrete-time systems. *Automatica* **43,** 111–116 (2007).
117. S. Gillijns & B. De Moor. Unbiased minimum-variance input and state estimation for linear discrete-time systems with direct feedthrough. *Automatica* **43,** 934–937 (2007).
118. E. Lourens, E. Reynders, G. De Roeck, G. Degrande & G. Lombaert. An augmented Kalman filter for force identification in structural dynamics. *Mechanical systems and signal processing* **27,** 446–460 (2012).
119. F. Naets, J. Cuadrado & W. Desmet. Stable force identification in structural dynamics using Kalman filtering and dummy-measurements. *Mechanical Systems and Signal Processing* **50,** 235–248 (2015).
120. S. E. Azam, E. Chatzi & C. Papadimitriou. A dual Kalman filter approach for state estimation via output-only acceleration measurements. *Mechanical systems and signal processing* **60,** 866–886 (2015).
121. S. P. Tchemodanova, M. Sanayei, B. Moaveni, K. Tatsis & E. Chatzi. Strain predictions at unmeasured locations of a substructure using sparse response-only vibration measurements. *Journal of Civil Structural Health Monitoring* **11,** 1113–1136 (2021).
122. K. Maes, A. Iliopoulos, W. Weijtjens, C. Devriendt & G. Lombaert. Dynamic strain estimation for fatigue assessment of an offshore monopile wind turbine using filtering and modal expansion algorithms. *Mechanical Systems and Signal Processing* **76,** 592–611 (2016).
123. K. Maes *et al. Continuous strain prediction for fatigue assessment of an offshore wind turbine using Kalman filtering techniques* in *2015 IEEE Workshop on Environmental, Energy, and Structural Monitoring Systems (EESMS) Proceedings* (2015), 44–49.
124. K. Tatsis & E. Lourens. in *Life-Cycle of Engineering Systems:* 209–216 (CRC Press, 2016).
125. S. Vettori, E. Di Lorenzo, B. Peeters, M. Luczak & E. Chatzi. An adaptive-noise Augmented Kalman Filter approach for input-state estimation in structural dynamics. *Mechanical Systems and Signal Processing* **184,** 109654 (2023).
126. E. Branlard, D. Giardina & C. S. Brown. Augmented Kalman filter with a reduced mechanical model to estimate tower loads on a land-based wind turbine: a step towards digital-twin simulations. *Wind Energy Science* **5,** 1155–1167 (2020).

127. A. Mehrjoo *et al.* Recursive Bayesian estimation of wind load on a monopile-supported offshore wind turbine using output-only measurements. *Mechanical Systems and Signal Processing* **224,** 112183 (2025).
128. N. Noppe, A. Iliopoulos, W. Weijtjens & C. Devriendt. *Full load estimation of an offshore wind turbine based on SCADA and accelerometer data* in *Journal of Physics: Conference Series* **753** (2016), 072025.
129. N. Noppe, K. Tatsis, E. Chatzi, C. Devrient & W. Weijtjens. *Fatigue stress estimation of offshore wind turbine using a Kalman filter in combination with accelerometers* in *Proceedings of International Conference on Noise and Vibration Engineering (ISMA 2018), International Conference on Uncertainty in Structural Dynamics (USD 2018)* (2018), 4693–6701.
130. P. C. Hansen. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM review* **34,** 561–580 (1992).
131. K. Tatsis, V. K. Dertimanis, C. Papadimitriou, E. Lourens & E. Chatzi. A general substructure-based framework for input-state estimation using limited output measurements. *Mechanical Systems and Signal Processing* **150,** 107223 (2021).
132. V. K. Dertimanis, E. Chatzi, S. E. Azam & C. Papadimitriou. Input-state-parameter estimation of structural systems from limited output information. *Mechanical Systems and Signal Processing* **126,** 711–746 (2019).
133. S. E. Azam, E. Chatzi, C. Papadimitriou & A. Smyth. Experimental validation of the Kalman-type filters for online and real-time state and input estimation. *Journal of vibration and control* **23,** 2494–2519 (2017).
134. M. Alvarez, D. Luengo & N. D. Lawrence. *Latent force models* in *Artificial Intelligence and Statistics* (2009), 9–16.
135. C. E. Rasmussen, C. K. Williams, *et al. Gaussian processes for machine learning* (Springer, 2006).
136. R. Nayek, S. Chakraborty & S. Narasimhan. A Gaussian process latent force model for joint input-state estimation in linear structural systems. *Mechanical Systems and Signal Processing* **128,** 497–530 (2019).
137. J. Hartikainen & S. Särkkä. *Kalman filtering and smoothing solutions to temporal Gaussian process regression models* in *2010 IEEE international workshop on machine learning for signal processing* (2010), 379–384.
138. J. Hartikainen & S. Sarkka. Sequential inference for latent force models. *arXiv preprint arXiv:1202.3730* (2012).
139. H. E. Rauch, F. Tung & C. T. Striebel. Maximum likelihood estimates of linear dynamic systems. *AIAA journal* **3,** 1445–1450 (1965).
140. J. Zou, E.-M. Lourens & A. Cicirello. Virtual sensing of subsoil strain response in monopile-based offshore wind turbines via Gaussian process latent force models. *arXiv preprint arXiv:2207.05901* (2022).
141. J. Bilbao, E.-M. Lourens, A. Schulze & L. Ziegler. Virtual sensing in an onshore wind turbine tower using a Gaussian process latent force model. *Data-Centric Engineering* **3,** e35 (2022).
142. S. Vettori, E. Di Lorenzo, B. Peeters & E. Chatzi. Assessment of alternative covariance functions for joint input-state estimation via Gaussian Process latent force models in structural dynamics. *Mechanical Systems and Signal Processing* **213,** 111303 (2024).
143. F. Naets, J. Croes & W. Desmet. An online coupled state/input/parameter estimation approach for structural dynamics. *Computer Methods in Applied Mechanics and Engineering* **283,** 1167–1188 (2015).
144. C.-B. Yun & M. Shinozuka. Identification of nonlinear structural dynamic systems. *Journal of structural mechanics* **8,** 187–203 (1980).
145. E. N. Chatzi & A. W. Smyth. The unscented Kalman filter and particle filter methods for nonlinear structural system identification with non-collocated heterogeneous sensing. *Structural Control and Health Monitoring: The Official Journal of the International Association for Structural Control and Monitoring and of the European Association for the Control of Structures* **16,** 99–123 (2009).
146. S. E. Azam, V. Dertimanis, E. Chatzi & C. Papadimitriou. *Output-only schemes for joint input-state-parameter estimation of linear systems* in *Proceedings of the 1st International Conference on Uncertainty Quantification in Computational Sciences and Engineering, UNCECOMP* (2015), 497–510.
147. V. Dertimanis, E. Chatzi, S. E. Azam & C. Papadimitriou. *Output-only fatigue prediction of uncertain steel structures* in *Vol. 2 of Proc., 8th European Workshop on Structural Health Monitoring (EWSHM 2016)* (2016), 869–879.
148. M. Impraimakis & A. W. Smyth. Input–parameter–state estimation of limited information wind-excited systems using a sequential Kalman filter. *Structural Control and Health Monitoring* **29,** e2919 (2022).
149. H. Ebrahimian, R. Astroza, J. P. Conte & C. Papadimitriou. Bayesian optimal estimation for output-only nonlinear system and damage identification of civil structures. *Structural Control and Health Monitoring* **25,** e2128 (2018).

150. T. Rogers, K. Worden & E. Cross. On the application of Gaussian process latent force models for joint input-state-parameter estimation: With a view to Bayesian operational identification. *Mechanical Systems and Signal Processing* **140,** 106580 (2020).
151. K. Maes, M. Chatzis & G. Lombaert. Observability of nonlinear systems with unmeasured inputs. *Mechanical Systems and Signal Processing* **130,** 378–394 (2019).
152. X. Shi & M. N. Chatzis. Lie symmetries of nonlinear systems with unknown inputs. *Mechanical Systems and Signal Processing* **188,** 110027 (2023).
153. M. S. Allen *et al. Substructuring in engineering dynamics* (Springer, 2020).
154. W. C. Hurty. Dynamic analysis of structural systems using component modes. *AIAA journal* **3,** 678–685 (1965).
155. S. Rubin. Improved component-mode representation for structural dynamic analysis. *AIAA journal* **13,** 995–1006 (1975).
156. J. Kullaa. Virtual sensing of structural vibrations using dynamic substructuring. *Mechanical Systems and Signal Processing* **79,** 203–224 (2016).
157. S. Vettori. *Development of virtual sensing techniques for dynamic virtualization* PhD thesis (ETH Zurich, 2023).
158. L. Caglio. *Kalman Filter-based nonlinear response estimation of structures subjected to unknown loads* PhD thesis (Technical University of Denmark, 2024).
159. T. M. Inc. *Simscape Multibody: Model and simulate multibody mechanical system.* Natick, Massachusetts, United States, 2025.
160. R. R. Craig Jr & M. C. Bampton. Coupling of substructures for dynamic analyses. *AIAA journal* **6,** 1313–1319 (1968).
161. M. O. Hansen. Aerodynamics of wind turbines. Earthscan. *James & James* **8,** 14 (2008).
162. B. Jonkman & L. Kilcher. TurbSim user's guide: version 1.06. 00. *National Renewable Energy Laboratory: Golden, CO, USA* (2012).
163. R. Dean & R. Dalrymple. *Water Wave Mechanics For Engineers And Scientists* (World Scientific Publishing Company, Singapore, 1991).
164. J. Morison, J. W. Johnson & S. A. Schaaf. The force exerted by surface waves on piles. *Journal of Petroleum Technology* **2,** 149–154 (1950).
165. K. Hasselmann *et al.* Measurements of wind-wave growth and swell decay during the Joint North Sea Wave Project (JONSWAP). *Ergaenzungsheft zur Deutschen Hydrographischen Zeitschrift, Reihe A* (1973).
166. M. J. Tucker, P. G. Challenor & D. J. T. Carter. Numerical simulation of a random sea: a common error and its effect upon wave group statistics. *Applied Ocean Research* **6,** 118–122 (1984).
167. J. Wheeler. Method for Calculating Forces Produced by Irregular Waves. *Journal of Petroleum Technology* **22,** 359–367 (1970).
168. M. A. Chella, A. Tørum & D. Myrhaug. An Overview of Wave Impact Forces on Offshore Wind Turbine Substructures. *Energy Procedia* **20,** 217–226 (2012).
169. NREL. *ROSCO. Version 2.3.0* 2021.
170. N. Abbas, D. Zalkind, L. Pao & A. Wright. A Reference Open-Source Controller for Fixed and Floating Offshore Wind Turbines. *Wind Energy Science Discussions* **2021,** 1–33 (2021).
171. E. Reynders, M. Schevenels & G. De Roeck. *MACEC 3.4: The MATLAB Toolbox for Experimental and Operational Modal Analysis* 2023.
172. T. M. Inc. *fminsearch: Find minimum of unconstrained multivariable function using derivative-free method* 2022.
173. A. Lye, A. Cicirello & E. Patelli. Sampling methods for solving Bayesian model updating problems: A tutorial. *Mechanical Systems and Signal Processing* **159,** 107760 (2021).
174. J. Ching & Y.-C. Chen. Transitional Markov chain Monte Carlo method for Bayesian model updating, model class selection, and model averaging. *Journal of engineering mechanics* **133,** 816–832 (2007).
175. W. Betz, I. Papaioannou & D. Straub. Transitional markov chain monte carlo: observations and improvements. *Journal of Engineering Mechanics* **142,** 04016016 (2016).
176. L. Fingersh, M. Hand & A. Laxson. *Wind turbine design cost and scaling model* tech. rep. (National Renewable Energy Lab.(NREL), Golden, CO (United States), 2006).
177. K. Ha, M. Bätge, D. Melcher & S. Czichon. Development and feasibility study of segment blade test methodology. *Wind Energy Science* **5,** 591–599 (2020).
178. A. Bech Abrahamsen *et al.* Method for estimating the future annual mass of decommissioned wind turbine blade material in Denmark. *Wind Energy* **27,** 165–178 (2024).

179. M. Henkel, W. Weijtjens & C. Devriendt. Fatigue stress estimation for submerged and sub-soil welds of offshore wind turbines on monopiles using modal expansion. *Energies* **14,** 7576 (2021).
180. D. De Klerk, D. J. Rixen & S. Voormeeren. General framework for dynamic substructuring: history, review and classification of techniques. *AIAA journal* **46,** 1169–1181 (2008).
181. B. Besselink *et al.* A comparison of model reduction techniques from structural dynamics, numerical mathematics and systems and control. *Journal of Sound and Vibration* **332,** 4403–4422 (2013).
182. R. H. MacNeal. A hybrid method of component mode synthesis. *Computers & Structures* **1,** 581–601 (1971).
183. R. R. Craig Jr & C.-J. Chang. A review of substructure coupling methods for dynamic analysis. *NASA. Langley Res. Center Advan. in Eng. Sci., Vol. 2* (1976).
184. S. Vettori *et al. Kalman-Based virtual sensing for improvement of service response replication in environmental tests* in *Model Validation and Uncertainty Quantification, Volume 3: Proceedings of the 38th IMAC, A Conference and Exposition on Structural Dynamics 2020* (2020), 93–106.
185. K. Maes *et al.* Design of sensor networks for instantaneous inversion of modally reduced order models in structural dynamics. *Mechanical Systems and Signal Processing* **52,** 628–644 (2015).
186. M. Matsuishi & T. Endo. Fatigue of metals subjected to varying stress. *Japan Society of Mechanical Engineers, Fukuoka, Japan* **68,** 37–40 (1968).
187. D. N. V. AS *et al.* Fatigue design of offshore steel structures. *Rev. 3* (2011).
188. L. R. Curtis & B. M. Ayyub. *A Meta-Analysis of Miner's Rule* in *2024 Annual Reliability and Maintainability Symposium (RAMS)* (2024), 1–6.
189. R. Gaiotti & B. S. Smith. P-delta analysis of building structures. *Journal of Structural Engineering* **115,** 755–770 (1989).
190. A. Hildebrandt, B. Schmidt & S. Marx. Wind-wave misalignment and a combination method for direction-dependent extreme incidents. *Ocean Engineering* **180,** 10–22 (2019).
191. X. Li, C. Zhu, Z. Fan, X. Chen & J. Tan. Effects of the yaw error and the wind-wave misalignment on the dynamic characteristics of the floating offshore wind turbine. *Ocean Engineering* **199,** 106960 (2020).
192. A. S. Verma, Z. Jiang, Z. Ren, Z. Gao & N. P. Vedvik. Effects of wind-wave misalignment on a wind turbine blade mating process: Impact velocities, blade root damages and structural SafetyAssessment. *Journal of Marine Science and Application* **19,** 218–233 (2020).
193. A. Arias. A measure of earthquake intensity. *Seismic design for nuclear plants*, 438–483 (1970).
194. K. E. Tatsis, V. K. Dertimanis & E. N. Chatzi. Sequential Bayesian inference for uncertain nonlinear dynamic systems: a tutorial. *arXiv preprint arXiv:2201.08180* (2022).
195. M. Briers, S. R. Maskell & R. Wright. *A rao-blackwellised unscented Kalman filter* in *Sixth International Conference of Information Fusion, 2003. Proceedings of the* **1** (2003), 55–61.
196. E. Konstantinidis & D. Bouris. Drag and inertia coefficients for a circular cylinder in steady plus low-amplitude oscillatory flows. *Applied Ocean Research* **65,** 219–228 (2017).
197. S. Bhattacharya, G. Nikitas, L. Arany & N. Nikitas. Soil–structure interactions for offshore wind turbines. *Engineering & Technology Reference* (2017).
198. L. Arany, S. Bhattacharya, S. Adhikari, S. Hogan & J. H. G. Macdonald. An analytical model to predict the natural frequency of offshore wind turbines on three-spring flexible foundations using two different beam models. *Soil Dynamics and Earthquake Engineering* **74,** 40–45 (2015).
199. E. Reynders, J. Houbrechts & G. De Roeck. Fully automated (operational) modal analysis. *Mechanical systems and signal processing* **29,** 228–250 (2012).
200. B. Peeters & G. De Roeck. Stochastic system identification for operational modal analysis: a review. *Journal of Dynamic Systems, Measurement, and Control* **123,** 659–667 (2001).
201. F. Magalhães & Á. Cunha. Explaining operational modal analysis with data from an arch bridge. *Mechanical systems and signal processing* **25,** 1431–1450 (2011).
202. S. Lloyd. Least squares quantization in PCM. *IEEE transactions on information theory* **28,** 129–137 (1982).
203. T. M. Inc. *Reduced Order Flexible Solid Documentation* Natick, Massachusetts, United States, 2024.
204. L. V. Andersen *et al.* Influence of characteristic-soil-property-estimation approach on the response of monopiles for offshore wind turbines. *Journal of Ocean and Wind Energy* **2,** 160–167 (2015).
205. G. Zorzi *et al.* Reliability analysis of offshore wind turbine foundations under lateral cyclic loading. *Wind Energy Science Discussions* **2019,** 1–25 (2019).

APPENDIX A

Comparison of Model-Predicted Modal Parameters for Chapter 4

Tables A.1 and A.2 present the frequency errors and modal assurance criterion (MAC) values of the model-predicted modal parameters for each blade-scaled model listed in Table 4.2. In each case, the macro-element foundation is defined using the *Maximum a Posteriori* (MAP) estimated compliance parameters. The frequencies and MAC values are computed relative to the baseline model (see Section 4.2), where the foundation is defined using $\mathbf{C}^0_{\text{macro}}$. Overall, the frequency errors and MAC values were found to be generally correlated with the errors in the MAP-estimated compliance parameters shown in Table 4.7a.

Table A.1: Comparison of MAP model-predicted modal parameters from each blade-scaled model (see Table 4.2) against those from the baseline model, based on updating with the **full mode-set**.

		MAP Frequency errors [%]							
Mode	**Baseline Frequency [Hz]**	**Compliance Parameterisation**				**γ-d Parameterisation**			
		Geom5	Geom10	BM5	BM10	Geom5	Geom10	BM5	BM10
SS1	0.155	-4.132	-7.177	5.054	-4.600	-10.190	-7.253	-1.504	-3.758
FA1	0.156	-2.714	-6.752	5.694	-4.378	-9.140	-6.826	-1.097	-3.514
SS2	1.014	8.602	-17.883	7.775	-16.077	4.384	-17.019	3.898	-14.834
FA2	1.045	37.134	1.923	32.864	2.103	22.082	2.090	19.689	2.237
FA3	1.937	24.527	-15.587	27.520	-13.466	-2.456	-14.250	-1.448	-12.329
SS3	2.298	26.723	-17.422	26.735	-15.300	2.194	-15.831	0.972	-13.942
FA4	5.372	2.966	-1.748	2.850	-1.243	1.420	-1.787	0.849	-1.394
FA5	7.269	-2.066	-1.227	0.498	0.274	-2.430	-1.952	-0.302	-0.507
FA6	9.954	-3.701	-1.205	-0.806	-0.360	-3.900	-3.760	-1.590	-2.618
FA7	12.311	-0.249	-0.101	0.138	0.029	-0.262	-0.511	0.086	-0.340
FA8	14.176	-1.833	-0.180	-0.231	0.250	-2.006	-2.040	-0.451	-1.295
		Mode Shape MAC Values [-]							
Mode	**Baseline Frequency [Hz]**	**Compliance Parameterisation**				**γ-d Parameterisation**			
		Geom5	Geom10	BM5	BM10	Geom5	Geom10	BM5	BM10
SS1	0.155	0.998	1.000	0.998	0.999	0.999	1.000	0.999	1.000
FA1	0.156	0.998	0.999	0.999	0.999	1.000	1.000	1.000	1.000
SS2	1.014	0.676	0.998	0.683	0.998	0.922	0.998	0.954	0.998
FA2	1.045	0.994	0.005	0.995	0.004	0.925	0.010	0.949	0.007
FA3	1.937	0.995	0.997	0.993	0.998	0.996	0.943	0.997	0.954
SS3	2.298	0.959	0.996	0.951	0.993	0.905	0.970	0.906	0.973
FA4	5.372	0.998	0.994	0.996	0.997	0.992	0.994	0.997	0.996
FA5	7.269	0.995	1.000	0.998	0.998	0.998	0.999	1.000	0.995
FA6	9.954	0.970	1.000	0.999	0.999	0.971	0.986	1.000	0.987
FA7	12.311	0.995	1.000	0.999	1.000	0.993	0.906	0.995	0.936
FA8	14.176	0.993	0.998	1.000	1.000	0.994	0.996	1.000	0.996

Table A.2: Comparison of MAP model-predicted modal parameters from each blade-scaled model (see Table 4.2) against those from the baseline model, based on updating with the **reduced mode-set**.

Mode	Baseline Frequency [Hz]	MAP Frequency errors [%]							
		Compliance Parameterisation				γ-d Parameterisation			
		Geom5	Geom10	BM5	BM10	Geom5	Geom10	BM5	BM10
SS1	0.155	-4.213	-1.001	-0.201	-0.564	-9.828	-3.726	-0.608	-0.555
FA1	0.156	-2.803	-0.329	0.249	-0.223	-8.758	-3.167	-0.172	-0.213
FA4	5.372	1.832	-0.500	1.396	-0.157	1.625	-0.386	1.243	-0.165
FA5	7.269	-2.488	-0.413	0.070	1.102	-2.309	-0.174	-0.001	1.094
FA6	9.954	-4.159	-0.791	-0.742	0.128	-3.472	-0.544	-0.786	0.121
FA7	12.311	-0.234	-0.115	0.169	0.013	-0.225	-0.121	0.175	0.012
FA8	14.176	-1.612	-0.369	-0.178	-0.004	-1.820	-0.518	-0.134	-0.002
Mode	**Baseline Frequency [Hz]**	**Mode Shape MAC Values [-]**							
		Compliance Parameterisation				**γ-d Parameterisation**			
		Geom5	Geom10	BM5	BM10	Geom5	Geom10	BM5	BM10
SS1	0.155	0.998	1.000	0.999	1.000	0.999	1.000	0.999	1.000
FA1	0.156	0.999	1.000	1.000	1.000	1.000	1.000	1.000	1.000
FA4	5.372	0.995	1.000	0.999	1.000	0.994	1.000	0.998	1.000
FA5	7.269	0.999	0.999	1.000	1.000	0.997	0.999	1.000	1.000
FA6	9.954	0.973	1.000	0.999	1.000	0.969	1.000	0.999	1.000
FA7	12.311	0.997	1.000	1.000	1.000	0.998	1.000	1.000	1.000
FA8	14.176	0.991	0.999	1.000	1.000	0.993	0.999	1.000	1.000

APPENDIX B

Tabulated Performance Metrics for Case Study 1 from Chapter 5

Table B.1 presents the performance metrics ($TRAC$, $FRAC$ and MAE) for Case Study 1 from Chapter 5. Metrics for selected displacements and rotations are compared in both the fore-aft and side-side directions.

Table B.1: $TRAC$, $FRAC$, and MAE for selected degrees of freedom in the fore-aft and side-side directions from Case Study 1 in Chapter 5.

Sensor Set-up	Metric	State Identifier							
		$\texttt{x0}_{\texttt{FA}}$	$\texttt{x0}_{\texttt{SS}}$	$\theta\texttt{0}_{\texttt{FA}}$	$\theta\texttt{0}_{\texttt{SS}}$	$\texttt{x15}_{\texttt{FA}}$	$\texttt{x15}_{\texttt{SS}}$	$\theta\texttt{15}_{\texttt{FA}}$	$\theta\texttt{15}_{\texttt{SS}}$
Dense: x-a	$TRAC$	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	$FRAC$	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	MAE	8e−6	1e−6	2e−6	2e−7	1e−5	1e−6	1e−6	2e−7
Dense: θ-a	$TRAC$	0.909	0.385	1.000	1.000	0.902	0.380	1.000	1.000
	$FRAC$	0.990	0.751	1.000	1.000	0.991	0.749	1.000	1.000
	MAE	0.300	5.346	2e−6	2e−7	0.300	5.346	1e−6	3e−7
Sparse: θ-a	$TRAC$	0.746	0.430	1.000	1.000	0.750	0.425	1.000	1.000
	$FRAC$	0.980	0.785	1.000	1.000	0.979	0.783	1.000	1.000
	MAE	13.300	27.064	6e−7	2e−7	13.300	27.064	2e−6	5e−7