



Coláiste na Tríonóide, Baile Átha Cliath
Trinity College Dublin

Ollscoil Átha Cliath | The University of Dublin

MACHINE LEARNING-ASSISTED DYNAMIC MODELLING OF WIND TURBINE BLADES FOR RESPONSE PREDICTION

SHUBHAM SHANKARSINGH BAISTHAKUR



A thesis submitted in partial fulfilment of the requirements for the degree of
Doctor Of Philosophy

in the

Department of Civil, Structural and Environmental Engineering
Trinity College Dublin, the University of Dublin

Ireland

March, 2025

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent / do not consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).

Signature	S.S.Baisthakur
	Shubham Baisthakur
	Civil, Struct. & Env. Engg.
	Trinity College Dublin
Date	30/04/2024

Machine Learning-Assisted Dynamic Modelling of Wind Turbine Blades for Response Prediction

Shubham Baisthakur

20331322

Abstract

The wind energy sector has seen significant growth in the last few decades in terms of the number of wind turbines constructed and the scale of these structures. Multimegawatt wind turbines on scales previously unimaginable are now routinely installed. Researchers are relentlessly pushing the boundaries of such structures, developing turbines with taller towers and longer blades to harness higher wind speeds and maximise energy capture. The increasing scale of these structures presents additional challenges in the form of reduced stiffness, increased loading, and larger deformations. In addition, to reduce the levelised cost of energy, all components of wind turbine systems are required to be light-weight and cost-effective but simultaneously satisfy performance and reliability requirements. These constraints require a more detailed analysis of the design, reliability, and fatigue of these structures.

The growing complexity of wind turbine designs, driven by the increasing scale of turbines, introduces significant challenges related to the stochastic nature of external loading, such as varying wind speeds and turbulence intensity. These environmental factors govern the dynamic loading on wind turbines, leading to variations in load distributions and uncertainty in achieving optimal performance. Since this uncertainty cannot be reduced, but must be quantified, traditional methods of analysis, which rely on static load envelopes, often fail to ensure accurate performance predictions across different installation sites. To address this challenge, a more dynamic, site-specific approach is needed, where the unique loading conditions at each installation location should be used for performance analysis. However, the high computational cost of such simulations, which involves modelling complex aerodynamic loads and deformations, makes this approach computationally challenging for large-scale use. This thesis aims to alleviate this issue by developing computationally efficient surrogate models to emulate the dynamic response of wind turbine.

In this thesis, a multi-body dynamic (MBD) model of a wind turbine is developed to model the onshore variant of the International Energy Agency's (IEA) 15-MW wind turbine. The high computational cost of performing dynamic analysis of wind turbines subjected to turbulent wind inflow via the MBD model stems

from two major factors: computation of aerodynamic loads using Blade Element Momentum (BEM) theory and computing the deformation response using a time stepping-based numerical integration algorithm. To address this, an Artificial Neural Network (ANN) model is developed using a data-driven and physics-informed approach to predict the aerodynamic loads. This model bypasses the computationally expensive BEM theory for quick estimation of aerodynamic loads. The resulting ANN-based surrogate model could predict the angle of attack at investigated blade nodes with the maximum Mean Absolute Error (MAE) of the order of 10^{-2} . Building up on this, a Long Short-Term Memory (LSTM) model is developed to predict the blade deformation response when subjected to time-varying aerodynamic loads. These models combined can substantially reduce the computational cost of performing dynamic analysis on wind turbine blades. A combined wrapper and filter-based approach to identify the most informative features and a learning rate scheduling approach to optimize the model training are developed to establish an efficient workflow for developing such models. In addition, these models offer the potential to serve as virtual sensors to digitally measure the aerodynamic loads and blade deformations. The insights from the feature selection approach can also guide sensor placement strategies. Further, an LSTM model is developed to predict the response of blade degrees of freedom (DOFs) when subjected to turbulent spatio-temporal wind fields. The Principle Component Analysis (PCA) and Discrete Cosine Transform (DCT) feature extraction methods are combined with the recursive feature addition (RFA) method to reduce the high dimensionality of the wind field and arrive at optimized input features for an LSTM model. The LSTM model developed using these features results in high-accuracy prediction at individual blade DOFs. In the validation stage, the LSTM model prediction of blade response achieved high accuracy with maximum Normalised Root Mean Square Error (NRMSE) attaining values of the order of 10^{-1} . This LSTM model can predict the blade response without explicit knowledge of blade properties. This method presents a particularly promising approach for analysing the performance of commercial wind turbine models where the blade properties can not be shared due to intellectual property concerns. Overall, the models presented in this thesis can help in enabling quick estimate of aerodynamic loads and blade response, which can aid in efficient site-specific analysis, quick fatigue-life estimation and faster design exploration.

Acknowledgements

To begin with, I am immensely grateful to Dr. Breiffni Fitzgerald for providing this opportunity and the freedom to explore my interests. His invaluable guidance has been instrumental in the successful completion of this challenging PhD work. I also extend sincere thanks to my doctoral committee members, Prof. Biswajit Basu and Prof. Brian Broderick, for their valuable suggestions in ensuring I focused on the right problems. Additionally, I acknowledge the invaluable discussions with Dr. Saptarshi Sarkar, which helped shape parts of this thesis.

I would like to thank my office mates Abdul, Sudipta, Sarvanan, Manish, Suhail and Hollie for all the valuable discussions during the coffee break. Special thanks to James for being my brainstorming partner and always supplying me with delicious caramel squares. You helped me tremendously, both academically and in making Dublin feel like home. And to Madhu and Madhavan, thank you for the countless adventures and Brewdog sessions. You truly made my PhD experience unforgettable.

In my personal life, I am eternally grateful to Mrunali for her unwavering love and support and for letting my ambitions take the driving seat. Her silent sacrifices and unwavering encouragement to chase my dreams have been a constant source of strength. My beloved friend, Ajay, thank you for always being a phone call away, ready to offer help, advice, or simply a listening ear whenever I needed it. And lastly the most important people, my family. My sister, Priyanka, for being my biggest motivation and giving me strength, Pallavi for always encouraging me and Komal for keeping me grounded. And the most important of all, my parents, their efforts and encouragement to always follow my dreams have made me the person I am today.

Contents

List of Figures	v
List of Tables	ix
Glossary	xi
1 Introduction	1
1.1 Challenges and Opportunities in Wind Turbine Design	6
1.2 The Role of Surrogate Models	9
1.3 Structure of the Thesis	11
2 Literature review	13
2.1 Introduction	13
2.2 Wind Turbine modelling	16
2.2.1 Structural Models	17
2.2.2 Aeroelastic Models	17
2.3 Major challenges in Wind Turbine Design and Operations	19
2.3.1 Reliability Analysis and Design Optimisation	19
2.3.2 Fatigue Analysis of Wind Turbines	22
2.4 Overcoming Computational Barriers in Wind Turbine Design: The Role of Surrogate Models	25
2.5 Commonly used surrogate modelling techniques in wind turbine . .	27
2.5.1 Response Surface Method	27
2.5.2 Gaussian Process Regression	28
2.5.3 Polynomial Chaos Expansion	30
2.5.4 Artificial Neural Networks	31
2.6 Summary	34

3	Multi-body dynamic model of a wind turbine	35
3.1	Motivation	35
3.2	Structural Elements of Wind Turbines and Modelling Strategies . .	36
3.2.1	Modelling rigid elements	36
3.2.2	Modelling flexible elements	37
3.3	Structural Elements and Degrees of Freedom	38
3.4	Aerodynamic Modelling of Wind Turbine Blades	40
3.5	Kane's Dynamics Approach for Deriving EOM for the MBD Model	43
3.6	Inflow wind generation	47
3.7	IEA-15MW Reference Wind Turbine	48
3.8	Limitations of the MBD model	52
3.8.1	Mode Shape Summation Approximation for Flexible Elements	52
3.8.2	Use of Steady BEM theory	53
3.9	Summary	54
4	Artificial Neural Network based surrogate model to predict aero- dynamic loads	55
4.1	Motivation	55
4.2	Background	56
4.3	Blade Element Momentum theory	57
4.4	Development of machine learning surrogate for BEM model	63
4.4.1	Training Data Generation	64
4.4.2	Definition of Objective Function	67
4.4.3	Hyperparameter Tuning	69
4.5	Numerical Results	70
4.5.1	Computational Advantage of Surrogate Models	73
4.6	Discussion	75
5	Long Short-Term Memory model to predict blade deformation response	79
5.1	Motivation	79
5.2	Background	80
5.3	LSTM Algorithm	81
5.4	Parameter Selection and Training Data Generation	83
5.5	Feature Selection	85
5.5.1	MRMR Algorithm	86

5.5.2	Recursive Feature Selection	88
5.6	Hyperparameter Tuning and Model Training	89
5.7	Numerical Results	91
5.7.1	Dealing with Noisy Input Signals	108
5.7.2	Analysis and Interpretation	110
5.7.3	Computational advantage of surrogate approach	114
5.8	Discussion	115
6	Multi-task learning Long Short-Term Memory model to emulate blade dynamics	117
6.1	Motivation	117
6.2	High dimensionality of the stochastic wind inflow	118
6.3	Dimensionality Reduction Techniques for Wind Field Data	121
6.3.1	Principal Component Analysis	123
6.3.2	Discrete Cosine Transform	125
6.4	Data Generation: Selection of Input Variables	128
6.5	Selecting the output parameter	130
6.6	Numerical Results	132
6.6.1	One-stage modelling for blade response prediction	133
6.6.2	Multi-stage modelling for blade response prediction	135
6.6.3	Predicting blade response	148
6.6.4	Uncertainty Quantification in Model Predictions	155
6.6.5	Uncertainties and Future Scope	156
6.6.6	Computational advantage of surrogate approach	159
6.7	Discussion	162
7	Conclusions and future work	165
7.1	Summary	166
7.2	Limitations and Future work	168
	Bibliography	171

List of Figures

1.1	Global average temperature anomaly(Source: Berkeley Earth	2
1.2	Impact of Paris accord on fossil fuels (Source: Climate Nexus) . . .	3
1.3	Impact of Paris accord on renewable energy (Source: Climate Nexus)	4
1.4	Global greenhouse gas emission by sector (Source: EPA)	5
1.5	Global crude oil prices over last ten years (Source: Macrotrends) . .	6
1.6	Paris agreement - Current Status and Consequences (Source: MSCI)	7
1.7	Increase in scale of wind turbines (Source: KPMG)	8
2.1	Various disciplines involved in achieving efficient analysis, design and operation of wind turbines (Source:Veers et al. [16])	14
2.2	Feed forward neural network architecture	32
2.3	Commonly used activation functions in ANN	33
3.1	Schematic diagram of the degrees of freedom (Source: [45])	40
3.2	Blade geometry visualisation	42
3.3	Component specific coordinate system (Source: Sarkar and Fitzger- ald [137])	45
3.4	The IEA 15MW reference wind turbine (Source: Definition of IEA- 15MW wind turbine [141])	48
3.6	Model Verification: Comparison of tower response	51
4.1	Control volume representation for momentum theory	57
4.2	Geometric representation of blade element theory	59
4.3	Variation of AOA with input features	65
4.4	Mean wind speed values sampled using Sobol sequence for generating the training data	66
4.5	Optimised weighing factor values for different blade nodes	70
4.6	Comparison of physics-based loss for DDNN and PINN surrogate .	72

List of Figures

4.7	Comparison of surrogate model estimate of AOA with BEM theory	73
4.8	Comparison of surrogate model estimate of Aerodynamic forces with BEM theory	74
5.1	LSTM Architecture (Source:[174])	83
5.2	Graphical representation of computing blade deformations from turbulent inflow using surrogate modelling approach	93
5.3	Feature score prediction using MRMR algorithm	94
5.5	Correlation matrix of features identified using MRMR algorithm . .	97
5.6	Performance of LSTM Model trained using features identified through MRMR algorithm	98
5.7	Effect of increasing input feature space on model performance . . .	99
5.8	Best performance achieved from LSTM Model trained using features identified through MRMR algorithm across different stochastic realisations	99
5.9	Best performance achieved from LSTM Model trained using features identified through MRMR algorithm across different mean wind speeds	100
5.10	Error analysis of the LSTM Model trained using features identified through MRMR algorithm across different stochastic realisations . .	100
5.11	Error analysis of the LSTM Model trained using features identified through MRMR algorithm across different mean wind speeds	101
5.16	Performance of LSTM Model trained using Exhaustive Feature Search across different mean wind speeds	106
5.17	Performance of LSTM Model trained using Exhaustive Feature Search across different stochastic realisations	106
5.18	Error analysis of LSTM Model trained using Exhaustive Feature Search across different mean wind speeds	107
5.19	Error of LSTM Model trained using Exhaustive Feature Search across different stochastic realisations	107
5.20	Sample loss curve demonstrating trends in training and validation loss	109
5.21	LSTM model predictions for noisy input data	110
6.1	Visual representation of TurbSim grid	119
6.2	Modelling wind speed data as image	126
6.3	Principal components of spatio-temporal wind field	134
6.4	5×5 submatrix of DCT representation of wind field	135
6.7	Multi-stage modelling approach for predicting blade response	138

6.10	LSTM Model architecture for predicting rotor speed	142
6.11	Normalised RMSE in rotor speed prediction using different principal components	142
6.12	Comparison of rotor speed predictions of LSTM model using PCA feature at constant TI and wind shear	143
6.13	Comparison of rotor speed predictions of LSTM model using PCA feature at random TI and wind shear	144
6.14	Visual representation of DCT components considered in RFA method	145
6.15	Normalised RMSE in blade pitch prediction using different DCT components	146
6.16	Comparison of blade pitch angle predictions of LSTM model using DCT features at constant TI and wind shear	147
6.17	Comparison of blade pitch predictions of LSTM model using DCT feature at random TI and wind shear	147
6.18	LSTM Model architecture for predicting blade response	149
6.20	Multi-task learning LSTM Model predictions of blade DOFs using PCA features	151
6.22	Multi-task learning LSTM Model predictions of blade DOFs using DCT features	154
6.24	Multi-task learning LSTM Model predictions of Blade DOF using multi-stage modelling	158
6.26	Multi-task learning LSTM Model predictions of Blade DOF using multi-stage modelling	161

List of Tables

1	List of symbols	xi
2	List of Acronyms	xii
3.1	Key parameters of the IEA-15MW wind turbine	49
3.2	Absolute error in prediction of mean response from MBD model and OpenFast at different mean wind speeds	50
3.3	Absolute error in prediction of extreme response from MBD model and OpenFast at different mean wind speeds	52
3.4	Normalised RMSE in prediction MBD model and OpenFast response at different mean wind speeds	52
4.1	Hyperparameters and exploration range	69
4.2	Data Fitting Errors	71
5.1	Environmental parameters used and their limits	84
5.2	Hyperparameter bounds considered for tuning	95
5.3	Effect of number of input force channels on NRMSE	99
5.4	Feature selection method and obtained performance	108
5.5	Effect of noise on LSTM model performance	109
6.1	Comparison of the effect of uncertainty level on rotor speed prediction	144
6.2	Comparison of the effect of uncertainty level on LSTM Model performance	147
6.3	Summary of hyperparameter exploration to optimise the number of hidden units in learning layers	149
6.4	Comparison of LSTM Model performance trained using PCA and DCT features	152
6.5	Comparison of PICP values for LSTM Model trained using PCA and DCT features	156

List of Tables

6.6 Comparison of Normalised RMSE for multi-stage modelling with ROSCO control output	159
--	-----

Glossary

Table 1: List of symbols

Symbol	Description
q_{Sg}	Foundation surge DOF
q_{Sw}	Foundation sway DOF
q_{Hv}	Foundation heave DOF
q_R	Foundation roll DOF
q_P	Foundation pitch DOF
q_Y	Foundation yaw DOF
q_{TFA1}	First tower fore-aft bending mode DOF
q_{TFA2}	Second tower fore-aft bending mode DOF
q_{TSS1}	First tower side-to-side bending mode DOF
q_{TSS2}	Second tower side-to-side bending mode DOF
q_{Yaw}	Nacelle yaw DOF
q_{GeAz}	Generator azimuth angle DOF
q_{DrTr}	Drive-train torsional flexibility DOF
q_{BiF1}	First flapwise bending mode for i^{th} blade DOF
q_{BiF2}	Second flapwise bending mode for i^{th} blade DOF
q_{BiE1}	First edgewise bending mode for i^{th} blade DOF
ϕ_{BiF1}	First flapwise mode shape for i^{th} blade
ϕ_{BiF2}	Second flapwise mode shape for i^{th} blade
ϕ_{BiE1}	First edgewise mode shape for i^{th} blade
$\delta_{B1_{op}}$	First blade out of plane deflection

Table 2: List of Acronyms

Acronym	Description
O&M	Operation and Maintenance
FE	Finite Element
MBD	Multi-body Dynamic
DOF	Degrees of Freedom
CFD	Computational Fluid Dynamics
FORM	First Order Reliability Method
SORM	Second Order Reliability Method
LIDAR	Light Detection and Ranging
SGD	Stochastic Gradient Descent
Adam	Adaptive Moment Estimation
GPR	Gaussian Process regression
GP	Gaussian Process
PCE	Polynomial Chaos Expansion
UQ	Uncertainty Quantification
ANN	Artificial Neural Network
NN	Neural Network
DDNN	Data-driven Neural Network
PINN	Physics Informed Neural Network
AOA	Angle of Attack
LSTM	Long Short-Term Memory
DCT	Discrete Cosine Transform
PCA	Principal Component Analysis
MAE	Mean Absolute Error
MSE	Mean Square Error
RBDO	Reliability based Design Optimisation
NRMSE	Normalised Root Mean Square Error

Chapter 1

Introduction

The global energy demand is ever increasing. Electricity, once a luxury, has become an essential element of modern life, powering homes, industries, and technological advancements. Access to electricity is increasingly recognised as a fundamental human right and living without it seems unimaginable. However, our dependence on fossil fuels for electricity generation comes at a significant environmental cost. Fossil fuels, mainly coal and natural gas, have dominated electricity generation for decades. In traditional power plants, fossil fuels are burnt to generate energy, which is converted mechanically to produce electricity. Although these power plants have offered reliable and relatively inexpensive energy for many years, the combustion of fossil fuels that power these facilities releases a large amount of pollutants into the atmosphere. The consumption of fossil fuels has been identified as the single largest source of air pollution [1]. According to the Lancet Commission on Pollution and Health, air pollution, in all forms, is responsible for more than 6.5 million deaths worldwide every year, which has increased significantly over the past two decades [2]. Air pollution is a significant factor that contributes to the development of chronic diseases such as asthma, bronchitis, and lung cancer [3, 4, 5]. Air pollution also causes significant ecological damage related to acid rains, soil degradation, bioaccumulation, and biomagnification [6, 7, 8]. In addition, fossil fuel-based power plants consume high volumes of water for operations, including cleaning, cooling, and other processes. According to International Energy Agencies (IEA), in 2021 alone, fossil fuel-based power plants consumed 189.6 billion cubic metres of fresh water, equated to approximately 5% of global freshwater consumption [9]. The presence of these




Figure 1.1: Global average temperature anomaly(Source:[Berkeley Earth](#))

power plants in water-scarce areas further aggravates this problem. The deployment of more water-efficient sources for electricity generation could significantly reduce the water consumption of this sector. In addition to pollution, the most concerning consequence of fossil fuel use is its impact on climate change. Burning fossil fuels releases greenhouse gases, mainly carbon dioxide, which trap the heat reflected from the atmosphere of the earth and cause an increase in global temperatures. The National Oceanic and Atmospheric Administration (NOAA) published a report [10] that points out that the temperature of the Earth has been increasing at 0.08°C per decade since 1880. However, in the past 40 years, the rate of increase has almost doubled to 0.18°C . The variations in the ambient surface temperature of the Earth, presented in Fig. 1.1, clearly highlight this trend. The rise in the temperature of the Earth's surface has induced extreme weather events such as droughts, wildfires, floods, and rising sea levels, threatening human populations and ecosystems worldwide. This phenomenon in a broader context is known as *climate change*. The impact of climate change has become more evident in recent decades. Disastrous floods in Venice, record-breaking wildfires in California, and the confirmation of 2023 as the hottest year on record are stark reminders of the situation's urgency. If left unchecked, climate change could lead to irreversible tipping points, such as the collapse of polar ice sheets or the disruption of ocean currents, drastically altering our planet's climate and ecosystems [11]. A 2023 European Environment Agency




Figure 1.2: Impact of Paris accord on fossil fuels (Source: [Climate Nexus](#))

(EEA) report estimates that economic losses from weather and climate-related extremes in the EU reached more than half a trillion euros between 1980 and 2021 [12]. This includes the cost of infrastructure damage from extreme weather events, agricultural losses due to changing weather patterns, and the mass displacement of people living in areas affected by rising sea levels. The scientific community and governments have recognised climate change as an existential threat, prompting a global effort to reduce dependence on fossil fuels. The United Nations Framework Convention on Climate Change (UNFCCC) is an international treaty among member countries to combat climate change. UNFCCC organises an annual convention, and the Conference of Parties (COP), which is the decision-making body of the convention, reviews the implementation and sets the targets for countering climate change. In COP21 held in Paris, a total of 196 nations agreed to a legally binding international treaty to hold the increase in global average temperature well below 2° C above preindustrial levels and pursue efforts to limit the increase in temperature to 1.5° C above preindustrial levels. This treaty is also known as the Paris Accord. This agreement has further strengthened the growth of the renewable energy sector. The impact of the Paris Accord on fossil fuels and renewable energy has been analysed by Climate Nexus, a non-profit strategic communication group dedicated to highlighting the wide-ranging impacts of climate change and clean energy solutions, and is presented in Fig. 1.2 and Fig. 1.3. These figures show a wide-scale move towards the renewable sector. A 2022 report from the Environment Protection Agency (EPA) shows the contribution of various sectors to global greenhouse




Figure 1.3: Impact of Paris accord on renewable energy (Source: [Climate Nexus](#))

gas emissions. These statistics are presented in Fig. 1.4, which shows the major emissions arising from the agriculture, transportation, energy, and manufacturing industries. Each of these sectors is going through pivotal changes and is moving toward more sustainable alternatives. In recent years, the transportation sector has largely adopted electric vehicles in an effort to reduce emission levels, while industries are exploring options such as green hydrogen and renewable energy powered microgrids. Electricity generation has also seen huge acceptance and significant investment in renewable energy generation using solar energy, wind energy, and hydropower.

Beyond environmental concerns, the reliance on fossil fuels extends to various economic and social aspects. Fossil fuels are a non-renewable source of energy, and their continued extraction raises concerns about resource depletion in the future. A 2002 report from the United States Geological Survey (USGS) predicts a potential shortage of fossil fuels within the 2050-2060 time frame [13]. The finite amount of resources combined with the changing geopolitical landscapes can lead to significant uncertainty in the energy security of a nation. In recent years, oil prices have seen a lot of volatility, partly impacted by the COVID-19 pandemic and major geopolitical events. A graph of crude oil prices in the last five years (shown in Fig. 1.5) highlights the volatility of the sector. The geopolitical landscape often dictates energy security, with countries dependent on fossil fuel imports from potentially volatile regions.

A transition to renewable energy sources offers greater energy security and




Figure 1.4: Global greenhouse gas emission by sector (Source: [EPA](#))

independence. Countries that have increased the share of renewable energy in their total energy consumption have seen significant positive impacts. So far, the move towards renewable energy has shown promising positive impacts on sustainability goals. Ireland, for example, has saved around 300 million euros on the imports of fossil fuels in the last ten years. In the same period, the nation has observed a 10% reduction in their greenhouse gas emissions. Similar encouraging examples are found in other nations such as Denmark, Norway, and the United Kingdom (UK). The growth of renewable energy has also given economic boost to these countries by creating a demand for high-skilled job opportunities and opening the doors to new businesses. Today, the wind energy sector directly and indirectly employs 3 million people globally [14]. In the UK alone, the wind sector contributes to more than 300,000 jobs [14].

Although the Paris accord has led to considerable investment in renewable energy sectors, the energy sector is still far from achieving the goals set in the Paris accord. MSCI, a leading provider of critical decision support tools to the global investment community through research and data analysis, has developed a report on current progress in achieving climate goals and the consequence of not achieving these targets (shown in Fig. 1.6). This report highlights the urgency to further accelerate the move towards renewable energy. The recently concluded Conference



Figure 1.5: Global crude oil prices over last ten years (Source: [Macrotrends](#))

of Parties (COP28) held in the UAE undertook the first global stocktake of the Paris agreement. This COP also calls for the tripling of renewable power and the double of energy efficiency by 2030 as crucial steps to achieving the goals set out in the Paris agreement. In general, this highlights the importance of increasing the contribution of renewable energy sources to overall energy consumption. Obstacles that hinder the advancement of wind energy can be systematically categorised into several distinct domains, including social, economic, and technical challenges [15, 16, 17]. Each of these areas presents various aspects of the limitations that must be addressed to facilitate the broader development of wind energy as a key component of the global energy mix. Some of the key technical challenges that hinder wind turbine advancement will be reviewed in the next section.

1.1 Challenges and Opportunities in Wind Turbine Design

The wind energy community has responded to the challenges and the sector has seen significant growth over the past two decades in terms of their contribution to

1.1. Challenges and Opportunities in Wind Turbine Design




Figure 1.6: Paris agreement - Current Status and Consequences (Source: [MSCI](https://www.msci.com))




Figure 1.7: Increase in scale of wind turbines (Source: [KPMG](#))

overall energy generation. Wind energy now contributes 5% to global electricity consumption and forms 23% of the overall production of renewable energy. In line with the growing energy production, the size of wind turbines has seen tremendous growth. In the period of 15 years (2000-2015), the production capacity of wind farms increased from a few tens of megawatts to hundreds of megawatts [18]. The results are feats of engineering that are almost hard to imagine, with towers taller than skyscrapers. A pictorial representation of the growth on the scale of wind turbines is shown in Fig. 1.7. Reducing the cost of electricity generation has been a major driving force in the recent developments of wind turbine technology. As wind turbines are becoming larger and more complex, a new class of design challenges are emerging. Increased size of wind turbine components translates to greater loads and deformations. However, the push to reduce levelised cost of energy (LCOE) requires all components of the wind turbine system to be light and cost-effective at the same time satisfying the performance and reliability requirements. Wind turbines are unique systems which are at the intersection of structures and machinery. The major components of a wind turbine such as tower, blades and foundation are designed using the concepts of structural engineering and the drive-train system, gearbox and generators require the concepts of mechanical and electrical design. The major conflict occurs during the design, where the wind turbine system as a whole is designed assuming standard conditions and manufactured in factories, but they are installed at various locations where their performance depends on the specifics of the installation location.

The wind turbine structures are subjected to dynamic loading governed by the environmental parameters such as wind speeds, turbulence intensity and wave

heights, to name a few. A combination of these parameters represent a unique load spectrum acting on a wind turbine. The external loading governed by these environmental parameters is inherently stochastic in nature, manifesting itself in the form of aleatory uncertainty which cannot be reduced but only identified and quantified. Along with the other challenges such as defects introduced during manufacturing process [19] and accurate modelling of highly flexible elements [20], the aleatoric uncertainty coming from external loading is one of the major challenges for optimum design and operation of wind turbines as the uncertainty levels are unique to the installation location [21]. This uncertainty in external loading will lead to a variation in load distribution statistics at each location, leading to a change in overlap between demand and load distribution giving rise to the uncertainty in the achieved safety factors at different installation location. To ensure consistent reliability, a site-specific performance evaluation of wind turbines is necessary, which involves evaluating the system performance for the loading envelope particular to the installation location. Due to the high number of uncertain variables governing the dynamic loads acting on a wind turbine, a large number of simulations are required to ensure a statistically significant analysis of the site-specific load demand. This problem is further compounded by the high computational cost of running the numerical models of wind turbine. These factors combined makes the detailed site specific performance analysis a computationally intensive task. As such, site-specific performance analysis of wind turbines is an open problem which is widely researched by the wind turbine community. Recent publications from the leading researchers in the field of wind turbine have specifically highlighted the importance of revising the design process and proposing the models suited to the requirements and limitations of each site [16, 21, 22]. In order to comprehensively include the impact of uncertainty at each site in the design and analysis of wind turbines, a computationally efficient approach for predicting the response at a particular site locations is required. Surrogate models offer a potential solution to this problem.

1.2 The Role of Surrogate Models

Surrogate modelling represents a pivotal strategy in dealing with computationally demanding models, providing a streamlined representation of the intricate relationships between inputs and outputs. At its core, surrogate modelling can be perceived as a regression task, where a functional approximation of a physics-based model is

obtained for a chosen input-output parameters. Having a low-order representation of computationally intense models allows us to run large sets of simulations at a fraction of the computational cost. Although developing accurate low-order representations of highly complex non-linear systems presents challenges, significant research has been done in this field to minimise the uncertainties and improve the accuracy of surrogate models. In this context, the emergence of machine learning (ML) models as a global functional approximator capable of modelling highly non-linear input-output relationships has further enhanced the surrogate modelling capabilities. The use of ML models presents a possible efficient approach for developing surrogate models capable of resembling the physics-based wind turbine models. The ML models are capable of learning point-to-point mappings as well as temporal relationships using the data from high fidelity models. Additionally, technological advancement in computations, such as the use of Graphical Processing Units (GPUs) for training and predictions, further enhance the capabilities of ML models in efficient processing of large volumes of data. These characteristics naturally make ML models very promising candidates for surrogating physics-based models.

This thesis aims to investigate the application of machine learning-based surrogate models to reduce the computational burden associated with the dynamic analysis of wind turbines. This research will develop machine learning models to bypass the computationally demanding tasks associated with the dynamic analysis of wind turbines, with the ultimate goal of bypassing the entire numerical model. Such models, which are capable of quickly estimating the dynamic response of wind turbines to turbulent wind inflow, offer multiple advantages, for example, quick site-specific performance assessments and fatigue life predictions. The thesis will evaluate and compare the effectiveness of physics-informed neural networks against data-driven neural networks specifically for wind turbine applications. The findings of this comparison will inform future research directions. In addition, sequence-based modelling will be used to predict dynamic responses, and a framework will be proposed to outline the algorithm to develop such models. The application of some of these models as a virtual sensor to digitally measure operating parameters which are difficult to accurately measure on site using physical sensors will also be explored.

1.3 Structure of the Thesis

Building on the information presented in this chapter, the remainder of the thesis will discuss these topics in further detail. Chapter 2 will provide a comprehensive review of various surrogate modelling techniques used in the domain of wind turbines. The methodology for the development of surrogate models and their applications in wind turbine design and analysis will be explored. Chapter 3 will detail the development and validation of a numerical model of a wind turbine which is used to generate the data required to train the surrogate models. Chapter 4, 5 and 6 will present the methodology used to develop various surrogate models and discuss their suitability in wind turbine design and operation. Finally, Chapter 7 concludes by summarising the findings and presenting future research directions.

Chapter 2

Literature review

2.1 Introduction

The inexhaustible potential of wind, minimal lifecycle emissions, and scalability solidify the position of wind energy as a critical component in the global transition toward clean and sustainable energy. However, with advances in wind turbine technology driven by the pursuit of greater energy production and reduced levelised costs, significant challenges emerge in the design, operation, and maintenance (O&M) of these systems. Design, analysis and O&M of wind turbine systems involves a multi-disciplinary approach, where an understanding of various disciplines is required. A representative sketch of the wind farm (adapted from Ref. [16]) highlighting various disciplines involved and their specific roles are shown in Fig. 2.1.

Building on the issues discussed in Chapter 1 regarding general trends and current challenges in wind engineering, this chapter identifies the major challenges that hinder the advancement of wind turbines and explores potential solution strategies. We draw on recent publications by leading researchers [16, 21, 22, 23] who have outlined these challenges and research needs for the further development of the wind energy sector. These publications highlight the challenges in the following categories:

1. The Atmosphere: A Domain of Stochastic Complexity

The atmosphere serves as the operational domain for wind turbines, and var-




Figure 2.1: Various disciplines involved in achieving efficient analysis, design and operation of wind turbines (Source:Veers et al. [16])

ious effects and parameters within it significantly impact their performance. These factors include turbulence, wake effects, oceanic flows (for offshore installations), and waves. The stochastic nature of these parameters presents a major challenge. An accurate mathematical representation of these phenomena is challenging, particularly with the increasing scale of wind turbines. As the dimensions of the turbine increase, some of these atmospheric parameters are likely to have a more pronounced impact, leading to greater uncertainty in the external loads acting on the structure. Developing robust mathematical or statistical models of the atmosphere is crucial to gaining a representative understanding of the external loads that a turbine can experience during operation. This understanding will ultimately help establish a modified design process that can more accurately account for these atmospheric effects.

2. The Turbine: Scaling Up Complexity

The Turbine category encompasses the wind turbine structure itself and its various components. A significant increase in the scale of these components has introduced additional complexities in their analysis, design, and installation. The increased scale also translates into increased loads and costs associ-

ated with O&M activities. Furthermore, the relentless pursuit of ever-larger turbines is limiting the opportunities to learn from the performance of the currently installed turbines. Since current large-scale turbines are still untested in real-world conditions, this raises concerns regarding the reliability of these rapidly evolving designs.

3. The Plant and Grid: Understanding Farm-Level Effects

Understanding the impact of wind farm-level flows on individual turbines and overall power output is crucial for efficient grid management and the larger energy system. Developing models for these farm-level flows, along with controllers designed to optimise total power output as opposed to individual turbine-level control, holds the potential for significantly enhanced power production. Optimising the wind farm layout while considering all these factors can further contribute to achieving higher overall efficiency and reduced loading at the grid level.

4. Environmental Co-Design: Sustainability in Focus

Environmental co-design is a crucial element for achieving sustainable development. This concept emphasises avoiding, minimising and compensating for the direct and indirect environmental impacts of wind turbine development. Furthermore, it advocates for incorporating environmental costs and benefits into the decision-making process. Integrating these principles ensures minimal negative environmental impact while fostering sustainable development.

5. Social Science: Addressing Societal Concerns

Social science considerations are equally important. By acknowledging the impact of wind turbine development on society and fostering open dialogue, a wider acceptance of wind turbines can be achieved within diverse communities.

The five challenges outlined here are all interconnected and critical to ensuring the continued growth of wind turbines as reliable and sustainable sources of energy generation. Advancements in any one of these areas are bound to influence all the other allied domains. Although all five areas have merit, this thesis will focus primarily on the second challenge: The Turbine. The main challenges in wind turbine modelling, design, and analysis, as well as their potential solutions, are presented next.

2.2 Wind Turbine modelling

Wind turbines play a crucial role in the transition towards renewable energy systems. However, ensuring their continued growth requires achieving a balance between two key parameters: safety and affordability. Wind turbines are intricate engineering systems that require the seamless integration of multiple components and disciplines to develop accurate numerical models. The interaction between aerodynamics, structural dynamics, and control systems with a highly stochastic loading environment requires sophisticated modelling approaches for reliable design and analysis. Some of the main factors contributing to the complexity of wind turbine modelling are:

1. **Diverse Materials and Intricate Geometries:** Wind turbines use a variety of materials, from fibre-reinforced composites in blades to steel and concrete in towers and foundations. Each material possesses unique properties and behaviours that must be accurately represented in numerical models. The complex geometries of components, particularly blades with their aerodynamically optimised shapes, further compound the modelling challenges.
2. **Highly Stochastic Loading:** Wind turbines operate in an inherently unpredictable environment. Fluctuating wind speeds, turbulence, and changes in wind direction introduce highly variable loads. Models must accurately capture these stochastic forces and their interaction with the flexible wind turbine structure to predict realistic dynamic responses.
3. **Multi-disciplinary Interactions:** Wind turbine design requires the integration of multiple disciplines of engineering. Aerodynamic forces dictate the loads experienced by the structure, while structural dynamics govern the response of the system. Control systems further modify the system's behaviour in real time. Capturing these complex interactions is essential for accurate modelling.

To address these multifaceted challenges, a variety of specialised modelling techniques have been developed, each tailored to provide specific insights. For the purpose of this thesis, wind turbine models are categorised into two broad classes based on the modelling purpose, namely, structural model and aeroelastic models. This terminology is used in a specific context for this thesis. Apart from this classi-

fication, various models are developed by researchers for addressing specific research objectives [24, 25, 26, 27, 28, 29, 30].

2.2.1 Structural Models

Structural models of a wind turbine are designed to provide an accurate analysis of the response of the structure subjected to external loading, analysing deformations, strains, and stresses. These models are derived using fundamental governing principles. Finite element (FE) models represent the state-of-the-art for developing detailed structural models of a wind turbine. In FE modelling, the complex geometry of a structure is discretised into smaller elements (e.g., beam, shell, or solid elements) based on the nature of the components and principles governing their deformation characteristics. The behaviour of each element is described using material properties and shape functions [31]. The governing equation of the system is derived by assembling the finite elements into global system matrices. By solving these equations across the system, the FE model predicts structural deformation, stress-strain response, and the dynamic behaviour of the structure [32, 33]. This information is crucial for analysing the strength of the components, optimising the designs and predicting the dynamic response under the highly stochastic load experienced by the components of wind turbines [34, 35, 36]. However, FE modelling of wind turbines poses unique challenges. Accurate modelling of the complex behaviour of fibre-reinforced composites requires specialised material models within the FE framework. Furthermore, the sheer size of the wind turbine components can lead to computationally demanding FE models, requiring careful consideration of mesh refinement and solution strategies. Although FE modelling provides a powerful tool for structural analysis, its focus on structural behaviour means that other aspects of wind turbine operation, such as detailed aerodynamics and control system interactions, are often simplified or excluded.

2.2.2 Aeroelastic Models

Unlike structural models that focus solely on the structural response, aeroelastic models provide a more holistic representation of the whole system. These models tend to employ simplified structural representations compared to FE models while integrating aerodynamics, structural dynamics, control theory, and power predic-

tion within a unified framework. Tools such as OpenFast [37], Bladed [38], and HAWC2 [39] are some of the popular tools used to model the wind turbine and its components. This category of models is referred to as multi-body dynamic (MBD) models throughout this thesis for simplicity. These tools have wide-ranging applications in wind turbine design and analysis. These models can accurately predict the time-varying aerodynamic and inertial loads acting on the entire system across a range of operational scenarios. Subsequently, these loads are used for detailed structural analysis. In addition, these tools are also used to develop and refine control strategies to optimise turbine performance, regulate loads, and prevent damaging operating states [40, 41]. Finally, these models play a crucial role in estimating the energy production potential, which is critical for the analysis of economic feasibility and the integration of the grid [42, 43]. OpenFast [37], an open source platform based on MBD principles, is a particularly valuable and widely used tool for wind turbine research.

The two categories of models serve unique roles. FE models are crucial for in-depth structural design and analysis, while MBD models focus on evaluating how operational factors affect loads, structural response, and power production. Although the computational cost varies with the fidelity level selected, FE models generally require significantly more computational time compared to MBD models. Despite requiring less computing power than FE models, in absolute terms, MBD models may still require considerable simulation time, particularly under highly turbulent external loads. The high computational expense is more pronounced in iterative analyses, posing a challenge for processes such as reliability analysis, design optimisation, and fatigue analysis, which necessitate numerous simulations. However, these computationally intensive analyses are important in ensuring the safety and cost-effectiveness of wind turbine design. In the following section, the importance of reliability and fatigue analysis in wind turbine design, as well as common modelling strategies in these applications and their related computational costs are discussed.

2.3 Major challenges in Wind Turbine Design and Operations

Wind turbines have advanced greatly in size and technology in the last few decades. Key design, analysis and performance challenges arise from their larger scale and the highly variable nature of external forces. Recent advances have markedly reduced the cost of wind energy generation [44]. This cost decrease is primarily due to efficient controls, cutting-edge technologies, and a deeper understanding of turbine operations [45, 40, 46, 41, 47, 48, 49]. However, further cost reductions are possible by optimising turbine designs and ensuring consistent operational reliability. These challenges are explored in more detail in the following section.

2.3.1 Reliability Analysis and Design Optimisation

The optimal design of wind turbines must ensure maximum performance at the lowest cost while maintaining structural integrity and reliability. Reliability-based design optimisation (RBDO) is a promising method to strike this balance. RBDO is an optimisation task that focusses on enhancing the system's reliability within a set of constraints, including safety, serviceability, and cost. It has proven to be effective in balancing these factors in various structural designs [50, 51, 52, 53, 54, 55]. However, implementing RBDO for wind turbine design is particularly challenging due to the high dimensionality of the system, with numerous interacting components and control systems. In addition, challenges arise from the complex modelling of factors that affect the loading and response of wind turbines, introducing uncertainties that complicate the optimisation. Addressing these challenges and developing ways to quantify structural reliability and potential optimisation have been an area of active research over recent decades.

2.3.1.1 Early efforts in Wind Turbine Reliability and Fatigue Analysis

The early years of multimegawatt wind turbine development (1990-2000) witnessed a growing emphasis on reliability analysis and optimised design. To address this, the researchers proposed various reliability assessment frameworks. Veers [56] was one of the earliest researchers who worked towards incorporating reliability principles in the fatigue analysis of wind turbines. Veers [56] presented a stepwise process for

obtaining fatigue reliability estimates for a vertical axis wind turbine blade joint. This study modelled fatigue damage as a function of seven random variables. The fatigue failure state function, defined as the difference between time to failure and target lifetime, is modelled as a function of these random variables, where the time to failure is calculated as the number of cycles to failure divided by the frequency of cycles. Both the first-order reliability method (FORM) and second-order reliability method (SORM) were employed, and their estimates were compared. This study simplified fatigue calculation by modelling the stress coming from the external wind as a random variable with a defined probability distribution. Building upon Veers' work, further research explored probabilistic methods for fatigue analysis of wind turbine components. Lange [57] in his thesis investigated methods for the probabilistic design of fatigue-sensitive components. This research group also delved into various aspects of probabilistic fatigue analysis [58, 59, 60, 61, 62, 63]. In a parallel study, Seebregts et al. [64] pointed out the limits of deterministic design as wind turbines grow in size. Inspired by aerospace, nuclear, and offshore industries, they suggested using a Probabilistic Safety Assessment (PSA) method, based on the mean value approach, for wind turbine engineering to better integrate reliability analysis. This method, along with deterministic rules, offers qualitative insight into uncertainties in failure-damage calculations. During this time, researchers worked on improving the reliability frameworks for wind turbines [65, 66, 67, 68, 69, 70]. Early studies mostly used numerical methods like FORM and SORM to estimate failure probabilities, but these methods have limitations due to assumptions such as Gaussian distributions and low-dimensional failure surface approximations. These may not capture the behaviour of non-linear systems with correlated, non-Gaussian uncertainties common in wind turbine analysis. With progress in modelling and computing, recent research is focussing on simulation-based reliability analyses, such as Monte Carlo Simulations (MCS) and its advanced versions.

2.3.1.2 Recent trend in reliability analysis and design optimisation

As reducing the cost of wind turbines and associated levelised cost of energy is a key focus of recent research and development, ongoing efforts have prioritised developing cost-effective, optimised designs that minimise material use while ensuring structural integrity under operational loads. Consequently, design optimisation has become a major focus for many research initiatives. Reliability analysis plays a vital role in ensuring structural integrity under highly stochastic external loading condi-

tions. Some studies employ numerical models for this purpose [71, 72]. For instance, Muskulus and Schafhirt [71] optimised monopile and jacket support structures by assuming the aerodynamic loads to be independent of the pile geometry and using linear elasticity theory and linear equation of motion for reliability estimation and Sichani, Nielsen, and Naess [72] explored the effects of non-linear controller models on failure probability estimation using enhanced Monte Carlo methods [73]. However, these numerical models are simplified to reduce the computational efforts and may not give an accurate estimation of structural response. Generally, most studies rely on FE models for analysing the structural responses, combined with either multi-body dynamics codes like OpenFast, Bladed, and HAWC2 or simplified approximations for calculating loads.

Dong et al. [74, 75] analysed fatigue reliability in welded multi-planar tubular joints of a jacket support structure, a complex task due to wind and wave interactions. To simplify, aerodynamic loads were calculated in HAWC2 using a monopile model, with wind and wave responses analysed independently. They accounted for uncertainties in environmental predictions, load computations, dynamic responses, and fatigue damage using a multiplicative model with a new random variable [76], computing failure probabilities with FORM. The study notes limitations from assumptions like constant turbulence but highlights the uncertainties and challenges in wind turbine reliability analysis. Hsu et al. [77] looked at turbine tower reliability using simplified load calculations based on wind speed and evaluated stress with an FE model for fatigue analysis. Using a Weibull model for wind speed data, they combined variability with fatigue and reliability theory to estimate tower design mean time to failure. In another study, Bacharoudis and Philippidis [78] assessed wind turbine blade strength under ultimate loads, focusing on fiber failure and ply-level buckling. They used 3D shell FE models to assess blade geometry, considering material and load uncertainties for probabilistic reliability analysis. To reduce FE model costs, Tsai-Hill criteria for composites and a Response Surface Method (RSM) to predict strain replaced expensive simulations in a Monte Carlo framework to check blade reliability against IEC 61400-1 [79]. Mardfekri and Gardoni [80] proposed a framework to assess offshore wind turbines under combined wind and seismic effects, increasing complexity. A nonlinear FE model simulated seismic events, using probabilistic models for wind and seismic uncertainties, requiring substantial computational power. Fragility curves showing damage probability under different hazard levels were supported by the Texas A&M Supercomputing Facility.

For the reliability analysis of wind turbine components, as earlier mentioned, a large number of simulations are often required, making the process computationally intensive. RBDO adds further computational demands due to its iterative nature while integrating reliability analysis into the design process to achieve optimal designs that meet both performance and reliability standards. Most studies related to reliability analysis or design optimisation of wind turbines have encountered a fundamental struggle, which is trying to find a balance between model fidelity and computational feasibility. While simplified structural models can be used, they may not capture the true behaviour sufficiently for reliable designs. The pursuit of accurate and reliable designs increasingly necessitates computationally intensive calculations involving numerous iterations and simulations. Studies like Hu, Choi, and Cho [81] demonstrate the significant computational time required for RBDO despite simplifying certain aspects of the computation, highlighting the need for alternative approaches. The major challenge is developing a computationally efficient model capable of closely approximating the behaviour of a high-fidelity model for a chosen input-output parameter combination. To address this challenge, the use of surrogate models offers significant promise.

2.3.2 Fatigue Analysis of Wind Turbines

Wind turbines are growing larger, posing unique structural challenges due to dynamic forces from blade rotations and unpredictable wind loads. The cantilever design of the blades results in significant bending moments at the root due to gravity. Longer and heavier blades in larger turbines increase these moments. Rotation causes the lever arm to change sinusoidally relative to the hub, creating repetitive loading that weakens materials over time, known as fatigue. Fatigue damage is more severe in large turbines, restricting their lifespan to about 20-25 years. In line with the IEC-61400-1 guidelines [79], fatigue evaluation of an onshore wind turbine focuses on operational wind speeds from cut-in to cut-out. Within this range, wind speeds are divided into small intervals (typically 2 m/s or less), assuming constant environmental conditions per interval. A minimum of six random wind field simulations per interval is advised to counteract wind variability. Fatigue damage quantification requires analysing fluctuating load series broken down into hysteresis cycles. The rain flow counting method [82] is commonly applied, characterising each cycle by mean stress and amplitude. Total fatigue damage evaluation relies on linear damage accumulation per Miner's Rule, summed over all cycles D using

Eq. 2.1.

$$D = \sum_i \frac{n_i}{N_i(L_i^{RF})} \quad (2.1)$$

where N_i denotes the number of cycles to failure, n_i the cycle count, and L_i^{RF} is the cycle's load range about a fixed load-mean value. The relationship between load range and cycles to failure (S-N curve) is modelled by Eq. 2.2

$$N_i = \left(\frac{L^{ult} - L^{MF}}{\frac{L_i^{RF}}{2}} \right)^m \quad (2.2)$$

where L^{ult} is the ultimate design load of the component, L^{MF} is the fixed load-mean, and the Whöler exponent, m , is specific to the component under consideration. The above formulations assume the fatigue cycles occur over a constant, or fixed, load-mean. However, the actual load cycles will occur over a spectrum of load means. Therefore, a correction must be made to the fatigue cycles' load ranges to analyse the data as if each cycle occurred with a fixed mean load. This correction, known as Goodman correction, is given in Eq. 2.3, where L_i^R is the i^{th} cycle's range about a load mean of L_i^m .

$$L_i^{RF} = L_i^R \left(\frac{L^{ult} - |L^{MF}|}{L^{ult} - |L_i^M|} \right) \quad (2.3)$$

Though six random seeds are typically suggested by design codes, Zwick and Muskulus [83] indicates this leads to notable uncertainty in predicting lifetime damage. To properly address this stochastic inflow uncertainty, more simulations than current guidelines suggest are needed. While increasing simulations leads to increased computational costs, it's essential for accurate fatigue damage estimation. The effect of limited random seeds is further amplified when other environmental conditions in each wind speed bin are treated as random variables, as shown by Häfele et al. [84]. The authors found that even with 2048 simulations, there was over a 10% error in the OC4 jacket damage estimate, due to the limited seeds combined with other probabilistic environmental influences. For a simplified approach, without considering the probabilistic nature of environmental conditions and performing only six simulations per bin, a substantial number of time domain simulations are still necessary. Offshore wind turbine analysis is more complex owing to additional environmental variables like wave height, spectral period, and wind-wave angles that contribute to fatigue damage and must be factored into the analysis. To accurately capture these, a detailed sampling method is critical, significantly increasing simulations, making the analysis potentially computationally impractical. Furthermore, varying

site-specific environmental distributions complicate performance analysis. Assessing site suitability through numerous simulations at each location demands considerable computational resources, as noted in studies highlighting the prohibitive costs of accurately characterising fatigue damage [85, 86, 87]. To bypass the high computational cost of fatigue analysis, various researchers have used simplification such as the use of semi-empirical methods [88], parametric models [89, 62] and Taylor series approximation[90]. Alternate solution techniques for a quick estimation of fatigue damage through the use of spectral methods were also explored [91, 92]. In all these methods, surrogate modelling has emerged as a popular alternative providing high accuracy at minimum computational cost [93]. A sensitivity analysis of the extreme and fatigue loads on the wind turbine shows that even though the uncertainties in the modelling of the turbine’s aerodynamic and structural properties are present, the uncertainty of environmental parameters has the highest impact on the total fatigue damage [94]. As such, various surrogate models are proposed to model the fatigue damage as a function of environmental parameters, which have been shown to significantly reduce the computational efforts required in fatigue estimation[95, 96, 97, 98, 99, 100].

From the presented literature, it is evident that high computational cost is a major factor limiting both accurate fatigue analysis and reliability analysis with design optimisation. To this end, surrogate models have been shown to efficiently approximate the behaviour of complex systems, drastically reducing computational time compared to high-fidelity simulations. This is particularly relevant to wind turbines, where a large number of design variables coupled with the highly stochastic nature of wind and wave loads introduces significant uncertainties that must be assessed through numerous simulations for efficient design. In this context, numerous researchers have focused on developing surrogate models for wind turbine design and analysis applications. The literature on surrogate modelling of wind turbines can be categorised into two major classes: surrogate models for approximating loads and surrogate models for approximating response. These surrogate models were proven to be highly efficient for quick estimation of loads and their effects on wind turbines, streamlining the design and reliability assessment process. Some of the commonly used surrogate modelling techniques are presented in the next section to summarise the research conducted so far and as guidance to establish the state-of-the-art in this field.

2.4 Overcoming Computational Barriers in Wind Turbine Design: The Role of Surrogate Models

In the domain of model development, particularly those involving high-fidelity models, computational cost often presents a significant barrier in conducting studies requiring a large number of model evaluations. In such applications, surrogate modelling has emerged as a powerful solution that offers a way to develop computationally efficient approximations for these intricate models. These surrogate models effectively emulate the behaviour of the original models, while requiring significantly less computational time for evaluation and often maintaining high accuracy.

Specifically, in the context of wind turbine modelling, MBD models provide a powerful tool for analysing their complex behaviour. These models simulate the interaction of various modules, including external loads, structural response, and control systems. The modelling process can be subdivided into three key tasks: computing aerodynamic loads, calculating structural response, and estimating control parameters. These modules are coupled through a main code, which is integrated to compute the overall non-linear dynamic behaviour and power response of the system. Although such models offer valuable information, the computational cost associated with these calculations can be substantial. This computational demand often stems from both aerodynamic load calculations and the time-intensive simulations required for dynamic structural response, posing a particular challenge for studies aimed at design optimisation, uncertainty analysis, and non-linear model-predictive control.

A detailed understanding of the aerodynamic forces acting on the blades is crucial to accurately predict the performance of the wind turbine and the resulting response. The aerodynamic loads acting on the blades are a function of the inflow wind speed, the aerodynamic shape of the airfoil and the operational parameters of the wind turbine. Various numerical methods such as Lifting Line Theory, Vortex methods, Blade Element Momentum (BEM) theory and Computational Fluid Dynamics (CFD) can be used for computing the aerodynamic forces. The choice of a numerical method to estimate aerodynamic forces is guided by the required fidelity of the solution and the computational cost of these models. To this end, BEM has

emerged as a widely used approach in MBD models, offering a favourable balance between accuracy and computational cost compared to higher-fidelity methods such as CFD [101].

Once the aerodynamic loads are calculated, the task of predicting the structural response of the wind turbine introduces further computational demands. The structural models of the components of wind turbines are based on the fundamental equations of dynamic equilibrium. To determine how these structures deform and vibrate under time-varying loads, numerical methods are employed to solve these equations in a time marching fashion. Constraints imposed by algorithm stability, the dynamic characteristics of the structure, and the high stochasticity of external loads necessitate very fine temporal discretisation, often on the order of milliseconds. Considering simulations typically aim for a duration of around 600 seconds to balance transient effects, statistical significance, and stochastic wind inflow, the number of computations required to determine the dynamic response, even for a single model evaluation, becomes significant. This computational burden limits the exploration of design optimisation, where numerous simulations are needed to evaluate performance and safety margins. The substantial computational cost of both aerodynamic load estimation and dynamic response prediction calls for computationally efficient alternatives such as surrogate modelling.

The core concept of surrogate modelling lies in identifying a functional form capable of capturing the relationship between the input and output parameters. For this purpose, a representative dataset is first created by running the high-fidelity model for various combinations of input parameters. This process is often guided by a sampling strategy. Following data generation, various surrogate model algorithms are employed to identify a functional form that accurately captures the relationship between the inputs and their corresponding outputs. In essence, this process can be viewed as solving a high-dimensional regression problem to identify a functional form that can model the input-output relationship represented by the data. The choice of the type of surrogate model depends on the underlying process being surrogated and the nature of input and output variables. A good understanding of the underlying model is important for selecting the appropriate surrogate modelling technique.

Various researchers have used surrogate models for applications related to wind turbines, such as site-specific response estimation [102], fatigue calculation

[95], reliability analysis [103], and design optimisation [81]. The next section will discuss common surrogate modelling techniques used in wind turbine research, along with their capabilities and limitations.

2.5 Commonly used surrogate modelling techniques in wind turbine

Surrogate modelling, offering computationally efficient approximations, is at the heart of analyses requiring large sets of simulations. Although many surrogate modelling techniques have been implemented in wind turbine-related research, most studies are based on four popular techniques: Response Surface Method (RSM), Gaussian Process Regression (Kriging), Polynomial Chaos Expansion (PCE), and Artificial Neural Networks (ANN). It's important to note that other classes of surrogate modelling dealing with sequence modelling or temporal data are not considered in this discussion. Details of RSM, Kriging, PCE, and ANN are presented next.

2.5.1 Response Surface Method

Response Surface Method (RSM) is a fundamental surrogate modelling approach that approximates input-output relationships using low-order polynomial functions. For a set of input output parameters, the coefficients of assumed functional form are computed such that it minimises the error between the output values and the model prediction. The mathematical formulation of the RSM can be explained using Eq. 2.4,

$$Y = \beta_0 + \sum_{i=1}^k \beta_i X_i + \sum_{i=1}^k \sum_{j=1}^k \beta_{ij} X_i X_j + \dots + \epsilon \quad (2.4)$$

where X_i represents the input variables, Y is the model prediction, β terms are the model coefficients, and ϵ represents the error term. This formulation can be represented in a more compact form as

$$Y = X\beta + \epsilon \quad (2.5)$$

The coefficients (β) in the mathematical models are estimated from the input-output dataset using regression techniques. Methods such as ordinary least squares regression or weighted least squares regression are commonly used to fit the models to the

data and estimate the coefficients. Toft et al. [95] investigated the use of RSM for estimating site-specific fatigue and extreme loads using two different RSMs generated using Taylor approximation and Central Composite Design method of experiment design as a function of mean wind speed, turbulence intensity, wind shear, air density and vertical inflow angle. Both the RSMs give a good estimate of site-specific loads with errors less than 5% observed in fatigue load estimation; however, much larger errors were observed in estimating extreme loads due to high non-linearity experienced in extreme loads. The methodology enables an efficient approach for load assessment based on wind climate at a specific site. Similar applications of RSM for damage assessment and estimating the sensitivity of loads to wind parameters are presented in Ref. [104] and Ref. [105]. RSM has also been used to estimate reliability and optimise design. Lee and Shin [106] used a second order RSM to optimise chord length, and twist angle distribution for efficient blade aerodynamic design. The authors used RSM to model the impact of design variables on various design constraints, which is subsequently used in the optimisation loop. In a very similar application, Li et al. [107] used RSM to optimise the lift-to-drag ratio of an airfoil at the design angle of attack. RSM offers advantages such as interpretability (with coefficients providing insights into input importance), computational efficiency, and relatively less data requirement. However, the limitations of RSM in handling highly non-linear and high-dimensional problems due to the polynomial formulation of the method limits its applications. In this context, Gaussian Process Regression (Kriging) is one such technique that offers greater flexibility in representing complex relationships within wind turbine analyses.

2.5.2 Gaussian Process Regression

Gaussian Process regression (GPR) is a surrogate modelling technique, where instead of assuming a functional form, an entire distribution of functions that are consistent with the observed data is modelled. GPR models the relationship between input and output variables as a Gaussian process (GP), which is a collection of random variables, any finite number of which are assumed to follow a joint Gaussian distribution. This flexibility makes GPR well-suited for capturing complex, non-linear relationships. Given a set of input-output pairs (X, y) , where X is the input matrix and y is the output vector, the objective of GPR is to predict the output y^* for a new input x^* . Due to the GP assumption, the predictive distribution

of the output y^* given x^* is a Gaussian distribution, such that

$$y^*|x^*, X, y \sim N(\mu^*, \sigma^{*2}) \quad (2.6)$$

where μ^* is the mean prediction and σ^* is the predictive variance, given by:

$$\mu^* = k_*^T (K + \sigma^2 I)^{-1} y \quad (2.7)$$

$$\sigma^{*2} = k(x_*, x_*) - k_*^T (K + \sigma^2 I)^{-1} k_* \quad (2.8)$$

where k_* is the vector of covariance between the new input x_* and the training inputs, and K is the covariance matrix of the training inputs. The covariance matrix K is constructed using the kernel function $k(x, x')$, which computes the covariance between pairs of training input vectors x and x' . Each element K_{ij} of the covariance matrix represents the covariance between the output values at the input points x_i and x_j , as determined by the kernel function. Similarly, the covariance vector k_* represents the covariance between the new input x^* and the training inputs. It is calculated using the kernel function $k(x_*, x)$, which quantifies the similarity between the new input and each training input. The choice of kernel function and its parameters determine the smoothness, flexibility, and correlation structure of the GP, ultimately affecting the predictions and uncertainty estimates in the GPR. Radial Basis Function, Matern, and Exponential are some of the commonly used kernel functions in GPR. Its ability to model complex relationships enables computationally efficient RBDO applications. For example, Meng et al. [108] developed a kriging-assisted approach to streamline the optimisation of offshore wind turbine towers, bypassing the need for repeated, computationally intensive physics simulations or FE analyses. Kriging also finds use in fatigue assessment, where surrogates can replace the time-consuming analysis of stress cycles in the time domain, significantly accelerating the process [109, 110, 111]. Furthermore, Kriging's uncertainty handling makes it well suited for reliability analysis in wind turbine design [112, 113, 114]. Researchers have also incorporated Kriging into the optimisation of wind turbine components for improved performance and robustness [115, 116, 117].

Although GPR provides notable benefits for wind turbine modelling, computational demands with extensive datasets and dependence on kernel choice prompt consideration of other surrogate modelling methods. Polynomial Chaos Expansion (PCE), with its ability to capture intricate data patterns and the promise of quicker

evaluation, emerges as a promising alternative, especially for large-scale wind turbine simulations.

2.5.3 Polynomial Chaos Expansion

PCE is a powerful technique in the realm of surrogate modelling. PCE constructs a functional to approximate the input-output relationship of a complex system by using polynomial expansion in terms of orthogonal polynomials. At the core of PCE is the use of orthogonal polynomials as basis functions. For a set of input-output pairs (X, y) , the prediction output of PCE y^* is given as:

$$y^* = \sum_{i=0}^n \beta_i \phi_i(x) \quad (2.9)$$

where $\phi_i(x)$ represents the orthogonal polynomial basis functions, and β_i are the expansion coefficients to be determined. The expansion coefficients β_i are estimated by fitting the polynomial expansion to the response of the system at the input points sampled. This can be done using regression techniques such as least-squares regression. Polynomial families, such as Legendre, Hermite, or Chebyshev polynomials, possess advantageous mathematical properties. These polynomials are orthogonal with respect to the probability density function associated with the input variables of the system. This orthogonality helps to efficiently represent the input-output relationship. PCE has been widely used in wind turbine-related research for uncertainty quantification (UQ). PCE finds diverse applications in wind engineering, where quantifying the impact of uncertainties is crucial. Beyaoui et al. [118] have employed PCE to investigate the dynamic response of wind turbine gearboxes under random aerodynamic torque excitations, demonstrating their value in predicting how input variability affects system behaviour. PCE has also proven effective in characterising both energy production and fatigue loads under realistic turbulent inflow conditions [119], where the author emphasises the creation of surrogates for the mean and standard deviation of outputs, providing insights for the estimation and optimisation of site-specific power. Furthermore, PCE facilitates the efficient prediction of fatigue damage in wind turbine blades, which is essential for designing and analysing useful life [120, 121].

Addressing a common challenge in the field with handling a large number of uncertain variables, Thapa and Missoum [122] coupled PCE with l1-minimisation

to perform UQ for composite blades considering numerous random variables in material properties, geometry, and loading. PCE's UQ approach is also applied to analyse the peak tension and reliability of the mooring cables used for offshore renewable energy devices [123, 124]. In a recent study, Haghi and Crawford [125] have developed an intrusive PCE surrogate to bypass the blade element momentum aerodynamic model to estimate the aerodynamic thrust and torque acting on the wind turbine rotor. As is evident from the studies cited above, PCE is widely used in the UQ field because of its efficiency in computing the statistical properties of the output and because the sensitivity analysis is obtained without any additional effort. This method's ability to efficiently represent moderately non-linear relationships and propagate uncertainties offers insights into how input variability translates to variability in loads, fatigue, and other performance metrics. Careful selection of the polynomial basis and input distributions is required to ensure accuracy, as PCE can struggle with extremely non-linear problems or those with many inputs. The main two limitations in the use of PCE for wind energy are: (1) the input atmospheric parameters are usually jointly distributed with several layers of dependency, and (2) some of the outputs have discontinuities and/or are restricted to certain values. While PCE has been used in applications related to wind turbine modelling, the need to handle complex non-linearities and potentially high-dimensional problems motivates the exploration of even more adaptable techniques. Artificial Neural Networks (ANNs), with their flexible structure and data-driven learning, present a promising alternative to address these challenges.

2.5.4 Artificial Neural Networks

Artificial neural networks (ANNs) are computational models composed of interconnected building blocks called "neurons." Each neuron performs a linear transformation, computing a weighted sum of its inputs along with a bias term. A non-linear activation function f is then applied to this weighted sum, producing the neuron's output. This mathematical operation is represented in Eq. 2.10 where x is the input to the neuron and y is the output, and w and b represents the neuron weight and bias, respectively.

$$y = f(wx + b) \quad (2.10)$$

Neurons are organised into layers: an input layer, one or more hidden layers, and an output layer (see Fig. 2.2). As the name suggests, in a feed-forward neural network




Figure 2.2: Feed forward neural network architecture

architecture, the information flows unidirectionally from the input to the output without any feedback to the previous layer. As such, the output of one layer acts as the input to the next layer. For the overall ANN, the training process then becomes an optimisation problem in that the weights and biases are tuned during the training process to minimise the error term. Usually, the mean square error (MSE) or mean absolute error (MAE) metrics are used in the training process. The activation function introduces non-linearity in the formulation that allows ANNs to model non-linear relationships. The commonly used activation functions in ANN are shown in Fig. 2.3. ANNs learn complex relationships through the adjustment of weights and biases within their interconnected structure, this formulation provides ANN the flexibility to approximate a wide range of functions. This makes ANNs particularly well-suited as global function approximators within the context of surrogate modelling. ANNs have found a wide-scale application in wind energy research. A review article [126] surveying the application of ANN in wind turbine research has classified the application into four main areas: forecasting and predictions, design optimisation, fault detection and diagnosis, and optimal control. Of these applications, forecasting wind speeds has been one of the most researched domains with the application of ANN. A review paper summarising the application of ANN to wind energy and wind power forecasting [127] demonstrated a trend of interest in using ANN for this task. ANN have been also used for design of airfoils [128, 129], multiple component wing sections [130], and three dimensional optimisation of the wing tip [131].




Figure 2.3: Commonly used activation functions in ANN

In the context of this thesis, similar to the surveyed application of other surrogate modelling techniques, a major focus is given on the application of ANN as a surrogate model for applications with high computational demands, such as reliability, fatigue analysis and site-specific load estimation. In this area of research, various studies looked at optimisation such as Mortazavi, Soltani, and Motieyan [132] used ANN to design the airfoil sections for the blades of a horizontal axis wind turbine where the ANN are trained using the results of computational fluids dynamics analysis to obtain a Pareto optimal set of solutions. The airfoil design was also considered by Ribeiro, Awruch, and Gomes [133], where the use of ANN results reduced the computational time by around 50%. All these studies report high accuracy achieved by ANN at a substantial reduction in computational cost. In the context of site-specific load estimation work at the Technical University of Denmark, a series of publications [102, 134] highlighted that ANN could produce site-specific loads more accurately and faster than the conventional surrogate methods presented in subsection. Another publication from this research group [135] further investigated the effect of uncertainty in input parameters on the ANN prediction by performing uncertainty propagation by means of a global sensitivity analysis. The results of this study demonstrate that the selected ANN architecture results in small output uncertainties, highlighting its effectiveness in surrogate modelling tasks. The additional capabilities, such as explicitly mentioning the error term and allowing for training

an ANN surrogate capable of satisfying multiple objective constraints, give this an edge over conventional surrogate modelling techniques. A recent publication (Ref. [136]) extending the application of ANNs to explicitly minimise the physical laws governing the system, giving rise to the Physics Informed Neural Networks (PINNs) have garnered huge attention from the scientific community. Subsequently, PINNs have been at the forefront of scientific machine learning, where the NN can make use of the physical laws of the system rather than relying only on data. Encoding such structured constraints founded on the physical knowledge of the system into ANN results in amplifying the information content of the training data, enabling a more efficient exploration of parameter space and better generalisation ability even when only limited training data is available. These innovations make ANN, in general, and PINN, in particular, an ideal choice for surrogate modelling techniques.

2.6 Summary

In the context of the discussion presented in this Chapter, ANN and PINN are used as the preferred surrogate modelling algorithm for further research in this thesis. Further capabilities available in the larger domain of machine learning such as sequence modelling through recurrent neural networks (RNNs), Long Short-Term Memory (LSTM) model and Convolutional Neural Network (CNN) extends the capability of machine learning models to model time history response. In this thesis, ANN models are used to develop time invariant surrogates, and LSTMs are used to emulate the dynamic response of wind turbine blades. The framework for developing an MBD model of wind turbines and computational bottlenecks in their dynamic analysis are presented in the next chapter.

Chapter 3

Multi-body dynamic model of a wind turbine

3.1 Motivation

Dynamic analysis of wind turbines under a wide range of operating conditions is essential for the analysis and design of these structures. Chapter 2 of this thesis discussed different types of wind turbine models classified according to the modelling strategy and degree of fidelity. The multibody dynamic (MBD) modelling technique is particularly effective in modelling the holistic performance of a wind turbine, as it can integrate various aspects of wind turbine operations such as aerodynamics, controller routines, structural behaviour, and power output. In this study, an MBD model of an onshore wind turbine is developed. The model includes controllers to regulate wind turbine operation and utilises BEM theory to calculate external loads. These features allow for the analysis of wind turbine structural responses, load prediction, and the development of efficient control strategies.

While the numerical model offers various capabilities, it suffers from the high computational cost of performing dynamic analysis. This high computational cost is particularly evident when analysing the response for a wide range of operational scenarios, motivating the development of computationally efficient surrogate models. These surrogate models aim to capture the input-output behaviour of the MBD model with significantly reduced computational cost.

This chapter presents the formulation of an MBD model that serves as the baseline for the development of subsequent surrogate models. The key components of the model, their interactions, and the types of dynamic response simulated are presented in this chapter from a conceptual standpoint, avoiding the intricate details of the kinematic derivations. Since numerous references with detailed discussions on modelling wind turbines using this approach exist, this chapter focusses on providing a conceptual overview of the modelling approach; interested readers are recommended to refer to Ref. [137] and [138] for an in-depth understanding of the modelling technique.

The MBD model of a wind turbine represents the entire system as an assembly of interconnected components described using multiple degrees of freedom (DOFs). DOFs represent a set of independent coordinates that can completely describe the motion of the component. The choice of DOFs is crucial to capture the overall dynamics of a system. As such, the choice of DOFs is guided by the desired level of accuracy and computational efficiency. The number of DOFs required to model a component depends on its geometry and the possible motions it can undergo during the operation. The various components of a wind turbine and the choice of DOFs used to model them in this study are explored in the next section.

3.2 Structural Elements of Wind Turbines and Modelling Strategies

Wind turbines are sophisticated systems with multiple components. From a structural behaviour perspective, components are classified as rigid or flexible on the basis of how they deform under the action of external loads.

3.2.1 Modelling rigid elements

A rigid body is defined as an object in which all the particles maintain fixed distances relative to each other, even when subjected to external loads. As such, the geometry of such objects is maintained during the entire operation. This assumption simplifies the analysis of mechanical systems by neglecting deformations and focussing solely on rigid body motions. For the MBD model of a wind turbine developed in this study, the nacelle and hub are modelled as rigid bodies.

3.2.2 Modelling flexible elements

Flexible bodies are physical objects that undergo deformation, leading to changes in their shape, when subjected to external loads. These components undergo significant deformations exhibiting bending, twisting, or other forms of deformation under typical operating loads, influencing the dynamic response of the wind turbine. For the MBD model of a wind turbine developed in this study the blades, tower, and drivetrain are modelled as flexible bodies. The modal summation method is implemented for efficient modelling of flexible elements.

An accurate representation of flexible structures often involves treating them as continuous systems with an infinite number of DOFs. This approach poses computational and modelling challenges, as the response at each point must be calculated to completely define the deformed configuration. The modal summation method provides a bridge to efficiently model flexible elements by discretising the structure and decomposing its complex deformations into a summation of simpler shapes called mode shapes. Each mode shape represents a unique fundamental pattern of deformation the structure can exhibit at its corresponding natural frequency. By combining a weighted sum of these mode shapes, the overall deformation of the flexible body under various load conditions can be effectively represented. Using this approach, the number of DOFs required to model the structure can be reduced to the number of mode shapes defined using their modal coordinate. The deformation of a flexible body using modal summation method can be described using Eq. 3.1:

$$u(x, t) = \sum_{i=1}^n \phi_i(x) q_i(t) \quad (3.1)$$

where $u(x, t)$ represent the deformation of a flexible body at a spatial location x at a given time t , $\phi_i(x)$ represent the i^{th} mode shape of the structure defined as a function of spatial coordinate x and $q_i(t)$ represents the i^{th} modal coordinate, a time-varying coefficient that determines the contribution of the i^{th} mode shape to the overall deformation at a given time t . The number of mode shapes ' n ' required to capture the dominant dynamics of the member is an important modelling consideration which determines the accuracy of the resulting model. Once the number of mode shapes ' n ' are defined, the system is solved for corresponding ' n ' modal coordinate. As such, the deformation at any point on the beam could also be defined using n

other functions ($\phi_j(x)$), called shape functions, using Eq. 3.2

$$u(x, t) = \sum_{j=p}^{n+p-1} \phi_j(x) q_j(t) \quad (3.2)$$

where $q_j(t)$ is the generalised coordinate associated with the shape function $\phi_j(x)$ and the parameter p is chosen based on the deformation patterns of the element and modelling assumptions. Since the shape function is not unique to a mode shape, the mode shape can be defined as a linear sum of shape functions using the Rayleigh-Ritz method, as shown in Eq. 3.3

$$\phi_i(x) = \sum_{j=p}^{n+p-1} C_{i,j} \phi_j(x) \quad (3.3)$$

In the MBD model, the mode shapes are assumed to be polynomial expression, therefore, the shape functions are defined using power function basis (see Eq. 3.4)

$$\phi_j(x) = \left(\frac{x}{X} \right)^j \quad (3.4)$$

Since the flexible members (tower and blade) are modelled as cantilever beam elements, the assumed shape function should be able to model the corresponding deformation pattern. A deformed shape of cantilever is characterised by zero slope at fixed end, a shape function model with $p = 2$ is chosen to capture the deformed configuration. Using this mathematical formulation, structural elements of wind turbine are modelled to capture the dominant dynamics of the structure. The methodology adopted for modelling individual structural elements are discussed in the next section.

3.3 Structural Elements and Degrees of Freedom

This section details the DOFs used to represent the various structural components of the wind turbine within the MBD model. Modelling choices are guided by the need to achieve a balance between accuracy in capturing key dynamic responses and computational efficiency. The DOFs used to model a component directly influence the dynamic behavior of the system. The main structural elements and their modelling assumptions are presented next:

- **Foundation:** The wind turbine foundation is modelled using six DOFs, representing three translations (q_{Sg}, q_{Sw}, q_{Hv}) and three rotations (q_R, q_P, q_Y). For this study focussing on onshore installations, the foundation is assumed to be fixed. However, soil-structure interaction models can be incorporated to model the flexibility of the foundation. Depending on the soil-structure interaction model, the foundation can be modelled as either a rigid or a flexible body.
- **Tower:** The tower is modelled as a flexible cantilever beam, allowing it to deform under wind-induced loads. To capture its predominant bending behaviour, four modes and the corresponding DOFs are included: two DOFs representing the behaviour in the fore-aft direction (q_{TFA1}, q_{TFA2}) and two in the side-to-side direction (q_{TSS1}, q_{TSS2}). The axial deformation under the action of external loads and torsional deformations are not modelled.
- **Rotor Nacelle Assembly (RNA):** RNA has two key DOFs: the azimuth angle of the generator (q_{GeAz}), representing the rotation of the blade, and the yaw of the nacelle (q_Y), which models the orientation of the nacelle to align the rotor with the direction of the wind.
- **Drivetrain:** The drive train is modelled using a single DOF (q_{DrTr}), which accounts for the torsional flexibility of the drive train.
- **Blades:** Each blade is modelled as a flexible cantilever beam fixed at its root. Three mode shapes are used to model the blade dynamics, namely: two flap-wise modes (q_{B1F1}, q_{B1F2}) and one edgewise mode (q_{B1E1}). These modes are crucial for predicting blade deflections under varying wind loads.

As such, a total of 22DOFs are used to define the structural model of the wind turbine. A schematic diagram showing various DOFs is presented in Fig. 3.1 The dynamic response of this structure subjected to various external loads is calculated for detailed design and analysis. For an onshore wind turbine, the aerodynamic forces represent the most significant loading factor governing the component design as well as energy generation. The following section focusses on modelling of aerodynamic loads.




Figure 3.1: Schematic diagram of the degrees of freedom (Source: [45])

3.4 Aerodynamic Modelling of Wind Turbine Blades

The interaction of airflow with the wind turbine blades generates aerodynamic forces and moments resulting in blade deformations and rotations. These forces govern the dynamic behaviour, power generation, and structural loads of a wind turbine. Accurate modelling of aerodynamic loads within the MBD framework is crucial in ensuring the physical relevance of the model response. The geometry of the blade, the material characteristics, and the aerodynamic properties are important aspects that influence the response of the wind turbine. These parameters dictate the aerodynamic forces generated, influencing the blade's structural behaviour and the turbine's overall dynamic response.

The shape of a wind turbine blade is defined through a combination of geometric properties defined along its length. The chord length (analogous to the width of the blade) decreases along the span, influencing the blade's ability to interact with the wind and capture power. The blade twist, which describes the progressive change in the angle of the cross-section relative to the plane of rotation, is designed to optimise the angle of attack experienced by each blade section across varying

wind speeds. Sweep and curvature refer to the backward (or forward) bending of the blade's centerline and its out-of-plane deviations, respectively. These attributes uniquely define the blade geometry and can be optimised to ensure efficient aeroelastic behaviour. In addition to geometry, the materials used to build wind turbine blades play an important role in their efficiency and dynamic response.

Modern wind turbine blades are primarily constructed from composite materials, such as glass fiber or carbon fiber reinforced plastics. These composites are used because of their superior strength-to-weight ratio, which ensures that the blades are capable of withstand significant loads and have a comparatively lower self weight. Specialised coatings or additional materials might be included in specific regions of the blade for purposes such as erosion protection, heat shielding, or lightning protection. The MBD model primarily accounts for the effect of these materials on the blade's mass and stiffness distribution. However, high-fidelity modelling of dynamic material behaviour, such as cracking and delamination, is typically beyond the scope of MBD models.

In addition to modelling the physical geometry and materials of the blade, accurately predicting the aerodynamic forces and moments acting on the blades requires careful aerodynamic modelling. This involves discretising the blade into spanwise elements and defining unique airfoil shapes along the blade's length, tailored to optimise a balance between aerodynamic efficiency and structural strength. These aspects of blade geometry are represented using Fig. 3.2 The modelling of aerodynamic performance characteristics for these airfoils includes defining the relationship between angles of attack and aerodynamic coefficients (lift, drag, and pitching moment). These factors, combined with wind conditions and blade deformations, influence the aerodynamic loads experienced by the structure.

3.4.0.1 Computing Aerodynamic Loads

Accurate modelling of both wind inflow conditions and the calculation of the resultant wind velocity at each blade element is crucial for the determination of aerodynamic loads. Wind inflow can be modelled on the basis of idealised wind profiles or highly detailed wind field data. The specifics of the inflow wind along with the velocities induced by the blade's rotation and overall turbine motion are essential in determining the resultant wind velocity vector at each blade element.




Figure 3.2: Blade geometry visualisation

Blade Element Momentum (BEM) theory then provides a framework to compute the aerodynamic forces and moments arising from this interaction.

BEM theory is a widely used method for calculating the aerodynamic forces and moments acting on wind turbine blades. BEM theory combines two key principles: blade element analysis, where the blade is divided into discrete elements, and momentum theory, which applies momentum conservation principles to airflow. BEM theory utilises each element’s geometry (chord, twist), airfoil characteristics (lift and drag coefficients), and local wind velocity to calculate the element-specific aerodynamic forces and moment. These element-wise loads are integrated along the blade span to determine the overall loads experienced by the wind turbine. Performing aerodynamic loads computed through BEM theory form a critical input to the MBD model. These loads directly influence the structural response of the turbine, including blade deformation, tower dynamics, and power generation. The load computations using BEM is one of the major computationally intensive processes in dynamic analysis using MBD models. The next chapter provides a more in-depth discussion of these limitations and explores the development of a surrogate model for load computations as an alternative approach.

Building on the information presented so far, a robust MBD model is developed using Kane’s dynamics, to accurately capture the dynamic response of the wind turbine under various operating conditions and wind inputs. The overview of Kane’s dynamics approach for deriving equations of motion (EOM) is presented next. A complete derivation of these equations of motion using Kane’s dynamics can be found in Ref. Sarkar and Fitzgerald [137]. Interested readers may refer to Ref. Kane and Levinson [139] for the underlying structural dynamic theory.

3.5 Kane’s Dynamics Approach for Deriving EOM for the MBD Model

The equations of motion for the MBD model of a wind turbine developed in this study are based on the principles of Kane dynamics [139]. Kane’s method was chosen specifically for its efficient handling of the complex interactions between various turbine components, enabling accurate representation of the overall dynamics. This method reduces the effort required to derive EOM and offers simpler computer


implementation compared to classical approaches like Euler-Lagrange and Newton's laws of motion. As mentioned in Section 3.3, a total of 22 DOFs are considered to accurately capture the dynamics of various components of wind turbine. Multiple reference frames are defined to describe the motion of different components of the system and to define their orientation with respect to the other members. Reference frames and corresponding coordinate systems for tower (**a**), nacelle (**d**) and blades (**g**) are presented in Fig. 3.3. These reference frames are orientated with respect to each other through the component's DOFs and geometric properties of the wind turbine model. These coordinate system orientations are used in defining position vectors for the key points on the structure, which are consistently taken into account through various vector calculations performed while deriving EOM. Using this model representation, the equilibrium equations for a simple holonomic multi-body system using Kane's approach is given by:

$$F_r + F_r^* = 0 \quad (3.5)$$


where F_r stands for the generalised active forces and F_r^* represent the inertia force. For the wind turbine model the mass of the tower, yaw bearing, nacelle, hub, blades, and generator contribute to the total generalised inertia forces as follows

$$F_r^* = F_r^*|_X + F_r^*|_T + F_r^*|_N + F_r^*|_H + F_r^*|_{B1} + F_r^*|_{B2} + F_r^*|_{B3} \quad (3.6)$$


Generalised active forces represent the external forces applied directly on the wind turbine system, forces ensuring constraint relationships between various rigid bodies and internal restoring forces generated in flexible members. The external forces directly applied on the wind turbine system include aerodynamic forces acting on the blades and tower; gravitational forces acting on the tower, yaw bearing, nacelle, hub, blades; generator torque and high speed shaft brake. In addition, yaw springs and dampers contribute to forces that enforce constraint relationships between rigid bodies. The internal forces within the flexible members include elasticity and damping in the tower, blades, and drivetrain. The modal damping model is used to represent the effect of the damping forces in the blades and tower. Thus generalised active forces can be given as



(a) Coordinate system defined for tower



(b) Coordinate system defined for nacelle



(c) Coordinate system defined for blades

Figure 3.3: Component specific coordinate system (Source: Sarkar and Fitzgerald [137])

$$\begin{aligned}
 F_r = & F_r|_{AeroT} + F_r|_{AeroB1} + F_r|_{AeroB2} + F_r|_{AeroB3} + F_r|_{GravT} + F_r|_{GravN} + F_r|_{GravH} \\
 & + F_r|_{GravB1} + F_r|_{GravB2} + F_r|_{GravB3} + F_r|_{YawStiff} + F_r|_{YawDamp} + F_r|_{Gen} + F_r|_{Brake} + \\
 & F_r|_{ElasticT} + F_r|_{ElasticB1} + F_r|_{ElasticB2} + F_r|_{ElasticB3} + F_r|_{DampT} + F_r|_{DampB1} + F_r|_{DampB2} + \\
 & F_r|_{DampB3} + F_r|_{ElasticDrive} + F_r|_{DampDrive} \quad (3.7)
 \end{aligned}$$

These forces can be expressed in terms of the kinematic quantities as

$$F_r = \sum_{i=1}^n {}^E v_r^{X_i} \cdot F^{X_i} + {}^E \omega_r^{N_i} \cdot M^{N_i} \quad (3.8)$$

$$F_r^* = - \sum_{i=1}^n {}^E v_r^{X_i} (m^{N_i} {}^E a^{X_i}) - {}^E \omega_r^{N_i} \cdot {}^E \dot{H}^{N_i} \quad (3.9)$$

where F^{X_i} is a force vector acting on the centre of mass of point X_i and M^{N_i} is the moment vector acting on the rigid body N_i . ${}^E v_r^{X_i}$ and ${}^E \omega_r^{N_i}$ are the partial linear and partial angular velocity of the point X_i and rigid body N_i , respectively. ${}^E \dot{H}^{N_i}$ is the time derivative of the angular momentum of the rigid body N_i about its centre of mass X_i in the inertial frame, given by the following equation

$${}^E \dot{H}^{N_i} = \bar{I}^{N_i} \cdot {}^E \alpha^{N_i} + {}^E \omega^{N_i} \times \bar{I}^{N_i} \cdot {}^E \omega^{N_i} \quad (3.10)$$

The final equation of the system is of the form

$$M(q, t) \ddot{q} + f(q, \dot{q}, t) = 0 \quad (3.11)$$

where $M(q, t)$ is the inertia matrix and $\ddot{q}$ is the acceleration vector, whereas $f(q, \dot{q}, t)$ is the force vector consisting of the external and restoring forces acting on the structure. This system of equations is solved using numerical methods. In this study the fourth-order Runge-Kutta method is used for solving the EOM.

The EOM presented in this section are solved for predicting the dynamic responses of wind turbine. The aerodynamic forces, in particular, are the most significant load component which directly influence blade bending, tower loads, and overall power generation. The accuracy of these aerodynamic load estimates relies on the nature of the wind inflow model used. A inflow model that effectively captures the complexities of real-world wind patterns is crucial for accurately representing the dynamic loads experienced by the wind turbine. The following section discusses

the modelling of inflow wind conditions for aerodynamic analysis of the MBD model.

3.6 Inflow wind generation

The ability of the MBD model to produce data representative of actual operations is heavily based on the quality of wind inflow representation. Modelling the wind field experienced by the wind turbine involves various techniques, ranging from simplified profiles to highly detailed wind data. This study uses TurbSim [140], a widely used stochastic wind simulator, to generate wind input data. Using this approach, steady, uniform, or turbulent inflow data can be generated.

TurbSim is a stochastic, full-field turbulent wind simulator developed by the National Renewable Energy Laboratory (NREL). TurbSim generates wind fields by defining a spatial grid that spans the domain of operation of a wind turbine. This spatial grid represents a set of discrete points where wind conditions are simulated. TurbSim uses spectral models in the frequency domain to represent variations in wind speed over time. Parameters like the turbulence intensity and the coherence functions control the relationships between the wind speeds at different spatial locations and at different frequencies. An inverse Fourier transform is then performed, converting this frequency domain representation into a time series of wind speeds for each point on the spatial grid, ultimately producing a dynamic and spatially correlated wind field. The stochastic approach implemented in TurbSim ensures that each generated wind field is unique while adhering to specified statistical properties that reflect realistic wind behaviour.

The spatio-temporal wind field data generated by TurbSim serve as a critical input to the MBD model. At each time step, the wind velocities at relevant points on the blade are interpolated from TurbSim's grid. These wind velocities, in combination with blade geometry, airfoil characteristics, and rotational motion, are then used by the MBD model to compute aerodynamic loads through BEM theory. The geometric and environmental parameters in TurbSim can be adjusted to simulate wind conditions tailored to the specific wind turbine design and the environmental characteristics of the intended installation site. A detailed discussion on the specific TurbSim parameters employed in this study and the rationale behind their selection is presented in Chapter 6. The next section will introduce the specifics of the onshore variant of the IEA Wind 15 MW reference wind turbine, which serves as




Figure 3.4: The IEA 15MW reference wind turbine (Source: Definition of IEA-15MW wind turbine [141])

the baseline for the development of surrogate models presented later in this thesis.

3.7 IEA-15MW Reference Wind Turbine

This study utilises the International Energy Agency’s (IEA) 15 MW reference wind turbine as the foundation for the developed MBD model. Established by the IEA Wind research task 37 [141], this reference model provides a well-defined and extensively documented wind turbine design. Its widespread use within the research community enables consistent evaluation, benchmarking, and the development of advanced modelling techniques.

The IEA-15 MW reference wind turbine is a three-bladed, upwind, variable-speed turbine. Key specifications include a rotor diameter of 240 meters, a hub height of 150 meters, and a rated power of 15 MW. Fig. 3.4, reproduced from the technical report on the definition of the IEA 15 megawatt reference wind turbine [141], provides a visual representation of a monopile-based offshore variant of this turbine model. The details of the operating parameters of this are presented in Table 3.1.


Parameter	Value
Hub Height	150 m
Rotor Diameter	240 m
Cut-in wind speed	3.00 m/s
Rated wind speed	10.59 m/s
Cut-out wind speed	25.00 m/s
Minimum rotor speed	5.00 rpm
Maximum rotor speed	7.56 rpm

Table 3.1: Key parameters of the IEA-15MW wind turbine


The MBD model developed in this study represents the key structural components and DOFs of the IEA-15 MW wind turbine. To ensure the precision and validity of the custom MBD model, benchmarking was performed against NREL OpenFAST [37], a widely used and experimentally validated open source wind turbine modelling platform. Fig. 3.5 and Fig. 3.6 illustrate the agreement between the MBD model and OpenFAST for the blade and tower response, which are critical quantities in the overall dynamic response of the structure.

The close agreement between the responses of the MBD model and OpenFAST results, as shown in these figures, reinforces confidence in the ability of the MBD model to accurately capture the dynamic behaviour of the IEA-15 MW wind turbine. To quantify the error between the two models, the absolute differences in the mean and extreme values are compared, as detailed in Table 3.2 and Table 3.3. Furthermore, the normalised root mean square error (NRMSE), which evaluates deviations over the entire time-history response, is reported in Table 3.4. This analysis shows that the MBD model is in good agreement with OpenFast for normal operating conditions. This analysis focusses on wind speeds below the rated level, as turbine response at above-rated wind speeds is heavily influenced by the pitch controller.

The controller routine is an important aspect of wind turbine operation. The reference controller for the IEA-15 MW model is based on the ROSCO controller [41]. ROSCO is an open source wind turbine controller that offers a variety of functionalities, including tip speed ratio tracking generator torque controller, variable pitch controller, minimum pitch saturation, wind speed estimation, and a smoothing algorithm in near-rated operation. These modern controller features were found to significantly reduce the rotor thrust experienced by the structure.



(a) Blade in-plane deformations




(b) Blade out of plane deformations

Figure 3.5: Model Verification: Comparison of blade response


Mean Wind Speed (m/s)	δ_{TFA} (m)	δ_{TSS} (m)	$\delta_{BldOopDef}$ (m)	$\delta_{BldIpDef}$ (m)
3	-0.0339	0.0017	0.1845	-0.0066
5	-0.0305	-0.0037	-0.1585	0.0027
7	-0.0275	-0.0063	-0.2202	0.0063
9	0.0197	-0.0117	-0.1801	0.0126

Table 3.2: Absolute error in prediction of mean response from MBD model and OpenFast at different mean wind speeds


3.7. IEA-15MW Reference Wind Turbine



(a) Tower FA deformation response



(b) Tower SS deformation response



(c) Fourier spectrum of Tower SS response

Figure 3.6: Model Verification: Comparison of tower response

Mean Wind Speed (m/s)	δ_{TFA} (m)	δ_{TSS} (m)	$\delta_{BldOpDef}$ (m)	$\delta_{BldIpDef}$ (m)
3	-0.0190	0.0082	0.1305	-0.0283
5	-0.0307	0.0031	-0.1097	-0.0257
7	-0.0259	0.0095	-0.1998	-0.0039
9	0.0303	0.0016	-0.0794	-0.0126

Table 3.3: Absolute error in prediction of extreme response from MBD model and OpenFast at different mean wind speeds

Mean Wind Speed	δ_{TFA}^{NRMSE}	δ_{TSS}^{NRMSE}	$\delta_{BldOpDef}^{NRMSE}$	$\delta_{BldIpDef}^{NRMSE}$
3	0.3154	0.1154	0.2569	0.4652
5	0.7035	0.0780	0.2590	0.1012
7	0.6040	0.0637	0.4420	0.4895
9	0.5522	0.0681	0.4859	0.5320

Table 3.4: Normalised RMSE in prediction MBD model and OpenFast response at different mean wind speeds

3.8 Limitations of the MBD model

While the developed MBD model provides an efficient framework to integrate aerodynamics, controller dynamics, and structural dynamics, it inherently introduces certain simplifications that limit its fidelity to capture intricate physical phenomena. These limitations arise primarily due to the use of mode summation for flexible elements and the reliance on steady Blade Element Momentum (BEM) theory for aerodynamic force computations. The limitations introduced by such simplifications are highlighted below:

3.8.1 Mode Shape Summation Approximation for Flexible Elements

The use of mode summation to model the flexibility of the tower and blades introduces several assumptions that affect the accuracy of structural dynamics predictions:

- **Limited Number of Modes:** The method captures only a finite set of dominant mode shapes, neglecting higher-order modes that may contribute significantly

to transient or high-frequency dynamic responses.

- **Linear Elasticity Assumption:** Mode summation assumes small deformations and linear structural behaviour, making it less accurate for capturing large deflections, non-linear material properties, or geometric non-linearities that occur in extreme operating conditions.
- **Inability to Capture Localised Effects:** Since the structural response is represented using a global modal basis, localised deformations, stress concentrations, and blade root-hub interactions are not explicitly resolved.

3.8.2 Use of Steady BEM theory

The aerodynamic forces in the model are computed using steady BEM theory, which introduces the following limitations:

- **Neglect of Unsteady Aerodynamic Effects:** The steady-state assumption ignores unsteady aerodynamic phenomena such as dynamic stall, flow separation, and wake-induced oscillations, which play a critical role in load fluctuations and transient response.
- **Empirical Corrections for Tip and Hub Losses:** BEM theory relies on empirical models (e.g., Prandtl's tip loss correction) to account for wake contraction and blade tip effects, which lack the accuracy of more advanced vortex-based or CFD approaches.
- **Lack of Wake Interaction Modelling:** The model does not account for wake interactions in multi-rotor configurations or turbine-turbine wake effects, which are crucial for wind farm simulations and controller design.

Despite its limitations, the MBD modelling strategy is arguably the only feasible approach to integrate system-level dynamics in a computationally practical manner. Although high-fidelity CFD or finite element methods offer detailed insights into specific physical phenomena, they lack the scalability and efficiency required for full-system simulations. The MBD model developed in this study can capture coupled interactions between aerodynamics, structural flexibility, and control dynamics in a way that is computationally manageable and sufficiently accurate for system-level analysis.

3.9 Summary

The MBD model integrates data on structural characteristics, the controller algorithm, and wind inflow to forecast both structural response and power generation. This model serves as a valuable tool for evaluating wind turbine performance under various operating conditions. However, conducting dynamic analyses with the MBD model is computationally intensive, primarily due to the high computational cost of calculating aerodynamic loads using BEM theory and predicting dynamic responses using numerical time integration methods. This thesis investigates ML-based surrogate development strategies to mitigate these computational costs while maintaining accuracy. The verified MBD model of the IEA-15 MW wind turbine is used to analyse the behaviour of the turbine. This model will be used to generate synthetic datasets necessary for developing efficient surrogate models. By creating ML-based surrogates, which demand extensive datasets across diverse operating conditions, the MBD model emerges as the optimal means of generating valuable training data and facilitating real-time dynamic response predictions. The following chapters will explore the methodologies and techniques used in developing these surrogate models. With reference to the simplifying assumptions made during the model formulation, the MBD model is used to predict the wind turbine response only for normal operating conditions, as such models may not be best suited to predict the response under extreme conditions. Furthermore, as normal operating conditions tend to govern the fatigue life of the wind turbine, this thesis focusses on developing models that can enable quick fatigue analysis.

Chapter 4

Artificial Neural Network based surrogate model to predict aerodynamic loads

4.1 Motivation

As the blades become longer and more flexible, challenges such as increased deformations, fatigue damage, and the loads experienced by the system become more pronounced. Aerodynamic loads significantly influence the design, performance, and fatigue life of wind turbine blades. The high cost of calculating aerodynamic loads using Blade Element Momentum (BEM) theory results in the longer simulation time required for running MBD models. This chapter introduces a surrogate modelling approach aimed at accelerating aerodynamic load calculations. This approach offers two key advantages:

1. **Reduced Computational Cost:** The surrogate model can provide computationally efficient load estimations compared to the conventional approach of numerically solving BEM equations
2. **Potential for Virtual Sensing:** The model can serve as a virtual sensor to digitally measure aerodynamic loads using wind speed data measured using devices such as anemometers, LIDAR, or pitot tubes. The surrogate model can predict the aerodynamic loads on the blades, which are difficult to measure

on-site.

Using the validated MBD model developed in Chapter 3, this chapter presents the methodology to create the surrogate model, its integration into the MBD framework, and the evaluation of its accuracy and computational efficiency.

4.2 Background

The BEM theory is widely used to estimate the aerodynamic loads on wind turbine blades in many design and analysis applications. BEM theory combines information from blade element theory and momentum theory to calculate the angle of attack (AOA) in an airfoil section, which is required to calculate the aerodynamic forces. The computational expense of the BEM model arises from the time-consuming root-finding approach used to solve the governing equations. Previously, Fluck and Crawford [142] and Haghi and Crawford [143] have presented a stochastic solution to estimate aerodynamic loads using modified forms of polynomial chaos expansions. In these studies, the authors developed a surrogate model of the loads' statistics acting on the wind turbine rotor. These stochastic models have been found to significantly reduce computational time in predicting the output distribution and can be useful in statistical studies. However, neither of these studies considers the effect of controller algorithms, which significantly impact BEM output and wind turbine operation. In addition, the analysis in these studies was performed assuming a constant rotation speed, which limits their application. The model developed by Haghi and Crawford [143] aims to provide overall accurate aerodynamic load statistics, but the accuracy of point-to-point force estimation is questionable.

To address these challenges and improve computational efficiency, the primary objective of this chapter is to develop an Artificial Neural Network (ANN) based surrogate model of the BEM aerodynamic model. A data-driven and physics-based approach is utilised to create the surrogate model. The physics-informed surrogate can incorporate the physics-based constraints in the training process [136]. A weighing factor is introduced to optimally balance the contribution of the data-fitting objective and physics-informed criterion during the optimisation process. The ANN-based surrogate model aims to improve computational efficiency by replacing the root finding process in solving BEM equations with high-dimensional regression. The model is trained to predict the AOA at each blade node as a function of inflow




Figure 4.1: Control volume representation for momentum theory

parameters. The surrogate model is trained using temporal data to map inputs to outputs at each time step. This enables the surrogate model to be implemented in existing wind turbine models for faster load estimation. The computational advantage gained by using the surrogate model in place of the BEM aerodynamic model is analysed in this chapter. A general overview of the BEM theory is presented in the next section.

4.3 Blade Element Momentum theory

The BEM theory is widely used in MBD models to estimate aerodynamic loads on wind turbine blades. This theory assumes that the rotor operates in an ideal flow field, neglecting viscosity, compressibility, and three-dimensional effects. BEM theory combines information from blade element theory and momentum theory to compute the angle between the resultant wind velocity and the chord line of the airfoil, commonly termed the AOA at an airfoil section, which is then used to compute the aerodynamic forces. The momentum theory assumes a control volume section of the wind inflow. It applies the principle of conservation of momentum to estimate the total thrust (dT) and torque (dQ) acting on an annular disk of radius r and thickness dr as shown in Fig. 4.1. The thrust and torque estimates from momentum theory are presented in Eq. 4.1 and Eq. 4.2

$$dT = \rho U^2 4a(1-a)\pi r dr \quad (4.1)$$

$$dQ = \rho U 4a'(1-a)\Omega \pi r^3 dr \quad (4.2)$$

where ρ represents the mass density of inflow, U is the inflow velocity far upstream of the actuator disk, and Ω is the rotor speed. The terms a and a' represent the axial and tangential induction factors, which are the measure of induced linear and tangential wind velocity at the rotor plane due to the presence of the actuator disk. Momentum theory is useful for estimating the overall thrust produced but does not provide detailed information about the forces acting on individual blade elements. Conversely, the blade element theory divides the rotor or propeller into small blade elements and analyses the forces and moments acting on each element. The blade element theory predicts the thrust (dT) and torque (dQ) as a function of the aerodynamic properties of the airfoil and the geometrical relationship between various forces acting on an airfoil section. The equation for thrust and torque at a distance r from the blade root on an airfoil of chord length c and radial length dr using blade element theory is given by Eq. 4.3 and Eq. 4.4 where, B is the number of blades in the rotor, U_{rel} is the relative wind velocity which is the vector sum of wind velocity at the rotor and the wind velocity due to blade rotation, and C_l and C_d are the lift and drag coefficients of the airfoil section corresponding to the inflow angle ϕ .

$$dT = \frac{1}{2} B \rho U_{rel}^2 (C_l \cos(\phi) + C_d \sin(\phi)) c dr \quad (4.3)$$

$$dQ = \frac{1}{2} B \rho U_{rel}^2 (C_l \sin(\phi) - C_d \cos(\phi)) c r dr \quad (4.4)$$

At a blade node, the inflow angle (ϕ) is related to the AOA (α) through Eq. 4.5, where θ_p is the sectional pitch angle.

$$\phi = \theta_p + \alpha \quad (4.5)$$

Eq. 4.6 gives the sectional pitch angle at a blade node, where θ_{p0} is the blade pitch angle, which is a time-dependent parameter determined by the controller algorithm based on the region of operation of wind turbine and θ_T is the twist angle which is a geometrical property of the blade which varies along the blade length.

$$\theta_p = \theta_{p0} + \theta_T \quad (4.6)$$




Figure 4.2: Geometric representation of blade element theory

The geometrical relationship between inflow angle, wind velocity and induction factors is defined in Eq. 4.7, where λ_r is the ratio of the tangential speed of blade tip to the inflow wind speed given by Eq. 4.8. The geometric relation between various parameters in blade element theory is presented in Fig. 4.2.

$$\tan \phi = \frac{U(1-a)}{\Omega r(1+a')} = \frac{1-a}{(1+a')\lambda_r} \quad (4.7)$$

$$\lambda_r = \frac{\Omega r}{U} \quad (4.8)$$

The prediction of thrust and torque values obtained using the momentum and aerodynamic principles are equated to consider both the global inflow characteristics and local aerodynamic properties. Combining the information from momentum theory with blade element theory ensures that the forces and moments acting on each blade element can be determined while also considering the overall effect of the rotor or propeller on the flow. Combining this information gives the BEM theory's prediction of the feasible lift coefficient (C_l) as a function of AOA (α) for a constant tip

speed ratio (λ_r) through Eq. 4.9.

$$C_l = \frac{4}{\sigma'_r} \sin \phi \left(\frac{\cos \phi - \lambda_r \sin \phi}{\sin \phi + \lambda_r \cos \phi + \frac{C_d}{C_l} (\lambda_r \sin \phi - \cos \phi)} \right) \quad (4.9)$$

where σ'_r represents the solidity ratio, which quantifies the area of the annular ring covered by the blades, given by Eq. 4.10

$$\sigma'_r = \frac{Bc}{2\pi r} \quad (4.10)$$

Also, from the principles of aerodynamics, the lift coefficient is an intrinsic attribute of the airfoil geometry and is expressed as a function of the AOA (α) and the Reynolds number (R_e). Employing the aerodynamic characteristics of an airfoil ($C_l(R_e, \alpha)$) and BEM theory's forecast of the lift coefficient as defined in Eq. 4.9, in conjunction with the geometric relationship outlined in Eq. 4.5, facilitates the iterative determination of the AOA (α). The BEM theory equations can also be described in the form of induction factors through Eq. 4.11 and Eq. 4.12

$$a = \frac{\sigma'_r C_n(\alpha, \theta)}{\sigma'_r C_n(\alpha, \theta) + 4 \sin^2 \theta} \quad (4.11)$$

$$a' = \frac{\sigma'_r C_t(\alpha, \theta)}{4 \sin \theta \cos \theta - \sigma'_r C_t(\alpha, \theta)} \quad (4.12)$$

where, the normal force coefficient (C_n) and tangential force coefficient (C_t) are related to lift and drag coefficient by Eq. 4.13 and Eq. 4.14

$$C_n(\alpha, \theta) = C_l(\alpha) \cos(\theta) + C_d(\alpha) \sin(\theta) \quad (4.13)$$

$$C_t(\alpha, \theta) = C_l(\alpha) \sin(\theta) - C_d(\alpha) \cos(\theta) \quad (4.14)$$

As BEM theory does not explicitly account for vortex shedding effects at the blade tip, certain correction factors are proposed to consider this effect. Due to the lower pressure on the suction side of a blade, air tends to flow around the tip from the lower to the upper surface, reducing the lift force and power production in this region. This effect can be introduced in BEM equations using the tip loss factor proposed by Prandtl, given by Eq. 4.15, where R represents blade length [144].

$$F_{tip} = \frac{2}{\pi} \left[\exp \left(-\frac{B(R-r)}{2r \sin \theta} \right) \right] \quad (4.15)$$

This correction factor quantifies the reduction in aerodynamic forces at a distance r along the blade span as an effect of the losses at the blade tip. Similarly, the hub loss is computed by Eq. 4.16, where R_{hub} is the hub radius.

$$F_{hub} = \frac{2}{\pi} \left[\exp \left(-\frac{B(R_{hub} - r)}{2r \sin \theta} \right) \right] \quad (4.16)$$

These two factors combined give the total vortex-based losses such that,

$$F = F_{tip} \times F_{hub}$$

The combined loss factor can then be used to compute axial and tangential equations using Eq. 4.18 and Eq. 4.17

$$a = \frac{\sigma'_r C_n(\alpha, \theta)}{\sigma'_r C_n(\alpha, \theta) + 4F \sin^2 \theta} \quad (4.17)$$

$$a' = \frac{\sigma'_r C_t(\alpha, \theta)}{4F \sin \theta \cos \theta - \sigma'_r C_t \alpha, \theta} \quad (4.18)$$

The BEM equations obtained by equating blade element theory and momentum theory are valid for normal wind turbine operating cases with $a < 0.5$. For higher induction factors, the flow patterns through the rotor become much more complicated than that predicted by the momentum theory [144]. To ensure BEM theory's validity in such cases, Galuert's empirical correction [145] can be applied for flows resulting in $a > 0.4$,

$$a = \frac{1}{F} [0.143 + \sqrt{0.0203 - 0.6427(0.889 - C_T)}] \quad (4.19)$$

where,

$$C_T = \begin{cases} 4a(1 - a), & \text{if } a \leq a_c \\ 4(a_c^2 + (1 - 2a_c)a), & \text{if } a > a_c \end{cases} \quad (4.20)$$

where, $a_c = 0.2$. This empirical relation is further modified by Buhl [146] to incorporate the tip loss effects in computing thrust coefficient using Eq. 4.21

$$C_T = \frac{8}{9} + \left(4F - \frac{40}{9} \right) a + \left(\frac{50}{9} - 4F \right) a^2 \quad (4.21)$$

The BEM model implemented in the MBD model accounts for both tip losses and high induction factor corrections to ensure the accuracy of the estimated aerodynamic loads. The steady BEM theory used in this thesis interpolates the wind speed at the blade nodes from the TurbSim generated 3D wind field to estimate aerodynamic loads based on the quasi-static aerodynamic properties (i.e., lift and drag) of the blades.

Though conceptually simple, the BEM equations can be challenging to solve reliably and efficiently with high precision. Many solution approaches exist for numerically converging the AOA and axial and tangential induction factors, but they generally suffer from a lack of robustness in some regions of the rotor blade design space or require significantly increased complexity to guarantee the convergence [147]. In this context, Ning et al. [148] has developed a robust method to solve the BEM equations with guaranteed convergence. The significant contribution of this work was to parameterise the BEM equations through a single unknown variable, local inflow angle (ϕ), by defining a residual function given by Eq. 4.22. Solving the residual equation by equating $\mathcal{R}(\phi) = 0$ gives the solution for the local inflow angle (ϕ).

$$\mathcal{R}(\phi) = \frac{\sin \phi}{1 - a} - \frac{1}{\lambda_r} \frac{\cos \phi}{1 + a'} \quad (4.22)$$

Using this approach, the BEM equations can be solved using one-dimensional root-finding approaches, which are more efficient. In this method, the BEM equations are segmented into three regions: momentum region, empirical region and propeller break region, and solutions for these operating regions are proposed. An overview of the approach presented by Ning et al. [148] is outlined below, while the interested reader can refer to Ref. [148] for the detailed algorithm.

1. Assume an initial inflow angle ϕ
2. Compute local AOA using Eq. 4.5
3. Estimate the lift and drag coefficients for the local AOA ($C_l(\alpha), C_d(\alpha)$)
4. Compute the normal and tangential force coefficients ($C_n(\alpha, \theta), C_t(\alpha, \theta)$)
5. Evaluate the axial and tangential induction factors (a, a')
6. Check the residual error using Eq. 4.22

7. Try another value of local inflow angle until residual error (Eq. 4.22) approaches zero (using root-finding method)

In the MBD model framework proposed in this work, the solution for the residual equation is found using Brent’s approach [149]. This iterative root-finding approach for finding the inflow angle brings about the high computational cost of solving BEM equations. An approximate relation between the inflow parameters and the AOA could exist for a given airfoil section, assuming constant geometry and deterministic curves for aerodynamic coefficient. A surrogate model developed to capture this relationship could replace the root-finding method with regression, significantly reducing the computational cost of obtaining the solution. In this work, an ANN-based surrogate is developed to achieve this objective. The mathematical formulation of ANN is discussed in the next section.

4.4 Development of machine learning surrogate for BEM model

An ANN model takes an input feature vector x and defines a suitable hypothesis $h_\theta(x)$ to match the output y . The NN hypothesis $h_\theta(x)$ is designed using a series of interconnected NN layers, where an NN layer is characterised by the number of neurons and an activation function. A neuron, a building block of an NN architecture, is characterised by a specific weight (w) and a bias (b) value. A neuron receives the input x and produces the output as the weighted sum of the input and the bias value activated by an activation function. This mathematical operation for a single neuron is presented in Eq. 4.23, where $f : R \rightarrow R$ is an activation function.

$$h_\theta(x) = f(wx + b) \quad (4.23)$$

The activation function $f(\cdot)$ introduces non-linearity in the NN formulation, allowing the network to model more complex input-output relationships and shape the decision boundaries. The activation function selection is essential as it controls the network’s performance and learning capabilities; tanh, sigmoid, and ReLU are the most commonly used activation functions. In the context of a NN layer, in Eq. 4.23, the NN layer receives an input matrix (x) with the dimension of $n_F \times n_S$ where n_F is the number of input features and n_S is the number of samples or observations and

produces an output ($h_\theta(x)$) of dimensions $n_N \times n_S$ where n_N denotes the number of nodes in the layer. The weight matrix and bias vector have dimensions of $n_N \times n_F$ and $n_N \times 1$, respectively. Multiple NN layers can be arranged in a sequence such that the output of one layer serves as the input to another layer. For a general M-layer NN with a non-linear activation function $f(\cdot)$ and a parameter matrix θ representing the learnable parameters, the NN hypothesis is defined in Eq. 4.24:

$$h_\theta(x) = f_M(\theta_M, f_{M-1}(\theta_{M-1}, \dots, f_2(\theta_2, f_1(\theta_1, x)))) \quad (4.24)$$


Adding additional layers to the NN can help the model capture more complex relationships between input and output; however, it is also more likely to overfit the training data, where the model becomes too specific to the training data and performs poorly on unseen data. To address this, a penalty or regularisation term is often added to the objective function, discouraging overfitting and promoting generalisability. The ANN training process revolves fundamentally around optimisation. A suitable metric that aligns with the desired model behaviour and captures the errors between predicted and actual outputs must be defined. The optimisation algorithm iteratively adjusts the weights and biases, optimising the model's performance based on the chosen loss function. Optimisation algorithms such as the stochastic gradient descent method (SGDM) [150] or adaptive moment estimation (Adam) [151] are commonly used to update the model parameters during the training process. Selecting a good set of input features and a suitable performance metric is not trivial.

4.4.1 Training Data Generation


In this study, the surrogate model designed to replace the BEM aerodynamic model is designed to predict AOA as output parameter. As discussed in section 4.3, the AOA is a function of inflow properties, control input, blade geometry, induction factors, and blade vibrations. Of all these variables, the surrogate model is designed to model the AOA at a blade node as a function of wind speed and control parameters (rotor speed and pitch angle) by assuming deterministic blade geometry and steady airfoil properties. The controller inputs (pitch angle and rotor speed) are explicitly considered as the input features as this input choice extends the applicability of the developed surrogate model for different controller algorithms than the ones used for generating the training data. When selecting the input features for the surrogate

4.4. Development of machine learning surrogate for BEM model


model, only the parameters that can be measured during the operation of the wind turbine are selected. This choice of input features will ensure that the developed model is useful in practical industry applications. Training data is generated using the numerical model presented in Chapter 3 that employs the ROSCO controller to regulate wind turbine operations [152]. To this end, the wind velocity in the x and y direction in the blade local coordinate system, rotor speed, blade pitch angle and generator azimuth angle are considered as the input features. The variation of the AOA with respect to the input features for steady inflow conditions is presented in Fig. 4.3. For generating the training data, fifty mean wind speed values ranging



(a) Variation of AOA with wind speed



(b) Variation of AOA with rotor speed



(c) Variation of AOA with blade pitch angle

Figure 4.3: Variation of AOA with input features

from cut-in to cut-out wind speeds are generated using the Sobol sequence. Sobol

sequence is used because it is known to have a low discrepancy and better distribution properties. The scatter plot of the mean wind speeds generated is presented in Fig. 4.4. Fig. 4.4 shows the efficiency of the Sobol sequence to effectively fill the pa-




Figure 4.4: Mean wind speed values sampled using Sobol sequence for generating the training data

parameter space. Random wind realisations are generated using the Turbsim module [140] for these sampled mean wind speed values. TurbSim is a stochastic, full-field, turbulent-wind simulator that uses statistical models to numerically simulate a time series of wind speed vectors in a two-dimensional vertical rectangular grid fixed in space. Wind speed time histories are generated for a turbulence intensity ranging from 5% to 10%, randomly selected for the wind speeds generated. This study assumes a normal turbulence model (NTM), defined in the IEC recommendations, and the power law exponent of 0.2 is assumed to correspond to a normal wind profile [79]. For this set of parameters, wind speed realisations are simulated considering two random seed values for 100 seconds. The response of the wind turbine for these conditions is evaluated, and the input and output features are recorded. For turbulent conditions, the response for the first 50 seconds is discarded to remove the transient effects. The response for the first random seed is selected for training, while the response corresponding to the second random seed is used for validation. The hold-out validation scheme is used during model training.

4.4.2 Definition of Objective Function

In training an ANN surrogate model, the objective function plays a key role. An objective function is a metric that quantifies the agreement between the model output and the surrogate model prediction. The one-half-squared error function is most commonly used in regression-based surrogate models. Following standard practice, the one-half squared error function presented in Eq. 4.25, is used as the objective function for training the data-driven ANN surrogate.

$$J_d(\theta) = \frac{1}{2} \|h_\theta(x) - y\|^2 \quad (4.25)$$

A regularisation term is added to the objective function to ensure generalisability and avoid overfitting. The objective function for a m set of input-output pairs is presented in Eq. 4.26, where $g(\theta)$ represents the regularisation term.

$$J_d(\theta) = \frac{1}{m} \sum_{i=1}^m J(\theta; x^{(i)}, y^{(i)}) + \lambda g(\theta) \quad (4.26)$$

In this study, the weight decay term is used for regularisation as it discourages the neural network (NN) from relying heavily on individual weights and encourages it to use a more distributed representation of the input data. This leads to a more robust model and makes the model less sensitive to these small perturbations in input values. The final objective function used for training the data-driven ANN surrogate model is shown in Eq. 4.27, where $\|\cdot\|$ represents the Frobenius norm of the weight matrix and n_L is the number of layers.

$$J_d(\theta) = \frac{1}{m} \sum_{i=1}^m \|h_\theta(x_i) - y_i\|^2 + \lambda \sum_{j=1}^{n_L} \|\mathbf{W}^{(j)}\|^2 \quad (4.27)$$

Although the objective function presented in Eq. 4.27 can effectively capture the complex patterns in the data, it does not use any physical knowledge or governing principles of the BEM theory in the optimisation process. This approach can be difficult in regions not covered in the training data or in areas where available data are sparse and noisy. In these situations, physics-informed neural networks (PINNs) have an advantage over the purely data-driven approach, as these models can leverage the prior knowledge of the system to guide the optimisation process to learn the physically more meaningful and reliable solutions. In the context of the present

problem, as presented in Section 4.3, axial and tangential induction factors are important parameters that govern the AOA in a blade section. These parameters are the latent variables of the system which cannot be explicitly measured. The PINN approach implemented in this study defines additional constraints using these latent system parameters. Explicitly considering these parameters in the optimisation process ensures a more accurate representation of the BEM theory's principles. The objective function to enforce the physics-informed constraint is defined in Eq. 4.28:

$$J_p(\theta) = \frac{1}{m} \sum_{i=1}^m \mathcal{R}^2(h_\theta(x_i)) \quad (4.28)$$

The loss function defined in Eq. 4.28 ensures that the residual term (see Eq. 4.22), which determines the accuracy of the BEM solution, is minimised during the model training. Although such physics-informed approaches appear to be straight-forward and have yielded remarkable solutions for various applications, the effect of additional residual terms is poorly understood and, in several cases, can lead to unstable or erroneous predictions [153, 154]. Therefore, to balance the interplay between data fit and physics-informed constraints, a weighing factor term (β) is introduced, which maintains a balance between the two loss terms in the final compositional form. The resulting optimisation functional form for the PINN training is presented in Eq. 4.29.

$$\operatorname{argmin}_{\theta} \left(\beta J_d(\theta) + (1 - \beta) J_p(\theta) + \lambda \sum_{j=1}^{n_L} \|\mathbf{w}^{(j)}\|^2 \right) \quad (4.29)$$

For a data-driven neural network (DDNN) surrogate, the optimisation functional form is presented in Eq. 4.30

$$\operatorname{argmin}_{\theta} \left(J_d(\theta) + \lambda \sum_{j=1}^{n_L} \|\mathbf{w}^{(j)}\|^2 \right) \quad (4.30)$$

In this study, the above optimisation problem is solved using the Adaptive Moment Estimation (Adam) optimiser [151]. The weighing factor (β) is modelled as the model hyperparameter, and its optimum value for each blade node is obtained through the hyperparameter tuning using the Bayesian optimisation approach. The hyperparameter tuning process is presented in detail in the next section.

Hyperparameter	Parameter bounds
Number of neurons in layer	[1,64]
Regularisation factor	[0, 1]
Learning rate	[1e-5 1e-2]
Weighing factor	[0 1]

Table 4.1: Hyperparameters and exploration range

4.4.3 Hyperparameter Tuning

In neural network models, the hyperparameters represent key model parameters that are determined prior to the model training process. Hyperparameter tuning is an important step in developing an efficient NN, as these parameters directly impact the performance and generalisation ability of the model. Effective hyperparameter tuning is critical to achieving an efficient NN model. This process involves selecting the optimal combination of hyperparameter values, which ensures high accuracy, faster convergence, and improved generalisation on unseen data. In this study, the Bayesian optimisation technique in Matlab® is implemented. Bayesian optimisation operates by constructing a probabilistic surrogate model of the objective function. This model facilitates informed decisions about the generation of new points within a predetermined range defined for the set of hyperparameters. This process involves an acquisition function that balances exploring new regions in parameter space and exploiting information at the current point.

This study considers a four-layer deep NN architecture with *tanh* activation function for each blade node. Hyperparameter tuning is performed across all fifty nodes corresponding to the IEA-15MW wind turbine blade, where each node represents a unique airfoil section. The number of neurons in each layer, the regularisation factor, the learning rate, and the weighing factor were chosen as hyperparameters. Table 4.1 details the ranges considered for these hyperparameters.

Within this optimisation framework, the *expected-improvement-plus* acquisition function is utilised. Notably, the results of hyperparameter tuning corresponding to the weighing factor show interesting insights regarding the physics-informed constraints. The optimised value for the weighing factor corresponding to the fifty blade nodes is presented in Fig. 4.5, which shows that the weighing factor has attained a range of values.




Figure 4.5: Optimised weighing factor values for different blade nodes

Some of the blade nodes have a value of the weighing factor closer to 1, suggesting that the data-fit constraint provided more valuable information than the physics-informed constraint. This indicates that the PINN approach may not always lead to better performance in all circumstances. A careful consideration of the weighing factor is required for the specific problem under investigation. The final NN model for the optimised set of hyperparameter combinations is trained using the Adam optimiser. The results of training using the DDNN and PINN approach are presented in the next section.

4.5 Numerical Results

This section presents a comprehensive analysis of the surrogate models developed using the methodology presented in previous sections. This section presents a quantitative assessment of the performance of the surrogate models in terms of computational advantage and the accuracy of the predictions. Nine nodes from the blade root to the tip section were selected, and the surrogate models corresponding to these nodes were designed. These nodes represent different operating conditions on a wind turbine blade and different shapes of airfoil sections. A PINN and DDNN surrogate was developed for each of the selected blade nodes, and their performance was compared against each other. The NN-based surrogate is validated against the

Blade Node	Surrogate Error PINN MAE	Surrogate Error DDNN RMSE	Surrogate Error PINN MAE	Surrogate Error DDNN RMSE
Node 05	0.0041	2.760e-05	0.0043	2.854e-05
Node 10	0.0017	4.623e-06	0.0018	1.130e-05
Node 15	0.0026	1.067e-05	0.0020	5.605e-06
Node 20	0.0015	3.331e-06	0.0017	4.617e-06
Node 25	0.0022	7.453e-06	0.0018	4.763e-06
Node 30	0.0028	1.302e-05	0.0020	6.361e-06
Node 35	0.0173	8.048e-04	0.0171	7.901e-04
Node 40	0.0038	3.757e-05	0.0039	3.648e-05
Node 45	0.0035	3.227e-05	0.0040	3.819e-05

Table 4.2: Data Fitting Errors

BEM aerodynamic model for steady and turbulent conditions. The accuracy of the developed surrogate models is quantified using Mean Absolute Error (MAE) and Mean Square Error (MSE) metrics. MAE is a robust indicator of the average magnitude of the errors between the surrogate model’s predictions and the corresponding outputs from the BEM aerodynamic model. By focussing on the absolute values of these errors, MAE provides a measure of the surrogate’s accuracy in approximating the BEM model’s behaviour. Additionally, the MSE quantifies the squared discrepancies between the surrogate’s predictions and the actual values of the AOA obtained from the BEM model. MSE, while sensitive to larger errors due to the squaring operation, offers valuable insights into the surrogate model’s precision and the extent to which it captures minor and major variations in the data. The choice of these particular metrics is underpinned by their effectiveness in assessing the surrogate model’s ability to replicate the behaviour of the BEM aerodynamic model across a range of operating wind speeds. A comprehensive analysis of the surrogate model’s performance is presented using MAE and MSE metrics. The magnitude of error at the selected blade nodes is given in Table 4.2 to highlight the accuracy of the surrogate model.

The maximum MAE values achieved using the PINN and DDNN surrogates are of the order of 10^{-2} , demonstrating that both models can predict the AOA at any blade node with a discrepancy of less than 0.5° . This shows the precision of the surrogate model developed in this study. These error metrics are evaluated for DDNN and PINN to compare their performance. From Table 4.2, it is observed that of the nine nodes selected in this study, five nodes performed better under




Figure 4.6: Comparison of physics-based loss for DDNN and PINN surrogate

the PINN approach, while the remaining four models performed better under the DDNN approach in terms of the data fitting objective.

To enhance model accuracy assessment, Fig. 4.6 displays the values of the physics-based loss function. This figure shows that five of the nine selected nodes exhibited significantly lower loss values, while two achieved comparable loss values defined using Eq. 4.28. In total, seven of the nine chosen nodes demonstrated improved accuracy in physics-based constraints using the PINN approach, with only two nodes surpassing the performance of the PINN method. Furthermore, the primary benefit of the PINN surrogate is its ability to generalise and find the most optimal solution satisfying the underlying physics-informed objective function. Fig. 4.6 shows that PINN has a lesser error in these latent system parameters, which are equally important to accurately predict aerodynamic loads.

To further demonstrate the surrogate model's capability, the PINN surrogate's time-history response is compared against the BEM model for a range of wind speeds from cut-in to rated wind speeds for both steady wind and turbulent wind with 10% turbulence intensity. Fig. 4.7 shows the ability of the developed surrogate to accurately predict AOA at each time step over the range of wind speeds in both steady and turbulent conditions. This capability of the surrogate model allows it to be integrated with the state-of-the-art wind turbine models used for dynamic




Figure 4.7: Comparison of surrogate model estimate of AOA with BEM theory

analysis of the wind turbine structure. In addition, the coefficient of determination (R^2) for a blade node is presented in Fig. 4.7. The R^2 value is a statistical measure of the goodness of fit of the regression model. An R^2 value of more than 0.99 is achieved, further demonstrating the accuracy of the surrogate model. Fig. 4.7 also represents the error distribution in the predicted and target values. A uniform error distribution is observed against the target values over the range of AOAs, which shows that the developed surrogate model does not suffer from heteroscedasticity commonly discussed in the literature [155]. The estimated values of aerodynamic forces using BEM theory, PINN surrogate model, and DDNN surrogate are compared in Fig. 4.8. For both surrogate models, the percentage error in predicting aerodynamic forces of less than 5% is observed in all cases.

The accuracy of the surrogate models and their capability in predicting the AOA and aerodynamic forces across different wind speeds and operating regions is highlighted through Fig. 4.7 and Fig. 4.8.

4.5.1 Computational Advantage of Surrogate Models

The main aim of surrogate models is to achieve high accuracy in prediction at a lower computational cost. The accuracy of the developed surrogate models is highlighted




Figure 4.8: Comparison of surrogate model estimate of Aerodynamic forces with BEM theory

in the previous section, while their computational advantage is discussed in the current section.

The NN-based surrogate models developed in this study can predict more than 80000 observations/sec, which is more than 40 times faster than the conventional way of predicting AOA through iterative root finding methods. The prediction capacity values are computed through the profiler analysis of the MBD presented in Chapter 3. The comparison of computational cost is made against Ning's method [147] of solving BEM equations, which is the most robust method and is also implemented in OpenFast [37]. Since the surrogate model can estimate the AOA at every time step, the model is integrated with the MBD model of the wind turbine presented in Chapter 3. The use of a surrogate model in place of BEM theory has resulted in a 35% reduction in the overall simulation time of the MBD model. The reduction in computational time is quantified by comparing the execution time of the numerical model required to predict the 600-second response to a turbulent wind inflow using Ning's method to solve the BEM equations and by using the surrogate models. The reduced computational time allows for running large sets of simulations for reliability estimation, fatigue analysis, or optimisation problems. The proposed surrogate model primarily benefits the site-suitability studies where the MBD model has to be evaluated for the load envelope at multiple sites to select

the most optimum installation location. All simulations in this study are performed on an 8-core Intel Xeon CPU with a clock speed of 3.8GHz using 32GB RAM and running on Microsoft Windows 10 Pro.

4.6 Discussion

In this chapter, an NN-based surrogate approach was presented aimed at estimating aerodynamic forces on wind turbine blades under stochastic inflow conditions. This methodology showcases the potential of combining data-driven and physics-informed techniques to train the NN models. The mean wind speeds were generated using the Sobol sequence, while TurbSim was used to generate turbulent wind speed realisations. The MBD model presented in Chapter 3 was used to generate the data for training the surrogate models.

This study extends the scope of surrogate modelling by considering physics-informed constraints and data-fit objective functions. A weighing factor that balances the relative importance of these two objectives was introduced and optimised using the Bayesian optimisation approach. A quantitative assessment of the surrogate models' accuracy showcased their ability to closely approximate the BEM aerodynamic model. In most cases, the physics-informed surrogate resulted in a more accurate representation of the BEM model predictions, offering an advantage over the data-driven methodologies.

One of the defining features of the proposed surrogate models is their ability to predict the AOA and, in turn, the aerodynamic forces at every time step. This capability of the surrogate model allows for integrating these models with the MBD models used for the dynamic analysis of wind turbines. The estimation of load time-history response enables these surrogates to be used for fatigue analysis, which is not possible with the surrogate models developed for statistical mapping of input and output domains as highlighted by Haghi and Crawford [143]. In addition, a significant computational advantage is obtained by using these surrogate models. In standalone aerodynamic load computations, a remarkable forty-fold speedup was observed compared to the conventional root-finding approach to solve BEM equations. Replacing the BEM model with the NN-based surrogate in the MBD model used for the dynamic analysis in this study led to an overall reduction of 35% in the computational time. Regarding computational efficiency, both data-driven and

physics-informed surrogates have a similar performance. While these surrogates are designed for airfoil shapes specific to the IEA-15MW wind turbine, the methodology is adaptable to any wind turbine design.

Mixed performance was observed between data-driven and physics-informed surrogate models in the investigation involving nine different blade nodes and airfoil shapes. The data-driven model exhibited superior performance in two cases, while the physics-informed surrogate performed better in seven cases. While PINN has demonstrated superior performance compared to DDNN in various works found in the literature [136, 156, 157], in our specific problem, it does not provide a notable advantage. This may be attributed to the fact that both the data-driven and physics-informed goals in this research tend to yield similar results since the data is sourced from a deterministic numerical model. PINN is generally more effective than DDNN for dynamic systems [158, 156] and experimental datasets [159]. This result highlights the crucial role of choosing the right type of neural network, whether data-driven or physics-informed, for surrogate modelling. It further emphasises that this selection should be tailored to the particular needs of each problem, rather than adopting a universal approach. While the findings highlight the versatility of data-driven networks in specific scenarios, it is imperative to recognise the enduring value of physics-informed networks within the broader context of aerodynamics and surrogate modelling. However, irrespective of the distinction between data-driven and physics-informed, the neural network-based surrogate approach presented in this study offers an advantage over the conventional surrogate modelling techniques in terms of computational time and accuracy.

One of the limitations of this study is the absence of stall correction and unsteady dynamic effects in the implemented BEM model. This restricts the model's ability to predict responses in highly turbulent flows, particularly in the stall region during flow separation. However, given the model's high accuracy in predicting the angle of attack for the chosen operating conditions suggests that such advanced phenomena could be captured if adequately represented in the training data. Future research could explore the use of more advanced BEM models for developing surrogate models.

In conclusion, the methodology proposed in this study lays the foundation for surrogate modelling of wind turbines. In particular, developing a surrogate model to predict the response of a wind turbine blade subjected to the aerodynamic loads can

bypass the computationally intensive MBD models with a more efficient data-driven surrogate. This aspect will be addressed in the next chapter.

Chapter 5

Long Short-Term Memory model to predict blade deformation response

5.1 Motivation

Aerodynamic forces are the main external loads driving the component design. The computational cost of computing aerodynamic forces in an MBD framework can be alleviated using the surrogate modelling approach presented in Chapter 4. Once the aerodynamic forces are computed, the next challenge is to predict the dynamic deformation response of the system. Traditionally, numerical integration schemes through time-stepping-based algorithms are used to calculate the deformation response. Often, these numerical schemes require small step sizes for convergence, resulting in a large number of computations. Consequently, the evaluation of the deformation response contributes significantly to the overall computational time required to analyse the structure. Since computing the deformation response needs to ensure temporal dependence, ANN-based surrogates are not suitable for such problems, as they treat each point independently. To address this, a Long Short-Term Memory (LSTM) model is used to model the dynamics of the system. LSTM models are particularly well-suited for capturing short- and long-term dependencies within sequential data, making them suitable for time series analysis. These models learn temporal patterns by adjusting fixed and recurrent weights. This approach

for predicting the deformation response offers a significant computational advantage compared to dynamic analysis through numerical integration, enhancing efficiency without compromising accuracy.

5.2 Background

The LSTM method has shown its effectiveness in solving a wide range of problems involving sequential data. Their ability to extract features from sequential data has been widely used in big data analysis [160], image processing [161], and speech recognition [162]. These models have also been implemented for wind turbine power forecasting [163, 164, 165] and damage detection [166, 167, 168]. Using a similar approach, the primary focus of this work is to predict the deformation response of the tip of a wind turbine blade when subjected to time-varying aerodynamic loads. Blade tip deformations are chosen as the target variable as it is one of the important design considerations, and it can also be used to predict the blade root moments using the moment-curvature relationship. In addition, accurately measuring blade deformations during wind turbine operation is a challenging task; these challenges in blade deformation measurement are highlighted by Liew et al. [169]. Dimitrov and Göçmen [170] in their study further addresses this issue and highlights the lack of studies addressing the challenges in the measurement and prediction of blade deformations. Dimitrov and Göçmen [170] have predicted blade deformation using the LSTM method, where the LSTM model is treated as a virtual sensor using the inputs from the supervisory control and data acquisition (SCADA) data and the load measurement system. To get the best performance out of an LSTM model, the optimum selection of input features and network architecture is of paramount importance. In their study Dimitrov and Göçmen [170], have used a wrapper type trial and error based approach for feature selection, which could be computationally prohibitive when a large number of input features are available. In this context, a mutual information-based minimum redundancy maximum relevance algorithm is used to select the optimum set of input features. The performance of the commonly used recursive feature selection method with exhaustive parameter space is compared, and insights into the computational cost of the respective methods and achieved accuracy are also presented. A combination of wrapper-type and filter-type approaches for optimal feature selection is presented. After feature selection, the Bayesian optimisation model is used to tune the model hyperparameters. A

plateauing validation loss-based learning rate scheduling approach is used, which ensures that an optimum convergence is achieved.

5.3 LSTM Algorithm

An LSTM model is part of the Recurrent Neural Networks (RNN) family, designed to work with sequential data and capture the long-term and short-term dependencies between the input and output variables. An RNN processes sequential data by maintaining the memory of previous inputs through an internal state, unlike traditional feedforward neural networks, which process each data point independently. RNNs employ feedback mechanisms, where the network's output at each step is fed back as input, facilitating sequential data modelling by incorporating past information into future predictions [171]. This recursive nature enables RNNs to capture the temporal dependencies inherent in sequential data. Traditionally, RNNs suffer from vanishing gradients, where the algorithm struggles to preserve the internal state for long-range sequences [172]. This phenomenon, known as a vanishing gradient, is widely encountered in machine learning. It refers to a state where the gradient of the loss function with respect to the learnable parameters becomes so small that it results in slow or ineffective learning. The LSTM architecture alleviates this problem by using various gates to control the flow of information [173]. These gates control the flow of gradients during backpropagation, allowing LSTM networks to propagate gradients over long sequences and through multiple layers effectively. An LSTM unit consisting of all these features is known as a memory cell. Selective data retention and elimination are performed as the data passes through these gates. Based on the task of each gate, the model architecture is defined through the following terminologies:

- **Input gate:** This gate receives the model features at the current time step and the hidden state predicted at the previous time step as input. This gate combines these two inputs to create a candidate for storing new information. Mathematically, this operation is represented in Eq. 5.1

$$i_t = \sigma_g(W_i \times x_t + R_i \times h_{t-1} + b_i) \quad (5.1)$$

- **Forget gate:** The input to the forget gate is the same as the input gate. This gate combines the information from input features at the current time step

with the hidden state from the previous time step and estimates what past information is no longer relevant. This operation is presented in Eq. 5.2

$$f_t = \sigma_g(W_f \times x_t + R_f \times h_{t-1} + b_f) \quad (5.2)$$

- Cell state: The cell state in an LSTM unit undergoes updates influenced by the output of the forget gate and the input gate, which determine the significance of new and existing memory. A cell state is the actual memory of the LSTM, which holds information across many steps. The cell state acts as the long-term memory of the unit. This operation is represented through Eq. 5.3 and Eq. 5.4

$$c'_t = \sigma_g(W_c \times x_t + R_c \times h_{t-1} + b_c) \quad (5.3)$$

$$c_t = f_t \cdot c_{t-1} + i_t \cdot c'_t \quad (5.4)$$

- Output gate: The output gate combines the cell state updated at the current time step with the previous hidden state, creating the current hidden state, predicting the network's output (hidden state) at the current time step. This gate acts as the short-term memory of the unit. Eq. 5.5 and Eq. 5.6 mathematically represent this operation

$$o_t = \sigma_g(W_o \times x_t + R_o \times h_{t-1} + b_o) \quad (5.5)$$

$$h_t = o_t \cdot \sigma_c(c_t) \quad (5.6)$$

In Eq. 5.1 to Eq. 5.6, x represents the feature vector of the LSTM model, while h denotes the hidden state, a crucial output of the LSTM unit. The subscript t denotes the temporal position of the vectors, where t refers to the current time step, $t - 1$ refers to the previous step, and so forth. The symbols i , f , c , and o correspond to the input gate, forget gate, cell state, and output gate, respectively. In each gate, W signifies the fixed weights, R denotes the recurrent weights, and b represents the bias term. The σ symbol signifies the sigmoid activation function, ensuring non-linearity in the input-output transformation and constraining the gate's output within the interval $[0,1]$. The sigmoid activation function is given by Eq. 5.7

$$\sigma(x) = \frac{1}{1 + e^{-x}} \quad (5.7)$$




Figure 5.1: LSTM Architecture (Source:[174])

Multiple LSTM units can be stacked together to capture complex relationships in the data. A sketch of a typical LSTM memory cell and data flow is presented in Fig. 5.1, adapted from Ref. [174].

5.4 Parameter Selection and Training Data Generation

This work aims to predict the dynamic time history of blade tip deformations. From structural dynamics theory, the deformation response can be deterministically evaluated for a linear deterministic system if the forcing function can be completely defined. To ensure generalisation ability and the model's ability to provide accurate results, force time history covering a range of possible values encountered during the operation should be simulated. During the operation of a wind turbine, the major dynamic forces arise from the aerodynamic loads acting on the system. Variations in other load components, such as gravity loads, remain relatively well defined. It is assumed that the effect of gravity forces will lead to a deterministic deformation pattern, which will be learned by the algorithm. The operational parameters, including the rotor speed, the azimuth angle of the generator, and the angle of the blade, are considered input features that capture this effect. In addition to these operational parameters, aerodynamic forces are also considered for the prediction of deformation.

The aerodynamic forces acting on the blade are a function of the turbulent inflow and the aerodynamic properties of the wind turbine blade. In the current

Parameter	Lower bound	Upper bound
Mean wind speed (m/s)	3	25
Turbulence intensity (%)	0.025	$\frac{0.18}{U} \left(6.8 + 0.75 \cdot U + 3 \left(\frac{10}{U} \right)^2 \right)$
Power law exponent	0	$\alpha_{ref,UB} + 0.4 \left(\frac{R}{z} \right) \left(\frac{U_{max}}{U} \right)$

Table 5.1: Environmental parameters used and their limits

investigation, the blade geometry and aerodynamic characteristics are assumed to remain constant during the operation. Therefore, the variation in aerodynamic forces can be captured by simulating a range of environmental parameters. Dimitrov et al. [102] has studied the effect of environmental parameters on aerodynamic loading, especially fatigue response. The main factors governing aerodynamic loads have been found to be the mean wind speed, the turbulence intensity, and the power law exponent. Also, these parameters are used by the IEC to define a particular site class, so selecting these input parameters will make this model applicable to all the different IEC site classes. In this work, the mean wind speed is selected as the only independent input parameter, and the turbulence intensity TI and the power law exponent α are modelled as the dependent variable of the mean wind speed. The mean wind speed values are first generated using the Sobol indices, and using the equations presented in Table 5.1, the parameter bounds for TI and α are determined. Here R denotes the rotor radius and z is the hub height. For the IEC-15MW wind turbine considered in this turbine, the parameters R and z assume values of 120m and 150m, respectively. Further, the parameter $\alpha_{ref,UB}$ denotes the reference upper bound of wind shear corresponding to the mean hub-height wind speed of 15m/s. This parameter is defined according to the IEC standards [79] and takes a value of 0.22. The random samples for these parameters are simulated using these parameter bounds along with Sobol indices. During the operation of wind turbines, the stochastic nature of wind inflow acts as one of the major sources of uncertainty. To take this into account, a total of 15 random seeds are considered in combination with environmental parameters. Of these 15 samples, nine are used for training and six are used for testing and validation.

5.5 Feature Selection

Optimum feature identification is of paramount importance in developing a robust machine learning surrogate. A carefully selected feature set is crucial in improving model performance, reducing chances of overfitting, and enhancing the interpretability of the model. By selecting a subset of informative features, computational resources are conserved, and model complexity is reduced, leading to more efficient and effective machine learning models. Two primary approaches for feature identification are feature extraction and feature selection. Feature extraction involves manipulating available input data through various operators to derive a set of features in alternative parameter spaces. Techniques such as Principal Component Analysis (PCA) [175], Linear Discriminant Analysis (LDA) [176], and Autoencoders [177] are commonly employed in feature extraction. For features identified through feature extraction methods that rely on specific operations or calculations involving input data from all channels, it is necessary to record data from all input channels after model deployment. Conversely, feature selection involves ranking available input features based on a scoring criterion within the same coordinate space. Based on the procedure used for feature ranking, these methods are classified into filter-type and wrapper-type methods. Filter-type methods utilise statistical parameters like correlation coefficient, covariance, and chi-square statistics to evaluate the relevance of each input feature with reference to the target variable. The input features are ranked based on their relevance. Conversely, wrapper-type methods integrate the machine learning model into the selection process, training it on various feature subsets to identify the subset yielding the best performance. While wrapper-type methods often yield more accurate results, they are computationally intensive due to the need to train the model for each feature subset. In this application, a hybrid approach combining filter and wrapper methods is employed to achieve optimal parameter combinations. Here, a mutual information-based minimum redundancy and maximum relevance (MRMR) algorithm is used to rank the features based on their relevance. Based on a thresholding criterion, the first few features are selected. A wrapper-type approach is implemented on these selected features to arrive at the feature subset yielding the best performance. Typically, when dealing with a large number of features, the recursive feature addition (RFA) method is often employed. This approach continually adds features to the model until their inclusion no longer enhances performance. Similarly, the recursive feature elimination (RFE) method starts with a full set of features and progressively removes them until dropping a

feature significantly affects prediction accuracy. However, both methods assume linear dependence and prioritise individual feature importance, overlooking feature interactions. While these methods have delivered good results in many applications, they may yield suboptimal results for non-linear systems where parameter combinations heavily influence system response. An exhaustive feature combination is performed to analyse the impact of interactions amongst input features on prediction accuracy. A comparative analysis of the performance and computational costs of RFA and an exhaustive parameter search is conducted. The hybrid approach of combining filter and wrapper methods leverages the strengths of both methods, exploiting statistical relationships between features while also considering the performance of the machine learning model. By balancing computational efficiency with model performance, the hybrid approach enables the identification of optimal feature combinations at a reasonable computational cost. The details of this method are presented in the next section.

5.5.1 MRMR Algorithm

MRMR algorithm selects a subset of input features so that each selected feature provides unique information about the target variable. This method is particularly useful in cases where a large set of correlated input features is available. In this work, mutual information is used as the criterion to quantify the relevance of an input feature with respect to the target variable. Mutual information is based on the concept of entropy, which measures the uncertainty associated with a random variable. The entropy of a discrete random variable X with probability mass function $p(X)$ is given by Eq. 5.8.

$$H(X) = - \sum_x p(x) \log p(x) \quad (5.8)$$

Here X is used to represent random variable and x is used to represent its realisation. Similarly, conditional entropy is a measure of average uncertainty present in a variable given the outcome of another random variable. Conditional entropy is defined in Eq. 5.9.

$$H(X|Y) = - \sum_{x,y} p(x,y) \log p(x|y) \quad (5.9)$$

where, $p(x,y)$ is a joint distribution of random variables x and y , and $p(x|y)$ is a conditional distribution of x given y . Using these concepts, mutual information

quantifies how much information one random variable gives about another. The formal definition of mutual information between two random variables X and Y is defined in Eq. 5.10.

$$I(X, Y) = \sum_{x, y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \quad (5.10)$$

Mutual information can be expressed in terms of entropy as presented in Eq. 5.11.

$$I(X, Y) = H(X) - H(X|Y) \quad (5.11)$$

In the context of feature selection, assuming X represents an input feature and Y represents the target variable, the MRMR algorithm's aim is to identify a subset of input features X_s which maximises the mutual information between X_s and Y and minimises the mutual information between the components of X_s . Mathematically, the MRMR algorithm is defined in Eq. 5.12.

$$\max X_s = \left\{ \frac{1}{|X_s|} \sum_{i=1}^{|X_s|} I(X_{s_i}, Y) - \frac{1}{|X_s|(|X_s| - 1)} \sum_{i=1}^{|X_s|} \sum_{j=1, j \neq i}^{|X_s|} I(X_{s_i}, X_{s_j}) \right\} \quad (5.12)$$

where $|X_s|$ represents the number of features selected in subset X_s , $I(X_{s_i}, Y)$ represents the mutual information between each individual feature X_{s_i} and target variable Y and $I(X_{s_i}, X_{s_j})$ represents the mutual information between pairs of selected features in subset X_s , with $i \neq j$ to avoid calculating mutual information of a feature with itself. The MRMR algorithm iteratively selects features based on their relevance to the target variable and their redundancy with respect to other selected features, ultimately resulting in a subset of input features that maximise predictive performance while minimising redundancy.

For implementation in the current application, *fsmrmr* feature in Matlab[®] is used. In this algorithm in Matlab[®], each feature is assigned a score using the mutual information quotient (MIQ) presented in Eq. 5.13, such that the feature with a higher score corresponds to greater relevance with the target parameter.

$$MIQ_x = \frac{I(x, y)}{\frac{1}{|x|} \sum_{z \in S} I(x, z)} \quad (5.13)$$

In this algorithm, a reduction in the feature score represents improved confidence

in the feature selection. Using this property, the difference in feature scores is computed, and the number of features with a threshold of reduction in feature score is selected for the further model training process. Using this approach, a small subset ($|X_s|$) of features is identified. However, these features are selected on the basis of statistical properties and may not lead to the optimum performance. To ensure optimum performance, the subset of the feature is evaluated using the wrapper-type method to select the final optimum subset most suitable for the specific machine learning algorithm considered for the application. Two possible wrapper-type methods based on the combination of subsets are presented in the next section.

5.5.2 Recursive Feature Selection

The concept of the recursive feature selection method is briefly discussed in Section 5.1. These iterative techniques operate by systematically exploring the feature space to identify an optimal subset of input features that maximise the predictive performance of the machine learning model. The starting point can be either an empty set or the complete set of features, and the selection process operates through sequential addition (forward selection) or removal (backward elimination) of features. Forward selection, also known as recursive feature addition (RFA), starts with an empty feature set (denoted as X_s) and iteratively incorporates features based on their individual relevance to the target variable (Y). This relevance is assessed by evaluating the impact of each feature on a chosen performance metric, such as the normalised root mean squared error (NRMSE) employed in this study. The feature leading to the lowest NRMSE is then added to X_s . This process continues by evaluating all remaining features in combination with the chosen ones, exploring the combination that further enhances the model performance. These steps are repeated iteratively until a pre-defined stopping criterion is met, such as reaching a desired number of features or achieving a satisfactory level of model performance. The method terminates when adding extra features to X_s does not improve the performance. Conversely, backward elimination, also known as recursive feature elimination (RFE), starts with the complete set of features (X_s) and iteratively removes the least informative features such that the removal does not significantly compromise model performance. RFE identifies and removes the least impactful members, ultimately retaining the core group of features that delivers optimal results. By formalising these concepts, the forward selection process can be expressed as $X_s = X_{max}$, where X_s denotes the evolving feature set, and X_{max} sym-

bolises the set of most impactful features. In contrast, RFE’s iterative refinement is captured by $X_s = X_s - X_{min}$, signifying the removal of the least informative feature set (X_{min}) from the complete set X_s . While RFS methods are effective in identifying an optimal subset leading to a robust feature identification, a major limitation of these models is the computational efforts required. Using RFA the number of simulations, N_{sim} , required to select n number of features from N available features is given by Eq. 5.14

$$N_{sim} = \frac{n(2N - n + 1)}{2} \quad (5.14)$$

The computational cost increases with an increase in available features ‘ N ’ and the final required number of features ‘ n ’. In this context, pairing this method with the MRMR algorithm and operating on a subset X_s , such that $|X_s| \ll |X|$ can significantly bring down the computational cost by reducing the total number of features ‘ N ’ from $|X|$ to $|X_s|$. One of the major disadvantages of this method is that it focuses on the importance of individual features and does not take into account feature interaction, potentially leading to sub-optimal solutions. This limitation can lead to a sub-optimal performance for non-linear systems where parameter interactions can significantly affect the system dynamics. To further investigate the impact of neglecting feature interactions, the performance of RFA on a reduced feature set (X_s) is compared with an exhaustive evaluation of all possible feature combinations within X_s . This will lead a total number of simulations given by Eq. 5.15

$$N_{sim} = \sum_{k=1}^n C(N, k); \quad C(N, k) = \frac{N!}{k!(N - k)!} \quad (5.15)$$

While exhaustive evaluation could be computationally intensive for larger datasets, pairing this method with MRMR leads to a reduction in total number of simulations required. Further, this approach enables a direct comparison and highlights the potential performance loss due to neglecting interactions while remaining computationally feasible. The findings of the current investigation demonstrate the broader applicability of RFA to non-linear systems.

5.6 Hyperparameter Tuning and Model Training

LSTMs offer exceptional flexibility and power in sequential data modelling, but their true potential hinges on the careful selection of architectural settings known as hy-

perparameters. Hyperparameters are key architectural parameters of a machine learning model that are defined prior to the model training, directly impacting the model performance and generalisation ability. Effective hyperparameter tuning is, therefore, paramount for realising an efficient and well-performing LSTM model. The hyperparameter tuning process aims to identify the optimal combination of hyperparameter values, potentially leading to increased accuracy, faster convergence, and enhanced generalisation to unseen data. In this work, the *Bayesian optimisation* technique implemented in Matlab[®] is used to optimise the hyperparameters. This approach constructs a probabilistic surrogate model of the objective function, guiding informed decisions about generating new parameter combinations within a predefined range. The optimisation process employs an acquisition function 'expected improvement plus', which strategically balances exploring new regions with exploiting promising regions identified so far.

Before hyperparameter tuning, a two-layer LSTM architecture was established through preliminary trials. This configuration was chosen based on its initial performance and suitability for capturing temporal dependencies within the data. For the implemented model architecture, the number of hidden units in each LSTM layer, dropout probability, and batch size are considered in the hyperparameter tuning process. The number of hidden units in each LSTM layer directly governs the model learning performance and its capability to capture the relationship between input and output parameters. The dropout probability refers to the percentage of data that would be randomly set to zero; this parameter governs the generalisation capability of the model and helps in avoiding overfitting. Batch size, controlling the data fed per training iteration, presents a trade-off between frequent updates, faster training with less memory (smaller batches), and smoother gradient updates with higher memory requirements (larger batches). Through the hyperparameter tuning process, an optimum combination of these parameters leading to the best performance (minimum NRMSE) is chosen.

Having established the optimal architecture, the LSTM model is trained using the Adaptive Moment Estimation (Adam) algorithm [151]. This efficient optimiser tackles the limitations of Stochastic Gradient Descent by incorporating decaying averages of past gradients, enabling it to escape local minima and converge towards the global optimum. To further fine-tune the learning process, a plateauing validation loss learning rate schedule is implemented. In this approach, the

learning rate was reduced by a factor of δ if the validation loss remained stagnant for n_p consecutive epochs, ensuring progress towards the minimum. In addition, a minimum learning rate was established to prevent stagnation and overly prolonged training. In this approach, after m reductions in the learning rate, if the performance plateaued again, the learning rate increases by a factor of ξ . This step ensures that the model does not get stuck in saddle points, which are plateaus in the loss landscape that are not the global minimum. This strategy, based on empirical observations, often produces superior results compared to a constant learning rate approach. In addition, during training, especially in deep networks, gradients can become very small, effectively slowing down the learning process. A controlled increase in learning rate can counteract this issue and accelerate progress, particularly when approaching the minimum. The performance of the presented learning rate schedule strategy depends on the chosen values for m , n_p , ξ , and δ for the current application the considered values of 5, 10, 1.5 and 0.9 were found to deliver good results. These values should be selected on the basis of the nature of the problem and the training data. The complete process of developing an LSTM surrogate for sequence-to-sequence regression is presented in Algorithm 1.

5.7 Numerical Results

Using the numerical model in combination with the training data generation approach, the numerical model was used to evaluate the wind turbine response to the turbulent wind inflow generated using Turbsim, NREL's turbulence modelling software. The time history of aerodynamic forces acting at each node of the wind turbine blade and the blade tip deformations were recorded. Aerodynamic forces were used as input features and the blade response represented the target variable.

Following the details presented in Algorithm 1, the first step was to perform the data pre-processing. Since, in this case, the output is generated from a numerical model, extensive pre-processing or outlier detection is not required. The only pre-processing carried out in this study is the downsampling of data. The numerical model uses a time step of 0.005 seconds to achieve convergence; however, such fine resolution may not be required in the development of the surrogate model. The downsampling frequency should be chosen so that the dynamics of the system is sufficiently captured without omitting any significant frequency. In the blade deformation response, the lower frequency response is contributed by the

Algorithm 1 Machine Learning Model Development Algorithm

Input: Raw data (X, y) **Output:** Optimised machine learning model

```

1: 1. Data Preprocessing
2: Outlier Removal: Apply a suitable outlier detection and removal technique
3: Data Downsampling: Use an appropriate downsampling technique to reduce the quantity of
   the data
4: 2. Feature Selection
5: Normalisation: Normalise features so that feature ranking is not impacted by the absolute
   value of the data
6: 2.1 MRMR Feature Ranking:
7: for  $X_i \in X$  do
8:   Calculate mutual information  $I(X_i, y)$  between  $X_i$  and  $y$  using Eq. 5.10
9:   for  $X_j \in X$  do
10:    if  $X_i \neq X_j$  then
11:      Calculate mutual information  $I(X_i, X_j)$  between  $X_i$  and  $X_j$  using Eq. 5.10
12:    end if
13:  end for
14:  Calculate MIQ score  $MIQ_x$  for  $X_i$  using Eq. 5.13
15: end for state Select a subset  $X_s$  of features based on the highest MIQ scores, accounting for
   a drop in scores for confidence
16: 2.2 Initial Hyperparameter Tuning state Define the LSTM model architecture
17: Select hyperparameters to tune
18: Perform initial hyperparameter tuning on  $X_s$  using bayesian optimisation
19: 2.3 Wrapper-based Feature Selection:
20: Recursive Feature Addition (RFA):
21:  $X_s^{temp} = \emptyset$ 
22: while improvement criteria met do
23:   for  $X_i \in X_s$  do
24:     Evaluate model performance with  $X_s$ 
25:   end for
26:   Select the  $X_i$  that yields the best performance
27:   Add the selected  $X_i$  to  $X_s^{RFA}$ 
28: end while
29: Exhaustive Parameter Search:
30: for all possible subsets of  $X_s$  do
31:   Evaluate model performance with the subset
32: end for
33: Select the subset  $X_s^{EPS}$  that yields the best performance
34: Choose the final refined subset:  $X_s^* = \text{argmax}(\text{performance}(X_s^{RFA}), \text{performance}(X_s^{EPS}))$ 
35: 3. Hyperparameter Tuning (optional)
36: Further refine hyperparameter tuning on the selected subset  $X_s^*$ 
37: 4. Model Training state Train the model using the selected hyperparameters and Adam
   optimiser with a proposed learning rate scheduling approach
38: 5. Model Evaluation state Evaluate the performance of the trained model on unseen data
   using appropriate metrics
Ensure: Optimised machine learning model is ready for deployment

```




Figure 5.2: Graphical representation of computing blade deformations from turbulent inflow using surrogate modelling approach

blade rotations and the low-frequency component of the wind. According to the operating parameters of the 15MW wind turbine, the cut-in rotor speed is 3 rpm, which corresponds to $\frac{1}{20}$ Hz. The frequency of the downsampled data should be greater than the lowest significant frequency to sufficiently capture the response at this frequency. From the perspective of model training and memory management, downsampling reduces the total amount of data, leading to faster training times and fewer memory requirements. The downsampling frequency also depends on the model deployment conditions. For these models, the data should be fed at the same frequency as the data considered during training. Assuming a hub-mounted Light Detection and Ranging (LIDAR) system to measure the incoming wind speed, the frequency of data used during training should match the frequency of data measurement. With reference to NREL's research on LIDAR wind speed measurements of evolving wind fields [178], the LIDAR sampling frequency is assumed to be 50 Hz. The data from the numerical model are sampled at 50 Hz.

The model proposed in this chapter can be applied to an operational wind turbine equipped with a LIDAR system. For such application, the surrogate model presented in Chapter 4 can be used to estimate aerodynamic forces using wind speed measurements. These aerodynamic load predictions can then be integrated into the LSTM model presented here to predict blade tip deformations. By combining these two surrogate models, a quick estimate of blade tip deformations can be obtained using the LIDAR measurements. A visual representation of this process is presented in Fig. 5.2.




Figure 5.3: Feature score prediction using MRMR algorithm

Assuming that the information on aerodynamic forces at each blade is known, the aim is to identify the subset of this data that can accurately predict blade deformations. For the 15MW wind turbine blade, a total of 50 such nodes are present, and at each node, the aerodynamic forces normal to the rotor plane and tangential to the rotor plane are recorded. For the purpose of surrogate model development, the 1st and the last blade node are not considered. Therefore, the complete set of aerodynamic forces considered for this investigation has a total of 96 input features. These features have been analysed using the MRMR algorithm to compute the MIQ quotient using Eq. 5.13. The results of feature scores predicted by the algorithm are presented in Fig. 5.3.

Using the drop in the predicted feature score as a criterion, the MRMR algorithm presents ten features that represent the subset X_s defined in Algorithm 1. The relevant nodes chosen by the MRMR algorithm are well distributed along the blade node, as shown in Fig. 5.3. The forces at these nodes can capture a unique component of the blade deformation response, as the nodes placed close to each other will exhibit a high correlation, leading to less unique information. The correlation matrix for the raw data and the data picked by the MRMR algorithm is presented in Fig. 5.4 and Fig. 5.5, respectively. These figures show that a smaller number of features with low correlation were chosen through the MRMR algorithm, which


Hyperparameters	Lower bound	Upper bound
Number of hidden units in LSTM	50	200
Dropout probability	0.1	0.6
Batch size	1	15

Table 5.2: Hyperparameter bounds considered for tuning


implies that these features are likely to hold distinct information about the response.

Hyperparameter tuning was performed to identify the best combination of hyperparameters using the features identified using the MRMR algorithm. The *bayesian optimisation* algorithm in Matlab® is used for this tuning process. The limits of the values for this optimisation are chosen based on the commonly adopted values in the literature and are presented in Table 5.2.

In the next step, using the set of optimised parameters, the model is trained using the RFA method. In all model training exercises for the below rated wind speeds, the azimuthal position of the blade and the rotor speed were kept constant characteristics of a subset X_s , while for the above rated operating conditions, the azimuthal position of the blade and the pitching angle are constant characteristics. For the RFA approach, the average NRMSE value across all wind speeds is considered as the performance metric. This algorithm continues to add features until the addition of features leads to a reduction in NRMSE. The general process for this is presented in Algorithm 1. The performance of the model developed using the features identified using the MRMR algorithm is presented in Fig. 5.6 and Fig. 5.7. Fig. 5.6 shows a qualitative representation of model performance for different subsets of the input feature set identified by the MRMR algorithm, where the size of the subset is linearly increased by adding the features according to their feature scores such that the feature with the highest score is added first and then the next feature with the second highest score and so on. However, increasing the number of input features may not always lead to improved performance; this fact is highlighted in Fig. 5.7, where the achieved NRMSE is plotted against the number of input features considered. From Fig. 5.7, it is evident that the best model performance is achieved when a total of five aerodynamic input features are considered. The results for a model trained on these parameters are presented in Fig. 5.8 and Fig. 5.9, while the histogram plot of model error is presented in Fig. 5.10 and 5.11



(a) Correlation matrix for forces normal to rotor plane



(b) Correlation matrix for forces tangential to rotor plane

Figure 5.4: Correlation matrix showing high correlation in feature set




Figure 5.5: Correlation matrix of features identified using MRMR algorithm

However, to ensure that an optimum subset of the input features is identified, wrapper-based feature selection methods are required. In wrapper-based selection, the RFA method is performed first on the subset of input features identified using the MRMR algorithm. The performance of the final model developed using this method is presented in Fig. 5.12 and Fig. 5.13.

To validate the consistency and generalisability of the trained model, an error analysis was performed through histogram analysis to ensure a comprehensive assessment of the model performance. Fig. 5.14 and 5.15 demonstrate that the error residuals follow a zero mean Gaussian distribution, suggesting that the error behaves as white noise without any noticeable heteroskedasticity or systematic patterns. This indicates the robust performance of the model across different inputs. Furthermore, the coefficient of variation (CV) for all trained models was calculated. CV was selected as the preferred metric for evaluating the accuracy of the model because of its ability to standardise the dispersion of errors relative to the mean prediction. Unlike absolute measures such as the standard deviation, which may not be directly comparable across models with different magnitudes of predictions, the CV normalises the variability by the mean value. This makes it an ideal metric to assess the relative consistency of model predictions, regardless of the scale of the




Figure 5.6: Performance of LSTM Model trained using features identified through MRMR algorithm

problem.

Following the RFA approach, the best model performance was observed considering the following input features:

- Normal force at Node 31
- Tangential force at Node 47
- Rotor speed
- Blade azimuth angle

Fig. 5.12 shows that the model is able to capture the accurate trend of the blade tip deformation using the information from the normal force at node 31. The results are further improved when the tangential force at node 47 is considered. The effect of the addition of more channels in terms of the average NRMSE value is presented in Table 5.3. Quantitative analysis of the loss values shows that the addition of information from node 47 results in a 19.11% reduction in the NRMSE. Since the model developed using information from only one node is able to capture the trend well, the need for the addition of more nodes depends on the expected performance of the model and the resources available at the installation location.




Figure 5.7: Effect of increasing input feature space on model performance




Figure 5.8: Best performance achieved from LSTM Model trained using features identified through MRMR algorithm across different stochastic realisations

Number of input features	NRMSE
1	0.081
2	0.068

Table 5.3: Effect of number of input force channels on NRMSE




Figure 5.9: Best performance achieved from LSTM Model trained using features identified through MRMR algorithm across different mean wind speeds




Figure 5.10: Error analysis of the LSTM Model trained using features identified through MRMR algorithm across different stochastic realisations





Figure 5.11: Error analysis of the LSTM Model trained using features identified through MRMR algorithm across different mean wind speeds


To study the effect of feature interaction, the model is trained using the exhaustive parameter search method. The results of the exhaustive parameter search are presented in Fig. 5.16 and Fig. 5.17, whereas their error analysis is presented in Fig. 5.18 and 5.19. The best performance for the exhaustive feature search method is observed for the following combination of input features:

- Normal force at Node 30
- Tangential force at Node 45
- Rotor speed
- Blade azimuth angle

To compare the performance of the three feature selection methods, the final NRMSE and CV obtained for the test data set using the LSTM model trained on the best feature subset identified by these methods are presented in Table 5.4. It can be seen that a CV value of less than 10% is achieved for trained LSTM models, indicating low relative variability, reinforcing the high precision and stability of




(a) Predictions of LSTM model trained using information from node 31




(b) Predictions of LSTM model trained using information from nodes 31 and 47

Figure 5.12: Performance of LSTM Model across different wind speeds using the features identified through RFA approach




(a) Predictions of LSTM model trained using information from node 31




(b) Predictions of LSTM model trained using information from nodes 31 and 47

Figure 5.13: Performance of LSTM Model across different stochastic realisations using the features identified through RFA approach




(a) Error analysis of LSTM model predictions trained using information from node 31




(b) Error analysis of LSTM model predictions trained using information from nodes 31 and 47

Figure 5.14: Histogram plot of LSTM model prediction error across different wind speeds trained using the features identified through RFA approach



(a) Error analysis of LSTM model predictions trained using information from nodes 31



(b) Error analysis of LSTM model predictions trained using information from nodes 31 and 47

Figure 5.15: Histogram plot of LSTM model prediction error across different stochastic realisations using the features identified through RFA approach




Figure 5.16: Performance of LSTM Model trained using Exhaustive Feature Search across different mean wind speeds




Figure 5.17: Performance of LSTM Model trained using Exhaustive Feature Search across different stochastic realisations




Figure 5.18: Error analysis of LSTM Model trained using Exhaustive Feature Search across different mean wind speeds




Figure 5.19: Error of LSTM Model trained using Exhaustive Feature Search across different stochastic realisations

Feature selection method	NRMSE	CV (%)
MRMR	0.114	7.026
MRMR + RFA	0.075	5.129
MRMR + Exhaustive Feature Identification	0.053	4.635

Table 5.4: Feature selection method and obtained performance

the model in making predictions. This measure therefore offers a clear and reliable indication of the accuracy of the model, confirming that the trained model consistently produces accurate and stable results across the dataset. The low CV value combined with the small magnitude of NRMSE highlights the precision and accuracy of the LSTM models developed. It can be seen that the exhaustive parameter search method achieves better results, but at the cost of increasing computational resources.

The loss curve depicted in Fig. 5.20 illustrates the progression of the loss values computed for both the training and validation data sets. As iterations increase, both loss values decrease, demonstrating the model's convergence. However, the decrease in loss slows with more iterations. The learning rate scheduling strategy suggested in this study monitors the average validation loss trend to adjust the learning rate, thereby enhancing convergence. This is presented in the graph that shows the variations in the average validation loss, highlighting the effectiveness of the proposed method.

5.7.1 Dealing with Noisy Input Signals

One of the possible applications of the LSTM model developed in this research is a virtual sensor to analyse the blade response in operating wind turbines. Due to its ability to handle sequential data, it is possible that this model could work online in conjunction with wind turbine controllers or be used for offline applications. In such scenarios, where the model works with real-world measurements, the model inputs will often be subject to uncertainty due to measurement noise or sensitivity limitations. To assess the model's robustness against this uncertainty, synthetic Gaussian noise is introduced into the simulated wind velocity data.

The noise levels chosen (10%, 20%, and 30%) represent scenarios with a high uncertainty of the measurement. Although generic, these values are chosen to be




Figure 5.20: Sample loss curve demonstrating trends in training and validation loss

Noise level	Achieved NRMSE	% Change in mean	% Change in maxima
10%	0.136	0.07	0.65
20%	0.211	0.94	2.63
30%	0.297	3.53	5.54

Table 5.5: Effect of noise on LSTM model performance

consistent with typical error ranges considered in the literature [170]. However, a more accurate representation would require accounting for the propagation of instrument-specific uncertainties in wind speed measurements through aerodynamic force calculations. This detailed uncertainty analysis is beyond the scope of this study, but remains an important consideration for real-world deployment.

The model's performance is evaluated under these noise levels by comparing its predictions with the noise-free ground truth. Fig. 5.21 visually shows the impact of noise on the predictions at a mean wind speed of 8 m/s, and Table 5.5 summarises the corresponding NRMSE and the percentage change observed in the mean and maximum value of the predictions with noisy data and the corresponding values with the zero noise case.

To ensure that the comparison of model performance provides statistically significant results, ten independent noise realisations are simulated for each noise




Figure 5.21: LSTM model predictions for noisy input data

level, and the average NRMSE, along with the percentage change in mean and maximum values, is presented in Table 5.5. The results demonstrate a good overall robustness to uncertainty in the input signal. The acceptable level of uncertainty in the input signals depends on the specific application domain. For applications such as fatigue analysis, where results rely primarily on aggregate statistics, comparatively higher noise levels could be considered acceptable. However, applications focused on extreme values, such as structural health monitoring, might be significantly affected by high noise levels. A further detailed analysis can be performed based on the expected application of the model.

5.7.2 Analysis and Interpretation


The results presented in this chapter highlight several key findings on feature selection, model performance, and noise handling capabilities. The filter type feature selection method, the MRMR algorithm in particular, provides an efficient approach to identify the most informative feature subset based on statistical properties of the data, drastically reducing the input feature size from 96 to a mere 10 feature subset (Fig. 5.3 and Fig. 5.5). The surrogate model trained on this subset further highlights that increasing the number of input features may not lead to improved performance (Fig. 5.7), highlighting the need to further refine the set of input features to extract

pertinent information while maintaining efficiency. Further refinement through RFA led to better results with a smaller number of input features. Although exhaustive parameter search yielded the best performance, it comes at the cost of increased computational burden. Therefore, the best choice between these methods is based on specific application requirements and resource constraints. This trade-off is crucial in practical applications where resource limitations cannot be ignored.


Analysing the chosen features' locations on the blade offered valuable insights. The locations of the features identified by these feature selection methods on the blade are represented in Fig. 5.22. This localisation of features through wrapper-type methods reveals the blade locations that capture most of the information about the blade tip deformation. Further analysis of these locations through a structural engineering perspective shows that the locations identified through the feature selection approach represent the areas with lower stiffness and higher deformation levels. Naturally, these regions will contribute substantially to the blade tip deformations. These regions and the corresponding mode shape and rigidity are presented in Fig. 5.23, where the areas are shaded using a rectangular patch.

The performance of the surrogate model has been demonstrated by presenting the predictive capacity of the model for a range of different mean wind speed values (Fig. 5.8, Fig. 5.12, and Fig. 5.16). The accuracy of these results shows that the model can accurately predict the blade response over a range of wind velocities that can be encountered during the operation of the wind turbine. Furthermore, since the stochastic nature of the inflow wind is the most significant source of uncertainty encountered during wind turbine operation, the performance of the model in different stochastic realisations further highlights the efficiency of the surrogate in handling the turbulent inflow encountered during the operation (Fig. 5.9, Fig. 5.13, and Fig. 5.17). The results presented comprehensively demonstrate the capability of this model to accurately predict the blade response at a negligible computational cost.


One of the main applications of this model could be as a virtual sensor to detect the response of the blade due to measured inflow conditions. Since this application will inherently require a model that can produce good results under noisy measurements, the performance of a model under noisy input is investigated (Fig. 5.21). The model exhibited good robustness against simulated Gaussian noise, even at levels as high as 30%. Interestingly, noise had a more pronounced effect on



(a) Nodes identified through exhaustive parameter search: 30 and 45




(b) Nodes identified through recursive feature addition: 31 and 47




(c) Nodes identified through MRMR: 2, 3, 31, 45 and 47

Figure 5.22: Nodes identified through different feature identification approaches



(a) Region identified by wrapper type method for aerodynamic force normal to plane



(b) Region identified by wrapper method for force tangential to rotor plane

Figure 5.23: Region of interest on blade span identified by wrapper type method

the extreme values compared to the mean responses. This behaviour is expected because fluctuations in extreme values are naturally more susceptible to distortion by noise. However, the acceptable noise level depends on the intended application. For example, in fatigue analysis, primarily concerned with aggregate statistics, higher noise levels can be tolerated than in applications like structural health monitoring, which rely heavily on precise extreme value estimations. In addition, deeper insight into the specific application can inform a more detailed analysis of noise tolerance. While this investigation considers robustness using simulated noise, it is crucial to acknowledge the limitations of this approach. Real-world noise characteristics can differ from assumed distributions and exhibit dynamic variations. Future work should explore the incorporation of uncertainty quantification techniques and calibration methods to bridge the gap between simulated and real-world scenarios. Additionally, investigating performance with actual sensor data from wind turbines will provide a more comprehensive understanding of the model's robustness in practical applications.

5.7.3 Computational advantage of surrogate approach

One of the main aims of surrogate models is to provide a computationally efficient alternative to MBD models to predict the dynamic response of wind turbines. Building on the high accuracy achieved through these models, a quantitative analysis of their computational advantage is presented in this section. While isolating the exact computational cost of the numerical integration scheme for deformation calculation implemented in MBD is challenging due to the multiple sub-modules implemented in this code, valuable insights can be gained through the comparison of overall execution times. The LSTM-based surrogate model can predict the blade deformation response to a 600-second random aerodynamic load in a tenth of a second. In contrast, simulating the same scenario using a complete dynamic analysis of the MBD model takes approximately 7 minutes. This comparison underscores the substantial efficiency gains achieved through the developed surrogate models. To ensure consistency, all simulations were performed on a system with an 8 core Intel Xeon CPU, 3.8GHz clock speed, 32GB RAM, and Microsoft Windows 10 Pro.

5.8 Discussion

In this chapter, the application of LSTM to predict the blade tip deformation response time series has been demonstrated. A detailed algorithm to develop a machine learning surrogate for sequence prediction addressing the most important degrees of freedom available was presented, and a comprehensive analysis of the model's accuracy and robustness against measurement uncertainty was demonstrated. The results presented in this chapter can be summarised as follows:

- Following the methodology presented in this chapter, accurate LSTM surrogate models can be developed for multivariate time series regression applications.
- The comparative analysis of feature selection methods provides a comprehensive analysis to identify the important features based on the required model performance and the available computational resources.
- The learning rate scheduling approach presented in this chapter leads to optimum parameter convergence.
- The application of the surrogate for wind turbine monitoring and sensing applications can reduce the number of sensors required and can produce satisfactory performance.

Chapter 6

Multi-task learning Long Short-Term Memory model to emulate blade dynamics

6.1 Motivation

The inherent computational cost of wind turbine response simulations, compounded by the highly stochastic spatio-temporal wind field, often hinders applications such as site-specific performance analysis, fatigue analysis, and reliability estimation. The models proposed in Chapter 4 and Chapter 5 can replace sub-modules for aerodynamic force computation and structural response estimation in MBD models, respectively. These models can also be valuable during the operational phase of a wind turbine by acting as virtual sensors. However, their application in a simulation environment adds additional complexities as these models require wind speed data at a specific point on a deforming and rotating blade. Generating wind speed data at points that evolve in both spatial and temporal dimensions introduces difficulties in their application in a simulation environment. The additional computational cost of such implementation will be compounded while performing large sets of simulations. To address this issue, this chapter proposes utilising the spatio-temporal wind field as the input to an LSTM model for predicting the blade response. This model design simplifies its implementation in the simulation environment and could potentially serve as a surrogate for the MBD model presented in Chapter 3, stream-

lining computationally intensive wind turbine analysis. The major challenge in such an implementation is introduced by the high dimensionality of the turbulent wind speed inflow due to the spatio-temporal evolution of wind speed. The nature of wind speed inflow and the possible ways to handle this high-dimensional input to arrive at an efficient LSTM model configuration are discussed in this chapter.

6.2 High dimensionality of the stochastic wind inflow

The dynamic analysis of wind turbines requires assessing how the structure responds to the wind conditions it encounters during operation. Given the increasing dimensions of wind turbine rotors, wind speeds can fluctuate significantly across the rotor plane, resulting in varying wind conditions experienced at various points within the rotor area. This effect will be more pronounced in turbines with larger rotor diameters. Furthermore, since wind speeds change over time, these spatial points within the rotor area are subjected to temporal variations in wind speed. Therefore, accurate prediction of a wind turbine's response to dynamic forces necessitates a realistic simulation of wind speed variations across its large rotor area and their evolution over time. The complexity of the spatio-temporal wind field introduces a fundamental challenge in wind turbine modelling in terms of a high dimensionality of the input space. To accurately estimate the wind turbine response, large sets of simulations must be generated during the design phase to cover the various possible combinations of wind flow patterns. TurbSim [140], a widely used turbulent wind field simulator, is a key tool for this purpose, but it is impacted by this dimensionality issue. TurbSim serves as a stochastic simulator for turbulent winds, generating full-field representations of wind dynamics. It operates on a statistical model to produce time series data comprising three-dimensional wind speed vectors across a fixed two-dimensional vertical grid. This grid remains stationary in space throughout the simulation process. To ensure relevance to specific wind turbine configurations and site conditions, TurbSim provides customisable parameters that include turbine geometry and meteorological factors that define the wind speed field. For detailed information on these parameters, the reader is referred to the TurbSim user guide [140]. Here, a summary of key variables crucial to this study is presented, with their definitions and significance. These parameters control the size and complexity of the wind datasets:

1. Random Seeds: TurbSim uses two random seeds (*RandSeed1* and *RandSeed2*) to create random phases, one per frequency per grid point per wind component, for velocity time series and for random number generation schemes. These seeds primarily ensure the reproducibility of the synthetic wind fields and generate unique stochastic realisations for specified environmental conditions. In addition, these seeds can be varied at constant meteorological parameters to produce various stochastic patterns with similar statistical representations.
2. Turbine/Model Specification: In the configuration of TurbSim, the parameters *GridHeight* and *GridWidth* define the total vertical expanse within which the wind turbine rotor operates. This region is then subdivided into segments defined using *NumGridY* and *NumGridZ*, establishing the spatial resolution of the computational domain. *HubHt*, another variable, designates the hub height of the wind turbine, serving as a reference point for grid placement. Horizontally, the rotor is assumed to be centrally aligned within the grid. Meanwhile, *HFlowAng* and *VFlowAng* govern the average horizontal and vertical angles of wind incidence relative to the vertical plane of the rotor. These angles remain constant throughout the grid. A representative image of the TurbSim grid that includes the wind turbine rotor is presented in Fig. 6.1




Figure 6.1: Visual representation of TurbSim grid

3. Meteorological Boundary Conditions: Within TurbSim, the variable *TurbModel* offers the flexibility to designate the spectral model used to generate the wind speed field. TurbSim offers multiple spectral models, including two IEC models (the IEC Kaimal model and the IEC Von Karman Isotropic

model), the Risø smooth-terrain model, and various NREL site-specific models. The choice of spectral model is guided by site characteristics and the type of analysis. In addition, users can specify turbulence types, such as normal and extreme turbulence, to suit their simulation requirements. The statistical properties characterising the inflow, notably the mean wind speed, the intensity of the turbulence, the length of the surface roughness, and the power law exponents, are defined through variables U_{ref} , IEC_{turb} , Z_0 , and PLE_{exp} , respectively. The mean wind speed is specifically defined at a designated reference height, denoted $RefHt$.

TurbSim employs spatial and temporal discretisation techniques to analyse the wind field comprehensively. It partitions the domain into grid cells, enabling the calculation of wind speeds and turbulence properties at each grid point. The temporal dimension of the spatial grid is subdivided into small intervals to accurately capture temporal fluctuations in wind speed and turbulence. Using the specified turbulence model and input parameters, TurbSim generates synthetic turbulent variations in wind speed along temporal and spatial coordinates across the domain. This process involves simulating the spatial and temporal evolution of the turbulence on the basis of statistical properties derived from the chosen turbulence model. The IEA-15MW reference wind turbine modelled in this thesis, characterised by a 240m rotor diameter and a 150m hub height, is modelled using a rectangular cross-section measuring 285m x 285m, featuring a 25 x 25 spatial grid layout. Considering these turbine specifications, a single 10-minute duration wind field, with a temporal discretisation of 0.05 seconds, yields a substantial 750,625 individual observations in the input space. However, the necessity for finer temporal discretisation and conducting numerous simulations to train a surrogate model further increases the number of data points forming a high-dimensional input space. Traditional machine learning algorithms may face challenges, such as overfitting and increased computational demands when dealing with such high-dimensional input data. To mitigate these issues, dimensionality reduction techniques are explored in the subsequent section, aiming to extract essential information from the input space while minimising the number of variables, facilitating streamlined and efficient model training and implementation.

6.3 Dimensionality Reduction Techniques for Wind Field Data

Dimensionality reduction refers to the process of transforming high-dimensional data into a meaningful representation with reduced dimensionality. Dimensionality reduction is often used as a pre-processing step before building surrogate models for high-dimensional input spaces. The reduced-dimensional representation facilitates faster model training and can improve the robustness of surrogate models.

Mathematically, let $\mathbf{X}$ denote a high-dimensional dataset, where each row corresponds to an observation and each column represents a feature. If p is the number of observations and N is the number of features, then $\mathbf{X}$ is an $p \times N$ matrix. If p is a large number, handling such high-dimensional data poses computational challenges and can lead to inefficiencies in analysis and modelling tasks. Dimensionality reduction techniques offer a mathematical framework to address these challenges by transforming the high-dimensional data into a lower-dimensional representation. Let $\mathbf{Z}$ represent the reduced-dimensional representation of the original dataset $\mathbf{X}$, where $\mathbf{Z}$ is an $p \times n$ matrix with, where $n \ll N$ being the reduced number of dimensions. The goal of dimensionality reduction is to find a mapping function $f : \mathbf{X} \rightarrow \mathbf{Z}$ that captures the essential information in the original data while reducing its dimensionality. This reduced representation of $\mathbf{X}$ should have a dimensionality that corresponds to the intrinsic dimensionality of the data, which captures the maximum information of the original dataset with a small number of intrinsic dimensions. Theoretically, the intrinsic dimension of data is the minimum number of parameters needed to account for the observed properties of the data. However, identifying the intrinsic dimension of high dimensional data is not a trivial task and becomes a parameter of the dimensionality reduction algorithm.

Dimensionality reduction has been an active area of research not just in the domain of surrogate modelling but also in applications related to processing speech signals [179], digital photographs [180], or medical imagery[181]. Specialised literature on various dimensionality reduction techniques and their comparative performance on standard datasets is presented by Van Der Maaten, Postma, Herik, et al. [180]. By identifying and retaining the most informative features, these techniques facilitate efficient analysis and modelling of wind turbine responses while reducing computational complexity and improving interpretability.

In the context of wind turbines, TurbSim simulations can generate wind fields consisting of thousands of data points across space and time. Dimensionality reduction is crucial to transform these complex wind fields into more manageable representations for efficient response prediction. Various dimensionality reduction techniques are explored in the literature for application to high dimensional problems, of all these techniques, Principal Component Analysis (PCA) is one of the most widely used approach. PCA is a statistical technique which focuses on capturing spatial correlations within the data by identifying principal components that capture the maximum variance in the dataset. PCA offers computational efficiency and linear mapping, making it suitable for handling large-scale wind datasets encountered in wind turbine modelling. Lataniotis [182] has shown although there is no “Jack of all trades” for dimensionality reduction, PCA consistently outperformed the other dimensionality reduction techniques implemented in their research in terms of the reconstruction error achieved, and the robustness of the results over different repetitions on the standard datasets used in dimensionality reduction problems. PCA implemented for wind speed forecasting [183, 184], wind turbine fault detection [185], and monitoring [186] has also delivered good results. Principal components can reveal meaningful patterns in the wind field in terms of recurring spatial and temporal trends in wind speed variations. The reduced-dimensional representation using PCA further helps in data handling and analysis. While PCA excels at capturing the variance within the data through linear relationships, it might overlook the presence of underlying non-linear patterns in complex wind fields. Also, while principal components represent significant variance, interpreting their physical meaning can be challenging for wind fields. To address these potential limitations, the Discrete Cosine Transform (DCT) has been chosen as a complementary technique.

The Discrete Cosine Transform (DCT) offers a powerful tool to analyse the dynamics of the wind field by decomposing the data into its underlying frequency components, revealing patterns of spatial and temporal variation. Unlike PCA, DCT emphasises the frequency-domain representation of signals or images, using its energy compaction property to highlight dominant frequency components. This data transformation allows DCT to capture the energy content of a signal or image in fewer coefficients by exploiting the characteristics of the data in the frequency domain. This property enables effective compression and representation of wind data, which can be exploited for dimensionality reduction by capturing dominant

spatial and temporal correlations inherent in wind speed fields. The components obtained through DCT (cosine functions with specific frequencies) correspond directly to spatial variations of different wavelengths or scales within the wind field data. This representation makes the DCT components physically interpretable. Recently, DCT has been applied by Schär, Marelli, and Sudret [187] for dimensionality reduction of the stochastic wind field and has been shown to deliver good results. Although DCT is well-suited for analysing frequency-domain characteristics, it implicitly assumes a degree of periodicity in the data. This might make it less effective in capturing localised, irregular, or non-periodic wind events.

Given the distinct strengths and limitations of both PCA and DCT for wind field representation, a comprehensive investigation is necessary to determine their suitability for specific analysis tasks. This thesis employs both techniques, enabling a thorough exploration of dimensionality reduction approaches and aiding in identifying the most effective method for the prediction of wind turbine response. This selection aligns with the distinct capabilities of these methods along the dimensionality reduction spectrum, providing an opportunity to leverage their complementary strengths. The mathematical formulation of these methods is presented below.

6.3.1 Principal Component Analysis

PCA is a mathematical tool designed to uncover the core patterns within complex, high-dimensional datasets. In the context of wind field analysis, it aims to identify a small number of dominant combinations of wind speed values across spatial locations and over time, which can capture the maximum information. These combinations, known as principal components, capture the most significant variations in the wind field while minimising the loss of information. Importantly, PCA transforms potentially correlated features into a smaller set of uncorrelated variables called principal components. This transformation maximises the variance explained by each component, thereby emphasising the most prominent patterns within the data. To achieve this, PCA normalises the data to have zero mean. This ensures that the principal components are focused on capturing variations or deviations from the average behaviour and not the absolute magnitude of the wind speeds. Following the original notations, assume $(\mathbf{X} \in \mathbb{R}^{p \times N})$ represent the original wind field dataset matrix, with ' p ' observations (time steps) and ' N ' features (spatial locations). The mean-centred dataset $(\mathbf{X}' \in \mathbb{R}^{p \times N})$ is obtained by subtracting the

mean wind speed at each spatial grid point from all its corresponding values across time:

$$\mathbf{X}' = \mathbf{X} - \bar{\mathbf{x}} \quad (6.1)$$

where $\bar{\mathbf{x}}$ is the mean vector of the dataset. PCA computes the covariance matrix to quantify the pairwise linear relationships between different features (i.e., wind speeds at different spatial locations) within the wind field data. The covariance matrix is a fundamental tool for identifying directions of maximum variance, which form the basis for the principal components. The covariance matrix (Σ) of the mean centred data is computed as:

$$\Sigma = \frac{1}{\mathbf{p} - 1} (\mathbf{X}'^T \mathbf{X}') \quad (6.2)$$

The covariance matrix Σ is a square matrix where each element (i, j) represents the covariance between the i -th and j -th wind speed features. A core step in PCA is the eigen-decomposition of the covariance matrix (Σ). Eigen-decomposition reveals the directions of maximum variance within the wind field data and quantifies the amount of variation captured along each direction. These directions are called eigenvectors ($\mathbf{V} \in \mathbb{R}^{N \times N}$ and their corresponding magnitudes are known as eigenvalues ($\Lambda \in \mathbb{R}^{1 \times N}$). The eigenvector matrix $\mathbf{V}$ with N columns and eigenvalue vector Λ with N entries can be expressed as:

$$\mathbf{V} = \begin{pmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \cdots & \mathbf{v}_N \end{pmatrix}$$

$$\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_N)$$

PCA sorts the eigenvectors in decreasing order of their corresponding eigenvalues. The eigenvector with the largest eigenvalue represents the principal component explaining the most variance in the wind field data. Subsequent eigenvectors, in decreasing order of eigenvalues, capture directions of progressively less significant variation, often orthogonal to the previously identified directions. A crucial step in PCA is determining the number of principal components to retain to represent the wind field data. Since the eigenvalues represent the variance explained by their corresponding eigenvectors, the most informative components can be prioritised by examining the cumulative explained variance. The cumulative explained variance is calculated as follows:

$$e_k = \frac{\sum_{i=1}^k \lambda_i}{\sum_{i=1}^N \lambda_i} \quad (6.3)$$

A common approach is to select the smallest number of principal components (k) that achieve a desired percentage of explained variance. Finally, the original high-dimensional dataset $\mathbf{X}$ is projected onto the new basis spanned by the selected principal components to obtain the reduced-dimensional representation $\mathbf{Z} \in \mathbb{R}^{p \times n}$, where n is the desired number of principal components to retain. The projection can be expressed as:

$$\mathbf{Z} = \mathbf{X}'\mathbf{V}_n$$

where $\mathbf{V}_n$ is the matrix containing the first n eigenvectors corresponding to the largest n eigenvalues. This projection transforms the original features representing wind speeds at different grid points into a reduced set of features. By selecting the n principal components with the largest eigenvalues, PCA ensures that these new features capture most of the essential variation within the wind field data.

6.3.2 Discrete Cosine Transform

The Discrete Cosine Transform (DCT) is a frequency-domain transformation technique closely related to the Discrete Fourier Transform (DFT). Unlike DFT, which uses complex exponentials (sines and cosines), DCT employs only real-valued cosine functions as a basis. This characteristic, along with its energy compaction property, makes DCT particularly valuable for dimensionality reduction and data compression. For a discrete signal $x(n)$ of length N , the 1D DCT is calculated using the Eq. 6.4

$$z(k) = \sum_{n=0}^{N-1} x(n) \cdot \cos\left(\frac{\pi}{N} \left(n + \frac{1}{2}\right) k\right) \quad (6.4)$$

where $z(k)$ is the k -th frequency component of the signal transformed using DCT. Due to the nature of wind fields generated using TurbSim, a mathematical correspondence can be established between a 2D spatial wind field grid and an image. Suppose a wind field is defined on a discrete grid of size $(N_x \times N_y)$, each point on the grid (i, j) , representing a spatial location in the wind field, can be assigned to a pixel in the image. The magnitude of the wind speed at a point in the grid, denoted by $(X(i, j))$, determines the intensity of its corresponding pixel (see Fig. 6.2). This conceptual mapping connects wind field analysis to image processing techniques. Treating wind fields like images offers a particular advantage in utilising powerful mathematical tools like 2D DCT to analyse wind speed variations across spatial scales and directions. A key property of the 2D DCT is its separability, where




Figure 6.2: Modelling wind speed data as image

the 2D DCT can be computed by applying a 1D DCT along each row of the wind field grid, followed by a 1D DCT along each column of the intermediate result (or vice versa). Separability significantly reduces the computational complexity of the transformation compared to calculating all frequency components simultaneously. The 2D DCT of this wind field (X), denoted by $Z(u, v)$, is calculated using the Eq. 6.5:

$$Z(u, v) = \alpha(u)\alpha(v) \sum_{i=0}^{N_x-1} \sum_{j=0}^{N_y-1} X(i, j) \cos\left(\frac{\pi(2i+1)u}{2N_x}\right) \cos\left(\frac{\pi(2j+1)v}{2N_y}\right) \quad (6.5)$$

such that,

$$\alpha(u) = \begin{cases} \frac{1}{\sqrt{N_x}}, & i = 0 \\ \sqrt{\frac{2}{N_x}}, & 1 \leq i \leq N_x - 1 \end{cases} \quad (6.6)$$

and

$$\alpha(v) = \begin{cases} \frac{1}{\sqrt{N_y}}, & j = 0 \\ \sqrt{\frac{2}{N_y}}, & 1 \leq j \leq N_y - 1 \end{cases} \quad (6.7)$$

where u and v are frequency indices representing spatial frequencies in the transformed domain and $\alpha(u)$ and $\alpha(v)$ are normalisation factors ensuring orthogonality

of the basis functions. The coefficients ($Z(u, v)$) obtained from the 2D DCT transformation represent the strength of different frequency components present within the wind field. These frequency components correspond to variations in wind speed at different spatial scales. Low-frequency coefficients, associated with low values of u and v , represent large-scale spatial patterns or trends across the wind field. These capture smooth variations over spatial domains in the wind field data. Conversely, high-frequency coefficients represent finer-scale, localised wind speed fluctuations and turbulent structures. These coefficients capture rapid, often less spatially organised changes in wind speed over short distances within the wind field. By analysing the energy distribution across different DCT coefficients, the dominant spatial scales of variation present in the wind field can be identified. A wind field with strong, large-scale patterns will exhibit significant energy concentrated in its low-frequency coefficients. In contrast, a wind field characterised by smaller-scale turbulent fluctuations will have more energy within its high-frequency coefficients. DCT's ability to concentrate a wind field's energy within a few low-frequency coefficients makes it a powerful tool for dimensionality reduction. By calculating the 2D DCT of a wind field and ranking the coefficients by their magnitudes, a subset of coefficients that captures the majority of the wind field's essential variations can be identified. Retaining only these significant coefficients and discarding the smaller, high-frequency coefficients results in a compressed representation of the wind field. This compression significantly reduces the computational complexity of subsequent analysis and modelling while preserving the dominant spatial patterns in the wind field.

Both PCA and DCT can transform a wind field grid into different representations, providing a different basis for capturing the spatio-temporal variations. However, the PCA and DCT transformations do not directly reduce the number of features, rather they just transform them into a different representation. Therefore, to achieve the best model performance, the selection of key components from the transformed data remains crucial. In line with the detailed analysis presented in Chapter 5 on feature selection for multivariate time series problems, the Recursive Feature Addition (RFA) method is used here to select an optimal subset of features from both the PCA and DCT representations. RFA is chosen for its proven ability (established in Chapter 5) to identify the strongest feature combinations while maintaining manageable computational costs. The remainder of this chapter presents the development of an LSTM model utilising the dimensionality reduction approach combined with the RFA feature selection method to predict wind turbine

response under realistic, stochastic spatio-temporal wind fields.

6.4 Data Generation: Selection of Input Variables

As a surrogate model learns the relationship between input and output purely based on the data passed during the training process, the generation of a representative dataset covering the various possible scenarios that can be encountered on-site cannot be overstated. To accurately predict wind turbine responses, it is essential to generate realistic wind fields that reflect the inherent variability of wind. Careful selection of wind field parameters and their distributions is crucial, as these define the conditions for which the surrogate model is applicable. While the surrogate model is capable of extrapolating beyond these bounds, its prediction accuracy in those regions is uncertain. Design standards (IEC 61400-1 [79]) recommend various design load cases for comprehensive load and response analysis, encompassing wind loading, wave loading, and control system scenarios. Since this thesis focuses on an onshore wind turbine, wave-loading cases are not considered. Additionally, control system-related load cases, often representing accidental or short-lived events, are to be evaluated only for limited scenarios, resulting in a small number of simulations that can be evaluated using detailed numerical models without significant computational cost. This study prioritises normal operating conditions, which significantly influence fatigue loading and power production. Analysing these frequent conditions for statistically significant assessment requires running a large number of simulations, which leads to significant computational costs.

Realistic wind field generation, capturing the spatio-temporal variability of natural wind patterns, is essential for ensuring that surrogate models accurately predict wind turbine loads and power production. While a number of parameters govern wind field generation, including mean wind inflow direction ($HFlowAng$, $VFlowAng$), air density, turbulence length scale, and anisotropy factor, researchers have established that the turbulence intensity (TI), mean wind speed (U), and power law exponent (α) have the highest impact on wind turbine loads under normal operating conditions. This focus aligns with the findings of Dimitrov et al. [102], where a sensitivity analysis of environmental factors demonstrates that mean wind speed and turbulence intensity dominate fatigue and extreme loading, while parameters like wind direction, turbulence length scale, and air density exhibit minimal influence within the parameter bounds considered in their analysis. As the focus

of this thesis is to accurately model the blade response, the inclusion of the power law exponent (α) is particularly important as wind shear induces cyclic changes in the wind velocity experienced by the rotating blades, significantly influencing blade root loads. These concerns are also highlighted in Section 5.4 of Chapter 5. To this end, the TI , U and α are selected for generating the wind fields. The parameter bounds used for generating the wind speed data are presented in Table 5.1. Using these parameter bounds along with the Sobol sampling method, a total of 50 wind speeds ranging from cut-in to cut-out wind speeds are generated.

The spatial correlations, temporal variability, and turbulence inherent in wind make it a fundamentally random process. This randomness stems from aleatory uncertainty due to the inherent variability of wind speed and epistemic uncertainty arising from limitations to perfectly model or measure all factors influencing wind behaviour. To train a surrogate model capable of accurately predicting wind turbine responses under these uncertainties necessitates the generation of multiple realisations (ensembles) for each set of environmental parameters. Quantifying these uncertainties is essential for making realistic risk assessments and ensuring reliable wind turbine designs. However, identifying the minimum number of seeds required so that the surrogate model can efficiently identify the underlying patterns in these wind fields is a crucial factor governing the amount of data required for training these models. In a different application, the IEC guidelines [79] recommend considering six random realisations for each mean wind speed for computing fatigue damage, implying that six seeds can effectively capture the variability in wind fields. Further, these guidelines recommend keeping environmental conditions constant within a wind speed bin. This approach, combined with the inherent variability in TI and α , leads to a large number of required simulations. A study by Hübler, Gebhardt, and Rolfes [188] demonstrates that assuming constant environmental conditions within a wind speed bin does not fully capture the uncertainty in fatigue damage. To address this uncertainty, the authors recommend scattering environmental conditions within a bin for a more comprehensive assessment. A combination of these sampling approaches is used in this study to simulate the wind fields required for training the surrogate model. Scattering environmental conditions within a wind speed bin also acts as a crucial data augmentation technique for the surrogate model, which prevents overfitting to a limited set of wind patterns and significantly improves generalisation. A well-generalised model is essential for accurate predictions when encountering real-world wind conditions that may dif-

fer from the specific realisations used during training. To this end, a total of 15 realisations are generated for each mean wind speed U in this study, assuming a random combination of $RandSeed1$, TI , and α . This ensures a broader sampling of wind patterns while considering computational constraints. As such, wind fields are generated using TurbSim with the following key settings:

1. Spectral Model: Kaimal spectral model (IECKAI)
2. Turbulence Model: Normal Turbulence Model (NTM)
3. Grid Size: 25 x 25 grid with a spatial extent of 285m x 285m
4. Temporal discretisation: 0.05-second

These parameter choices ensure that the generated wind fields match the scale and characteristics relevant to the IEA-15MW wind turbine model being studied. This level of spatial and temporal resolution is necessary to capture wind patterns that significantly influence loads. The MBD model of the IEA-15MW wind turbine is evaluated for these wind fields, and the 0.005 second time step is used to ensure accuracy and numerical stability. The structural response is then downsampled to align with the temporal resolution of the wind field data. The wind field is represented by a grid 25×25 for a duration of 600 seconds, resulting in a large input matrix 12000×625 . This high dimensionality underscores the importance of dimensionality reduction techniques. PCA and DCT dimensionality techniques, in combination with the RFA feature selection algorithm, are employed to reduce the feature space and identify the most influential parameters, as discussed earlier in this chapter.

6.5 Selecting the output parameter

The surrogate model's ability to provide comprehensive insights depends on strategically selecting output parameters. While Section 6.4 outlined the reasoning behind the selection of input parameters, this section focusses on identifying the outputs that directly represent the fundamental dynamics of the blade and can be used to derive other relevant quantities. Building upon the focus on blade deformations in Chapter 5, this section expands the capabilities of the surrogate model by target-

ing the prediction of fundamental parameters related to blade degrees of freedom (DOF).

Creating surrogate models to predict the primary variables of interest has become a conventional practice. Typically, these variables are related to load components, such as fatigue damage or reactions in the tower and blades. However, these load components are derived quantities as the numerical models are essentially employed to determine responses at individual DOF, which are subsequently used to derive the load reactions. Generally, the mapping from DOF response to load reactions is straightforward and does not necessitate extensive computational resources. For example, in this study, the numerical model adopts the normal mode summation method to simulate the flexible components (tower and blades). The blade deflection at any point located at a distance x from the blade root at any time instance t is given by Eq. 6.8

$$\delta_{B1_{op}}(x, t) = \phi_{B1F1}(x)q_{B1F1}(t) + \phi_{B1F2}(x)q_{B1F2}(t) + \phi_{B1E1}(x)q_{B1E1}(t) \quad (6.8)$$

The symbols are explicitly defined in the list of acronyms. By incorporating this definition alongside the moment-curvature relationship (see Eq. 6.9), the bending moment can be calculated using Eq. 6.10.

$$\frac{d^2\delta_{B1_{op}}}{dx^2} = \frac{M(x)}{EI(x)} \quad (6.9)$$

$$M(x) = EI(x) \left(\phi_{B1F1}(x)'' q_{B1F1}(t) + \phi_{B1F2}(x)'' q_{B1F2}(t) + \phi_{B1E1}(x)'' q_{B1E1}(t) \right) \quad (6.10)$$

As mode shapes represent inherent features of a structural element and remain constant throughout the analysis, the curvature of a mode shape at any given point on the blade remains constant. Given this, if the response at the blade's DOFs is established, the bending moments can be subsequently computed through this relationship. Following this rationale, the objective shifts from predicting derived quantities to forecasting the responses of individual DOFs. This methodology ensures that the surrogate model is versatile enough to compute various derived quantities and conduct the corresponding analyses.

The process of developing a numerical model involves assuming various DOFs to capture structural behaviour, the number of DOFs depending on the chosen mod-

elling technique and the desired accuracy of the solution. To obtain complete information about the entire system, the response is required at each DOF. However, creating a surrogate model for each DOF would require the development of multiple surrogate models. As the number of DOFs increases, the computational burden associated with training and implementing individual models can increase significantly. To address this, multi-output learning, also known as multi-task learning in the context of machine learning, offers a compelling solution to these challenges [189, 190]. By leveraging the inherent relationships between different outputs within a system, multi-task learning enables the development of a single, unified model capable of simultaneously predicting multiple outputs for a given input. This streamlines the modelling process and enhances the model's capacity to capture complex interactions and dependencies within the system. One of the key advantages of multi-task learning is its ability to exploit shared information across related outputs. In the context of wind turbine blade analysis, the responses of different DOFs are often governed by common load elements, reflecting the holistic behaviour of the structure. By jointly modelling these responses within a multi-task learning framework, the model can leverage shared features and patterns, leading to improved generalisation and predictive performance. Referring to the numerical model developed in this study, the blades are modelled using three DOFs: q_{B1F1} , q_{B1F2} and q_{B1E1} . The multi-output learning paradigm is employed to develop an LSTM model to forecast the time history responses of these three DOFs. Furthermore, to be consistent with the results presented in Chapter 5, multi-task learning is also used to predict in-plane and out-of-plane blade deformations. The results obtained through this approach are presented in the next section.

6.6 Numerical Results

This section presents results that demonstrate the efficiency of LSTM models in predicting the dynamic response of wind turbine blades under turbulent wind conditions. Initial models using dimensionality reduction techniques (PCA and DCT) were explored, followed by a multi-stage approach for enhanced accuracy. Two modelling scenarios were considered to assess model performance under varying levels of uncertainty and build confidence in their applicability to real-world situations:

- Scenario 1: Constant TI (10%) and power law exponent (0.2), consistent with

the normal wind profile model, with variable mean wind speed. This scenario isolates the effect of the variations in mean wind speed and offers a baseline for model performance.

- Scenario 2: Variations in all parameters (mean wind speed, TI , wind shear and random seed) to capture a wider range of realistic wind conditions. This scenario tests the robustness of the model under greater uncertainty by introducing the additional complexities associated with fluctuating turbulence and wind shear.

The performance of the model in these scenarios was analysed and discussed.

6.6.1 One-stage modelling for blade response prediction

In the initial stage of the analysis, the models were developed to predict the blade deflection responses using the turbulent wind field data. PCA and DCT techniques were employed to reduce the dimensionality of the wind data. Although these techniques can capture some aspects of the blade response, their advantages and limitations will be explored in the next section. PCA converts the raw wind speed data into “principal components” ranked according to the amount of explained variance. A sample figure showing the first ten principal components of a wind speed of 12 m/s with a TI of 10% and $\alpha = 0.2$ is shown in Fig. 6.3. This figure shows that each principal component has a unique mean value and frequency-capturing patterns in spatio-temporal velocity fields. Similarly, the DCT transforms the features based on frequency content in the horizontal and vertical directions. The results of DCT are shown in Fig. 6.4, presented as a 5×5 matrix where the spatial frequency in the horizontal direction increases from left to right, and the frequency in the vertical direction increases from top to bottom.

The blade deflection results considering the first five principal components are shown in Fig. 6.5, while predictions with the 3×3 DCT components are presented in Fig. 6.6.

These figures highlight both the capabilities and limitations of the models trained using information from the wind speed data alone. Although information from wind speed data can effectively capture the mean trends in blade deflection, it is difficult to model higher-frequency responses. This limitation stems from the





Figure 6.3: Principal components of spatio-temporal wind field

choice of input features, which lacks direct information about the control parameters (rotor speed and blade pitch angle), which significantly impact the dynamics of blade response. During the circular rotations of the blade, the wind speed acting on the blade varies significantly with the angular position of the blade due to the wind shear effect. In one revolution, the wind speed experienced by the blade varies roughly in a sinusoidal fashion. This cyclic change in wind speed in each rotation creates a cyclic variation in the deformation response to a mean component governed by the low-frequency component of the wind speed. Without incorporating these control parameters, the models cannot fully capture the resulting dynamic fluctuations around the mean response.


In conclusion, these results demonstrate the inherent limitations of relying solely on wind speed data for accurate modelling of dynamic wind turbine blade responses, highlighting the need to explicitly account for data on control parameters, specifically the rotor speed and blade pitch angle. These parameters play a crucial role in driving higher-frequency fluctuations and stochastic effects on the blade. To address these challenges and improve model accuracy, a multi-stage modelling approach presents a promising avenue. By first predicting control responses and then incorporating this information into blade response models, the model can effectively capture the complex dynamics governing wind turbine behaviour. Recently, multi-

Figure 6.4: 5×5 submatrix of DCT representation of wind field


stage modelling has been used to gain incremental information in the emulation of dynamic systems, offering the potential for significantly enhanced predictive capabilities for complicated dynamic systems, including wind turbines [187]. A similar approach has been followed in this work. The potential of a multi-stage approach for modelling the blade dynamics is explored in the next section.

6.6.2 Multi-stage modelling for blade response prediction

As highlighted in the previous section, accurate prediction of wind turbine blade responses requires information beyond the wind speed data. Two key control parameters, rotor speed and blade pitch angle, play a critical role in governing the blade response. Rotor speed is a control parameter regulated through the generator torque controller for power output maximisation in below-rated wind conditions. The pitch angle of the blade, in turn, controls the aerodynamic forces on the blades and is actively managed to maintain a constant power output for the above-rated wind speed. Knowledge of these control parameters is essential to develop an effective predictive model of blade response. For this study, the ROSCO controller [41, 152] is used to optimise the performance of the wind turbine by regulating the rotor speed and the angle of rotation of the blade. ROSCO has been developed to provide a modular reference wind turbine controller that represents industry standards and




(a) Comparison of blade out-of-plane deflection predictions




(b) Comparison of blade in-plane deflection predictions

Figure 6.5: LSTM Model predictions of blade deflection using PCA features



(a) Comparison of blade out-of-plane deflection predictions



(b) Comparison of blade in-plane deflection predictions

Figure 6.6: LSTM Model predictions of blade deflection using DCT features




Figure 6.7: Multi-stage modelling approach for predicting blade response

performs comparable to or better than existing reference controllers. Since different controller algorithms can be applied to optimise the performance of wind turbines, which significantly impacts blade dynamics, the surrogate models for predicting control parameters depend on the underlying controller logic. Therefore, while the specific models developed in this study are applicable only to the ROSCO controller, the general methodology can be easily adapted to other controllers by training on the relevant datasets. To capture the impact of these control parameters, a multi-stage modelling approach is adopted. In this approach, an LSTM model is trained on the wind speed data to predict the rotor speed and blade pitch angle. These results are then combined with the wind speed data and act as input to a second LSTM model, which aims to predict the blade response. In essence, the multi-stage modelling approach divides the model development into smaller sub-tasks, which provide incremental information about the final output through a series of sub-tasks. This concept is demonstrated in Fig. 6.7


In this section, the use of LSTM models to predict the rotor speed and the blade pitch angle is discussed using wind speed data as input. Initially, a multi-task learning approach was explored, such that the wind speed data was used to simultaneously predict rotor speed and blade pitch angle. Subsequently, separate models for both control parameters were developed and found to be more effective. The results for both approaches are presented in this chapter, highlighting the advantages and limitations of the multi-task learning approach. In the initial stage of the investigation, a multi-task learning approach was explored to concurrently predict both rotor speed and blade pitch utilising a unified LSTM architecture. A fundamental prerequisite for multi-task learning is the identification of a shared

set of input features capable of holding sufficient information to accurately predict all targeted model outputs. Consequently, a multi-task learning LSTM model was trained based on the results of PCA and DCT techniques. The results of this multi-task learning framework are presented in Fig. 6.8 and Fig. 6.9. These results showed that while rotor speed variations are effectively captured by the PCA output, the dynamic behaviour of blade pitch was more accurately predicted from the results obtained using DCT.


These results illustrate the distinct contribution of each dimensionality reduction technique to the predictive performance of the models. In light of these results, an individual LSTM model was developed for each control parameter separately. This approach is utilised in the next section.

6.6.2.1 Predicting rotor speed response using PCA

Since the PCA approach was found to be more efficient in predicting rotor speed, an LSTM model was trained using principal components of the wind speed data to predict the rotor speed variations. The model architecture is shown in Fig. 6.10. The sequence input layer inputs sequential data to the network, where the size of the sequence layer is equal to the number of input features. The LSTM layer is used to learn long-term dependencies between time steps and extract temporal patterns from sequential data. This layer performs additive interactions, which can help improve gradient flow over long sequences during training. Since the model was trained for a wide range of mean wind speeds, a normalisation layer was used, which normalises a mini-batch of data across all channels for each observation independently. The use of the normalisation layer after the learning layers helped speed up the training and was found to improve the model performance. A fully connected layer was then used to combine the temporal patterns and transform these extracted features. Although a fully connected layer extracts the learnt information from the LSTM layers, it also increases the risk of overfitting. To avoid this, a dropout layer was added to the model. A dropout layer randomly sets input elements to zero with a given probability, which helps avoid overfitting the training data and ensures the network's better generalisation ability. In this model, a constant dropout probability of 10% was assumed for each dropout layer. The first fully connected layer was used to aggregate the information from the LSTM layer, while the second fully connected layer was used to represent the number of output parameters. To




(a) Comparison of blade pitch angle predictions




(b) Comparison of rotor speed predictions

Figure 6.8: LSTM Model predictions of control parameters using PCA features



(a) Comparison of blade pitch angle predictions



(b) Comparison of rotor speed predictions

Figure 6.9: LSTM Model predictions of control parameters using DCT features




Figure 6.10: LSTM Model architecture for predicting rotor speed

this end, the final fully connected layer in the model matches the number of target variables and has only one neuron for the model used to predict the rotor speed. For multi-task learning models, the number of neurons in the final layer will be greater than one. This model was trained using the Adam optimisation algorithm [151]. Initially, the first five principal components, which cumulatively capture 50% of the variance in the wind speed data, were selected to capture the dominant patterns in the wind speed data. This feature set was further filtered using RFA by iteratively adding features and evaluating the performance of the model to find the optimal feature set that yields the best results.

The model was first trained for the first uncertainty scenario in which the wind speed data was generated with constant turbulence intensity (TI) and wind shear. Fig. 6.11 shows the normalised RMSE obtained from validation data using different principal components as input features. This figure shows that the first




Figure 6.11: Normalised RMSE in rotor speed prediction using different principal components




Figure 6.12: Comparison of rotor speed predictions of LSTM model using PCA feature at constant TI and wind shear

component corresponds to the lowest prediction error for the validation data set. Adding more input features using the RFA method did not increase the accuracy of the model predictions. To this end, the first principal component of the spatio-temporal wind field is used as the preferred input feature to predict the rotor speed. A visual representation of the predictions of the LSTM model to the actual rotor speed across different mean wind speeds is presented in Fig. 6.12, demonstrating a close match that emphasises the model's ability to capture rotor speed dynamics using solely the first principal component of the wind speed data. The overall efficiency of this approach can be contextualised by the fact that the proposed approach reduces 625 input features to a single variable capable of capturing the overall dynamics. This underscores the importance of dimensionality reduction techniques such as PCA to simplify complex data sets without sacrificing crucial information.

Next, the model was trained for the higher uncertainty scenario using data that incorporate variations in all four wind parameters: TI , wind shear, mean wind speed, and random seed (adding stochasticity). This enhanced data complexity leads to a reduction in model accuracy, reflected in Fig. 6.13. A quantitative comparison of normalised RMSE achieved in modelling scenarios assuming constant and variable TI and wind shear is presented in Table 6.1. An increase in normalised




Figure 6.13: Comparison of rotor speed predictions of LSTM model using PCA feature at random TI and wind shear

Modelling Scenario	Normalised RMSE
Scenario 1	0.210
Scenario 2	0.264

Table 6.1: Comparison of the effect of uncertainty level on rotor speed prediction

RMSE of 26% was observed when uncertainty in TI and wind shear are considered. This expected result is attributed to the higher level of uncertainty in the output data. A larger training dataset may be required to address the higher uncertainty contributed by variations in TI and wind shear. Since the prediction of rotor speed is one of the sub-tasks of the multi-stage approach aimed at predicting the blade response, the impact of how uncertainty in this stage propagates through the multi-stage blade response model will be explored in later sections of this chapter. In addition, for the final model training, the rotor speed was transformed into the angular position of the blade so that it could more effectively capture the frequency of rotation. Since the angular position varies within the well-defined limits $[0, 2\pi]$, the uncertainty at this modelling stage may not significantly impact the final predictions.




Figure 6.14: Visual representation of DCT components considered in RFA method

6.6.2.2 Predicting blade pitch response using DCT

The LSTM model used to predict the blade pitch response shares the same architecture as shown in Fig. 6.10. Wind speed data, transformed using DCT, served as the input, and RFA was used to determine the most informative subset of DCT features for optimal prediction accuracy. The model was trained using the Adam optimisation algorithm.

The wind speed grid at each time step was treated as an image and processed using the 2D DCT. The 2D DCT decomposes the image into different spatial frequency components, allowing one to extract patterns within the wind field that influence blade pitch response. The process of modelling wind speed data as an image and processing it through 2D DCT inherently distributes relevant information across the off-diagonal terms of the output matrix. As such, the off-diagonal terms in the DCT matrix share mutual information about spatial frequencies in the horizontal and vertical directions. To capture this information, DCT components were grouped to form feature subsets, each representing an incremental level of detail. A mathematical definition of a feature subset z_i is given by Eq. 6.11

$$z_i = Z_{n \times n}(1 : i, 1 : i) - Z_{n \times n}(1 : i - 1, 1 : i - 1) \forall 1 \leq i \leq n \quad (6.11)$$

A visual representation of component groups for a 5×5 DCT matrix is shown in Fig. 6.14. Using this feature definition along with the RFA scheme, the LSTM model was trained to predict the blade pitch response.

The model was initially trained on wind speed data with constant turbulence intensity (TI) and wind shear. The initial 5×5 sub-matrix of the 25×25 matrix obtained through 2D DCT was selected as the dominant subset of input features as it represents the lower frequency patterns which can more effectively capture the general trends in the wind speed data. Fig. 6.15 presents the normalised RMSE




Figure 6.15: Normalised RMSE in blade pitch prediction using different DCT components

obtained for the LSTM model trained on each DCT component group (presented in Fig. 6.14), showing that the first component, representing the lowest frequency patterns extracted by the DCT, achieves the best performance in predicting the blade pitch angle. This signifies that low-frequency fluctuations in wind speed play a dominant role in driving the blade pitch response. The findings of the RFA showing the best predictive capacity of the LSTM under the low frequency component can be explained using control theory. The control system governing blade pitch angle is designed to optimise turbine performance and stability by filtering out high-frequency disturbances and adjusting the blade pitch angle in a more deliberate and controlled manner. This intentional smoothing of blade pitch variations serves to mitigate the impact of sudden changes in wind speed and maintain the overall stability and efficiency of turbine operation. Fig. 6.16 further reinforces this connection by comparing the actual blade pitch response to the LSTM model predictions at different wind speeds. A good agreement in the blade pitch response across different mean wind speeds showcases the generalisation ability of the LSTM model.

Next, this approach was extended to data with variations in all four wind parameters: TI , wind shear, mean wind speed, and random seed. As shown in Fig. 6.17, an increase in error was observed, likely due to the higher uncertainty of the output data when considering these wider variations. Table 6.2 offers a quantitative comparison of normalised RMSE for both uncertainty scenarios. A 23% increase in normalised RMSE was observed, considering variations in all four parameters. A similar behaviour was observed when predicting the rotor speed response. The impact of this reduced accuracy on the final blade response prediction




Figure 6.16: Comparison of blade pitch angle predictions of LSTM model using DCT features at constant TI and wind shear




Figure 6.17: Comparison of blade pitch predictions of LSTM model using DCT feature at random TI and wind shear

Modelling Scenario	Normalised RMSE
Scenario 1	0.142
Scenario 2	0.175

Table 6.2: Comparison of the effect of uncertainty level on LSTM Model performance

will be explored in later sections.

6.6.3 Predicting blade response

In this section, a two-layer LSTM model architecture was implemented to predict wind turbine blade response. To explore the effectiveness of different dimensionality reduction techniques, two sets of input features are considered:

1. Wind turbine data processed through PCA, combined with the angular position of the blade and blade pitch angle predictions from the LSTM models described in Sections 6.6.2.1 and 6.6.2.2.
2. Wind speed data processed through DCT, also combined with the angular position of the blade and blade pitch angle predictions from the LSTM model.

The results of these two approaches were compared to assess the capabilities of PCA and DCT in the context of blade response prediction. A multi-task learning approach was employed to simultaneously predict the blade response. Initially, the model was trained to predict in-plane and out-of-plane deflections, which was subsequently extended to predict the response of individual blade DOFs in the flapwise and edgewise directions.

6.6.3.1 Blade response prediction with PCA features

Following the approach described in the previous sections, the analysis first focusses on the results obtained under constant turbulence intensity (TI) and wind shear. The architecture of the model for blade response prediction is detailed in Fig. 6.18. The LSTM architecture is designed through an iterative process, and hyperparameter optimisation is used to identify the most optimal configuration. The number of hidden units in the learning layers was used as a hyperparameter in this study. The performance of the model is evaluated using normalised RMSE between the predictions in the validation data and the actual results. Table 6.3 summarises the hyperparameter exploration. A progressively smaller size is used for the range of hyperparameters, as consistently decreasing the number of hidden units was found to deliver better results. Using the angular position of the blade and blade pitch angle as constant features, the principal components of the wind speed data were se-




Figure 6.18: LSTM Model architecture for predicting blade response


Hyperparameter	Parameter range	Optimised values
LSTM layer - 1	[100 - 200]	138
LSTM layer - 2	[50 -100]	63
Fully connected layer - 1	[1-50]	44

Table 6.3: Summary of hyperparameter exploration to optimise the number of hidden units in learning layers

lected using the RFA algorithm. This approach enables accurate predictions of both in-plane and out-of-plane blade deformations (Fig. 6.19). In particular, the second principal component, when paired with the angular position of the blade and blade pitch angle, yielded the best performance. This suggests that the second principal component likely captures patterns more relevant to predicting blade deformations than the first principal component, despite the fact that the first component captures the maximum variance in the dataset. These findings further reinforce the importance of using feature selection approaches in obtaining a simplified model representation with a small set of the most informative features. This method is further extended to predict the response in individual blade DOFs, which are the most fundamental quantities in numerical modelling, and the corresponding results are presented in 6.20. Using the fifth principal component as the input feature in conjunction with the blade angular position and the blade pitch angle produced the lowest normalised RMSE to predict blade deflections, while the first principal component in conjunction with control parameters delivered the best performance for predicting the response in the blade DOFs. Here, it can be seen that a higher level of accuracy was obtained in the prediction of DOFs q_{B1F1} and q_{B1E1} as compared to q_{B1F2} . This can be attributed to the higher frequency of the mode shape for q_{B1F2} .




(a) Comparison of blade out of plane deflection predictions




(b) Comparison of blade in-plane deflection predictions


Figure 6.19: Multi-task learning LSTM Model predictions of blade deflections using PCA features



(a) Blade first flapwise DOF response



(b) Blade second flapwise DOF response



(c) Blade first edgewise DOF response

Figure 6.20: Multi-task learning LSTM Model predictions of blade DOFs using PCA features

Output Parameter	Normalised RMSE	
	PCA	DCT
Blade IP Deformation	0.074	0.063
Blade Oop Deformation	0.085	0.076
Blade DOF $-q_{B1F1}$	0.086	0.081
Blade DOF $-q_{B1F2}$	0.095	0.091
Blade DOF $-q_{B1E1}$	0.067	0.065

Table 6.4: Comparison of LSTM Model performance trained using PCA and DCT features


Further, the absolute magnitude of q_{B1F2} shows that this DOF does not contribute to the total deformation as much as q_{B1F1} . However, it is evident from Fig. 6.20 that a good approximation in predicting all the DOFs is achieved by capturing the governing dynamics in all the parameters.

While these results were promising under constant TI and wind shear conditions, extending the model to scenarios with variations in all four wind parameters presented a greater challenge. This higher uncertainty aspect and its impact on prediction accuracy will be explored in subsequent sections.


6.6.3.2 Blade response prediction with DCT features

Building upon the PCA approach, the use of DCT features for blade response prediction is explored in this section. The LSTM model follows the same model architecture presented in the previous section (Fig. 6.18). Following the methodology adopted in Section 6.6.3.1, the model architecture was trained using DCT features. Fig. 6.21 and Fig. 6.22 demonstrate the accuracy of DCT in predicting the blade response, exhibiting a close match to the true response. The best performance for blade response prediction was obtained using the second group of DCT components presented in Fig. 6.14 as input features, while the blade DOF predictions achieve higher accuracy under the combination of the second and third groups of DCT components.

Table 6.4 provides a comparison of the normalised RMSE obtained for models developed using the PCA and DCT features, highlighting the superior performance of DCT in predicting the blade response. This aligns with the expectation that DCT, by focussing on spatial frequency patterns, is well suited to capture the dy-




(a) Comparison of blade out of plane deflection predictions




(b) Comparison of blade in-plane deflection predictions


Figure 6.21: Multi-task learning LSTM Model prediction of blade deflections using DCT features



(a) Blade first flapwise DOF response



(b) Blade second flapwise DOF response



(c) Blade first edgewise DOF response

Figure 6.22: Multi-task learning LSTM Model predictions of blade DOFs using DCT features

namic variations in wind fields that drive the blade responses. To further test its capabilities, the DCT model is trained to predict blade response under varying TI and wind shear. The results of this case are discussed in the next section.

6.6.4 Uncertainty Quantification in Model Predictions

In data-driven modelling, quantifying the uncertainty associated with model predictions is essential to ensure reliability and robustness. Traditional point predictions often fail to capture the inherent variability in the predictions, making it necessary to estimate confidence intervals that provide a probabilistic measure of uncertainty. Confidence intervals allow for a more informative representation of model predictions, offering a range within which the true values are expected to lie with a given probability. One of the widely used metrics for assessing the reliability of confidence intervals is the *Prediction Interval Coverage Probability* (PICP). The PICP quantifies the proportion of actual observations that fall within the predicted confidence intervals of the model, serving as an objective measure of the model's ability to estimate uncertainty. It is particularly useful in determining whether the underlying model errors exhibit Gaussian behaviour, which is a key assumption in many probabilistic forecasting methods. Mathematically, PICP is defined as

$$PICP = \frac{\sum_{i=1}^N \mathbb{I}(y_i \in [\hat{\mathbf{y}}_i^L, \hat{\mathbf{y}}_i^U])}{N} \times 100 \quad (6.12)$$

where N represents the total number of predictions, y_i is the true value at the time step i , and $\hat{\mathbf{y}}_i^L$ and $\hat{\mathbf{y}}_i^U$ denote the lower and upper bounds of the predicted confidence interval, respectively. The indicator function $\mathbb{I}(y_i \in [\hat{\mathbf{y}}_i^L, \hat{\mathbf{y}}_i^U])$ takes the value 1 if the actual value y_i falls within the confidence interval and 0 otherwise. For a model whose residual errors closely follow a Gaussian distribution, the PICP should be approximately equal to the expected confidence level (e.g., 95% for a 95% confidence interval). If the computed PICP is significantly lower than this threshold, it suggests that the model is underestimating uncertainty, leading to overly confident predictions that may not generalise well. In contrast, if the PICP is substantially higher than the expected level, the model may be overestimating uncertainty, generating excessively conservative predictions with confidence intervals that are wider than necessary. The deviation of the PICP value from the expected confidence level indicates the deviation from the Gaussian behaviour of the model

Output Parameter	Average PICP	
	PCA	DCT
Blade IP Deformation	94.54%	94.59%
Blade Oop Deformation	95.11%	95.03%
Blade DOF $-q_{B1F1}$	95.41%	95.51%
Blade DOF $-q_{B1F2}$	94.75%	94.17%
Blade DOF $-q_{B1E1}$	94.92%	94.97%


Table 6.5: Comparison of PICP values for LSTM Model trained using PCA and DCT features

prediction error. Since standard confidence intervals are typically derived under the assumption of normally distributed errors, a PICP close to the nominal confidence level indicates that this assumption holds reasonably well. A notable deviation from the expected PICP could imply that the residuals exhibit heteroscedasticity or non-Gaussian behaviour, necessitating alternative approaches such as quantile regression or non-parametric prediction intervals to more accurately capture uncertainty. The PICP values obtained for the model predictions based on the characteristics of PCA and DCT are presented in Table 6.5. Most values are close to 95%, indicating that the model errors are approximately Gaussian distributed and that the uncertainty is well captured by the confidence intervals. The suitability of the confidence intervals of the model achieved and the propagation of this uncertainty in subsequent tasks, such as reliability analysis or fatigue evaluation, remains the scope of future investigations.


6.6.5 Uncertainties and Future Scope

To address the challenges identified in the previous sections, this section focusses on analysing the impact of uncertainty in predicting the speed of the rotor and the angle of inclination of the blade on the prediction of the response of the blade. A drop in accuracy was observed when extending the model to scenarios with variations in all four wind parameters using a multi-stage modelling approach (see Fig. 6.23 and Fig. 6.24).

This can be attributed to errors in the predictions of the rotor speed (and, in turn, the angular blade position) and the pitch angle presented in previous sections. To further validate this assumption, the trained model was used to predict




(a) Comparison of blade out of plane deflection prediction




(b) Comparison of blade in-plane deflection predictions


Figure 6.23: Multi-task learning LSTM Model prediction of blade deflections using multi-stage modelling



(a) Comparison of Blade first flapwise DOF response



(b) Comparison of Blade second flapwise DOF response



(c) Comparison of Blade first edgewise DOF response

Figure 6.24: Multi-task learning LSTM Model predictions of Blade DOF using multi-stage modelling

Output Parameter	Multi-stage modelling	ROSCO control input
Blade IP Deformation	0.307	0.073
Blade Oop Deformation	0.186	0.093
Blade DOF $-q_{B1F1}$	0.203	0.096
Blade DOF $-q_{B1F2}$	0.171	0.102
Blade DOF $-q_{B1E1}$	0.363	0.072

Table 6.6: Comparison of Normalised RMSE for multi-stage modelling with ROSCO control output


the blade response by providing an accurate response corresponding to the rotor speed and pitch angle obtained from the numerical model. In this case, the results demonstrated improved performance in predicting both the deformation and the DOF responses (see Fig. 6.25 and Fig. 6.26).

A quantitative analysis of model performance is presented in Table 6.6; This shows a consistent increase in normalised RMSE where the model relies on the rotor speed and azimuth angle predicted using LSTM models as opposed to when the estimates of these parameters were used as input through the ROSCO controller algorithm.


This trend highlights that the LSTM model is capable of predicting the blade dynamics; however, the accuracy of predictions depends on the accuracy of input features achieved in the intermediate stage of the multi-stage modelling approach. This calls for a further improvement of the LSTM model that predicts the control parameters, which will be a focus of future work. However, in the current stage, the model can efficiently capture the dynamics, which could be used for quick preliminary estimates of fatigue analysis and site-specific performance. The next section focusses on analysing the computational advantage of the developed surrogate models.

6.6.6 Computational advantage of surrogate approach

The need to reduce the computational cost in predicting the dynamic response of a wind turbine is at the heart of the surrogate model development in this thesis. Having established the accuracy of the developed surrogate model, this section focusses on the computational advantages of using these models. A direct comparison




(a) Comparison of blade out of plane deflection prediction




(b) Comparison of blade in-plane deflection predictions


Figure 6.25: Multi-task learning LSTM Model prediction of blade deflections using ROSCO control input



(a) Comparison of Blade first flapwise DOF response



(b) Comparison of Blade second flapwise DOF response



(c) Comparison of Blade first edgewise DOF response

Figure 6.26: Multi-task learning LSTM Model predictions of Blade DOF using multi-stage modelling

of execution times was performed to quantify the degree of computational efficiency of these models. Both the full MBD model and the surrogate model were used to predict the blade response for a 600-second turbulent wind inflow generated using TurbSim. To ensure a fair comparison, only the essential blade DOFs and generator azimuth angle were activated within the MBD model. The surrogate modelling approach developed in this study offers a significant computational advantage in predicting blade response. The execution time comparison shows a 75-fold speed improvement using the surrogate model compared to the corresponding MBD model. This comparison highlights the computational efficiency of the developed surrogate models. The computational advantage allows for conducting a large number of simulations required for site-specific performance analysis. Furthermore, the ability of the surrogate model to quickly predict blade deformations and loads under varying wind and control inputs makes it a useful tool for comprehensive fatigue analysis of wind turbine blades. While quantifying the computational gains, all the simulations were performed on a system with an 8-core Intel Xeon CPU with a clock speed of 3.8GHz using 32GB RAM and running on Microsoft Windows 10 Pro.

6.7 Discussion

In this chapter, an approach that combines multi-stage and multi-task learning with dimensionality reduction techniques and a feature selection algorithm was presented. This combined method aims to improve the efficiency of developing LSTM models to predict blade response. The major contributions from this chapter can be summarised as follows.

- **Dimensionality Reduction:** The effectiveness of PCA and DCT in simplifying wind field data while retaining crucial information for prediction tasks was demonstrated. Both PCA and DCT, particularly when combined with recursive feature addition, helped to achieve efficient model configuration and improved prediction accuracy.
- **Rotor Speed and Blade Pitch Prediction:** LSTM models were developed for prediction of the rotor speed (using PCA features) and prediction of the blade pitch (using DCT features). These models achieved varying degrees of accuracy in capturing the dynamics of their respective target parameters at different uncertainty levels in wind conditions. Accurate rotor speed and blade

pitch information was identified as critical parameters for subsequent prediction of blade response.

- **Blade Response Prediction:** A multi-stage modelling approach was employed, where predicted control parameters were fed into an LSTM model to forecast blade deformations and DOF responses. Models trained using DCT features showed higher accuracy as compared to PCA features for this task, indicating DCT's ability to capture spatial frequency patterns driving blade dynamics.
- **Uncertainty and Future Directions:** The prediction accuracy of the models decreased when extended to scenarios with higher uncertainty in the wind parameters. Further analysis identified errors originating from inaccuracies in the control parameter prediction stage significantly impact blade response predictions. Future work will prioritise refinements to the LSTM models responsible for rotor speed and blade pitch prediction to enhance overall system performance across a broader range of operational conditions.

These models demonstrate good accuracy in capturing trends in the dynamic behaviour of wind turbine blades under normal operating conditions. Thus, their application is more suitable for fatigue analysis. They may not be optimal for reliability analysis, which emphasises extreme responses. Future research should aim to quantify the accuracy of the models in performing reliability and fatigue analyses.

Chapter 7

Conclusions and future work

The rapid evolution of wind turbine technologies, driven by the goal of achieving greater energy capture and reduced costs, has led to structures of unprecedented scale and complexity. Although larger turbines offer increased power potential, they also introduce new challenges stemming from reduced stiffness, increased loads, and a greater need for design optimisation. Traditional design approaches, based on partial safety factors, must be improved through the use of advanced analysis tools to ensure the reliability and cost-effectiveness of these next-generation turbines. The design of a wind turbine is governed by the external loads acting on the structure, which are inherently random in nature. One of the fundamental challenges in the design process lies in accurately quantifying the impact of uncertain environmental factors on the loads experienced by the turbines. Aerodynamic loads, in particular, are highly stochastic in nature as they are governed by the turbulent spatio-temporal wind fields. Wind flow patterns can vary significantly between installation locations, which impacts the performance of these structures. The uncertainty coming from the stochastic inflow represents an aleatoric form of uncertainty which cannot be eliminated, its impact can only be quantified. Quantifying the aleatoric uncertainty in the wind and performing the uncertainty propagation to examine its effect on the design loads requires extensive simulations. However, conventional high-fidelity multi-body dynamic (MBD) models often become computationally prohibitive for such large-scale analyses.

This thesis addresses the critical challenge posed by the high computational cost of MBD models by developing computationally efficient surrogate models ca-

pable of accurately predicting wind turbine responses. These surrogate models offer significant advantages in both the design and operational phases. In this chapter, a summary of this thesis is presented in the form of key contributions to this research and its limitations, together with directions for future work.

7.1 Summary

The surrogate models developed in this thesis offer a compelling solution to the computational bottlenecks of the MBD model and analysis workflows. Across all the methodologies explored, significant computational gains were achieved. A reduction in execution time of up to 75 times was demonstrated for certain stages of the analysis process. The computational efficiency of these models, coupled with the consistently high accuracy demonstrated, presents a high potential to extend the current design and analysis methods to address the high stochasticity contributed by the wind inflow more effectively.

Chapter 4 addressed a fundamental bottleneck in an MBD model by developing artificial neural network (ANN) based surrogate models to replace the computationally intensive BEM theory. Replacing the iterative root-finding method with high dimensional regression resulted in a 40x speed improvement for aerodynamic load calculation and a 35% reduction in the overall MBD simulation time. The ANN surrogates exhibited high accuracy in predicting the angle of attack at any blade node with a discrepancy of less than 0.5° . Implementing the physics-informed neural network (PINN) approach by enforcing constraints on latent variables through a physics-based loss function resulted in higher accuracy in most scenarios. The weighing factor used to balance the data-fitting objective and physics-based loss function presented insights into the optimum contribution of these constraints required for efficient model learning. Outside the MBD framework, these surrogate models can act as virtual sensors to gain insights into the loads experienced by blades during operations. A fundamental limitation of this approach lies in the need to develop airfoil-specific surrogates to accurately model loads across the entire blade span. This arises due to the unique aerodynamic responses of different airfoil profiles and can become prohibitive for modern blades with multiple airfoil sections. Chapter 6 directly addresses this challenge by exploring advanced multi-task learning methods that have the potential to streamline this process.

Building upon the efficient load calculations enabled by ANN surrogates, Chapter 5 develops a surrogate model to predict the dynamic response of wind turbine blades using Long Short-Term Memory (LSTM) networks. An innovative approach derived from combining the filter-based minimum redundancy maximum relevance and wrapper-type recursive feature addition feature selection approach was found to efficiently identify the most informative features from a high dimensional feature space. This method efficiently identifies the most informative features influencing blade deformation, enabling optimised LSTM models. Importantly, this analysis also has practical implications, suggesting that accurate deformation prediction is possible with data from only a few strategically placed sensors. The learning rate scheduling methodology implemented for optimal convergence further enhances model performance. The models trained following the methodology presented in Chapter 5 result in efficient LSTM models that achieve high accuracy with a fraction of the computational cost of an MBD model. Further, the robustness of the developed models to noisy input data highlights the possibility of virtual sensing applications.

Chapter 6 builds upon the previous work with a multi-stage, multi-task LSTM model trained to predict blade responses at selected degrees of freedom (DOFs) directly from spatio-temporal wind field data. The Principal Component Analysis (PCA) and Discrete Cosine Transform (DCT) methods used for reducing the high dimensionality of the spatio-temporal wind fields result in a highly efficient representation of wind fields using only 5-10 variables as opposed to 625 variables in the wind speed data generated using TurbSim. Combining these dimensionality reduction techniques with a feature selection algorithm resulted in an efficient surrogate model that can rapidly predict critical responses without requiring iterative aerodynamic calculations and numerical integration for full dynamic simulations. These models consistently achieved low normalised root mean square error of the order of 10^{-1} to 10^{-2} across different uncertainty levels, highlighting their accuracy in emulating blade dynamics. The high accuracy in predicting the response with a 75x speed up compared to the MBD model highlights the ability of these models to present an innovative workflow for evaluating wind turbine response, enabling faster design exploration. Since the model can be trained solely using the simulation data without requiring explicit knowledge of the model configuration, this approach presents an exciting solution for performing site-specific load and response analysis for commercial wind turbine models where the model details cannot be shared due

to intellectual property concerns. The ability to assess performance under a wide range of wind conditions is a significant advantage of this method.

7.2 Limitations and Future work

While the surrogate models developed in this thesis offer significant computational advantages and exciting potential applications, it is essential to acknowledge their fundamental limitations. The reliance on training data generated using deterministic structural geometry and material properties limits the application of these surrogates to reliability analysis and design exploration and may not have a direct application in design optimisation processes. Further, since the models are trained using the data under normal operating conditions, these models may not offer high accuracy in estimating response under extreme conditions. As such, the models can facilitate quick fatigue estimation as they are governed by the normal operating conditions however the surrogate models may not be best suited for reliability analysis which is governed by the extreme response. Nevertheless, the surrogate models developed in this study provide an efficient approach for investigating the system response governed by normal operating conditions, especially fatigue analysis. The LSTM models are trained to predict the dynamic response of blade deformations and response at different DOFs, which can be used to estimate blade root bending moments and corresponding stresses. By generating stochastic wind fields using a range of mean wind speed, turbulence intensity and wind shear, the corresponding dynamic response can be predicted using the LSTM model. This response can be used to perform fatigue analysis. Due to the low computational cost of running these surrogate models, the system response can be analysed for large parameter combinations, enabling detailed fatigue analysis of wind turbine blades. To this end, these models can accurately quantify the fatigue damage for different load distributions. Subsequently, these models can also enable reliability analysis accounting for uncertainty in external loads across various installation locations.

Future research can address the limitations of the developed surrogate models by accounting for the effect of system geometry on the loads and corresponding responses. Such models can be directly applied to efficiently predict structural response under random external loads at various system geometries. Successfully building these models can significantly reduce the computational cost of reliability-based design optimisation. Further, the advancements in the domain of machine

learning, such as transfer learning, present high potential in building up on the learnings of these models and supplementing it with the data generated from high fidelity models such as Computational Fluid Dynamics (CFD) for aerodynamics analysis and Finite Element (FE) models for the structural analysis to further improve the accuracy of these surrogates. Exploring the use of Neural Ordinary Differential Equations (Neural ODEs) could pave the way for developing models which learn from the underlying differential equations. These models are known to deliver better results when extrapolating the training data. Finally, these approaches could be extended to more challenging problems, such as floating offshore wind turbines and wind farm models, to realise the capabilities of the surrogate models to analyse and optimise these systems.

The insights and methodological advances presented in this work have implications beyond the immediate focus on wind turbine blades. The demonstrated success in surrogate modelling for complex flexible structures lays a foundation for extending these techniques to other engineering domains where computational efficiency and fast response prediction are critical.

Bibliography

- [1] Vohra, Karn et al. “Global mortality from outdoor fine particle pollution generated by fossil fuel combustion: Results from GEOS-Chem”. In: *Environmental research* 195 (2021), p. 110754.
- [2] Fuller, Richard et al. “Pollution and health: a progress update”. In: *The Lancet Planetary Health* 6.6 (2022), e535–e547.
- [3] Manisalidis, Ioannis, Stavropoulou, Elisavet, Stavropoulos, Agathangelos, and Bezirtzoglou, Eugenia. “Environmental and health impacts of air pollution: a review”. In: *Frontiers in public health* 8 (2020), p. 505570.
- [4] Liang, Huaying et al. “Association of outdoor air pollution, lifestyle, genetic factors with the risk of lung cancer: A prospective cohort study”. In: *Environmental Research* 218 (2023), p. 114996.
- [5] Mishra, Pabitra, Behera, Jiban Kumar, Jena, Anway Kumar, Bhattacharya, Manojit, and Behera, Bhaskar. “Risk analysis of air pollution correlates with socioeconomic and heart diseases”. In: *Spatial Modeling of Environmental Pollution and Ecological Risk*. Elsevier, 2024, pp. 87–119.
- [6] Liu, Dongyue. “Value evaluation system of ecological environment damage compensation caused by air pollution”. In: *Environmental Technology & Innovation* 22 (2021), p. 101473.
- [7] Grennfelt, Peringe et al. “Acid rain and air pollution: 50 years of progress in environmental science and policy”. In: *Ambio* 49 (2020), pp. 849–864.
- [8] Veltman, Karin, McKone, Thomas E, Huijbregts, Mark AJ, and Hendriks, A Jan. “Bioaccumulation potential of air contaminants: Combining biological allometry, chemical equilibrium and mass-balances to predict accumulation of air pollutants in various mammals”. In: *Toxicology and applied pharmacology* 238.1 (2009), pp. 47–55.
- [9] Agency, International Energy. *Global water consumption in the energy sector by fuel and power generation type in the Stated Policies Scenario, 2021 and 2030*. 2023. URL: <https://www.iea.org/data-and-statistics/charts/global-water-consumption-in-the-energy-sector-by-fuel-and-power-generation-type-in-the-stated-policies-scenario-2021-and-2030>.
- [10] Lindsey, R. and Dahlman, L. *Climate change: Global temperature*. NOAA, 2021.
- [11] Council, National Research et al. *Abrupt impacts of climate change: Anticipating surprises*. National Academies Press, 2013.

- [12] EEA. *Assessing the costs and benefits of climate change adaptation*. European Environment Agency, 2023.
- [13] Ahlbrandt, Thomas S. “Future petroleum energy resources of the world”. In: *International Geology Review* 44.12 (2002), pp. 1092–1104.
- [14] *The socioeconomic impacts of wind energy in the context of the energy transition*. KPMG, 2019.
- [15] Van Kuik, GAM et al. “Long-term research challenges in wind energy—a research agenda by the European Academy of Wind Energy”. In: *Wind energy science* 1.1 (2016), pp. 1–39.
- [16] Veers, Paul et al. “Grand challenges in the science of wind energy”. In: *Science* 366.6464 (2019), eaau2027.
- [17] Kirkegaard, Julia Kirch et al. “Tackling grand challenges in wind energy through a socio-technical perspective”. In: *Nature Energy* 8.7 (2023), pp. 655–664.
- [18] Anaya-Lara, Olimpo, Tande, John Olav, Uhlen, Kjetil, and Merz, Karl. *Off-shore wind energy technology*. John Wiley & Sons, 2018.
- [19] Kong, Kyungil, Dyer, Kirsten, Payne, Christopher, Hamerton, Ian, and Weaver, Paul M. “Progress and trends in damage detection methods, maintenance, and data-driven monitoring of wind turbine blades—A review”. In: *Renewable Energy Focus* 44 (2023), pp. 390–412.
- [20] Wu, Xianyou, Feng, Kai, and Li, Qing'an. “A Numerical Method for the Dynamics Analysis of Blade Fracture Faults in Wind Turbines Using Geometrically Exact Beam Theory and Its Validation”. In: *Energies* 17.4 (2024), p. 824.
- [21] Veers, Paul et al. “Grand challenges in the design, manufacture, and operation of future wind turbine systems”. In: *Wind Energy Science Discussions* 2022 (2022), pp. 1–102.
- [22] Veers, Paul et al. “Grand Challenges: wind energy research needs for a global energy transition”. In: *Wind Energy Science Discussions* 2022 (2022), pp. 1–8.
- [23] Veers, Paul et al. *Grand Challenges Revisited: Wind Energy Research Needs for a Global Energy Transition*. Tech. rep. National Renewable Energy Laboratory (NREL), Golden, CO (United States), 2023.
- [24] Murtagh, PJ, Basu, B, and Broderick, BM. “Along-wind response of a wind turbine tower with blade coupling subjected to rotationally sampled wind loading”. In: *Engineering structures* 27.8 (2005), pp. 1209–1219.
- [25] Dinh, Van-Nguyen, Basu, Biswajit, and Nielsen, Søren RK. “Impact of spar-nacelle-blade coupling on the edgewise response of floating offshore wind turbines”. In: *Coupled Syst. Mech* 2.3 (2013), pp. 231–253.
- [26] Fitzgerald, Breiffni and Basu, Biswajit. “Structural control of wind turbines with soil structure interaction included”. In: *Engineering Structures* 111 (2016), pp. 131–151.
- [27] Harte, M, Basu, B, and Nielsen, Søren RK. “Dynamic analysis of wind turbines including soil-structure interaction”. In: *Engineering structures* 45 (2012), pp. 509–518.

-
- [28] Waris, Muhammad Bilal and Ishihara, Takeshi. “Dynamic response analysis of floating offshore wind turbine with different types of heave plates and mooring systems by using a fully nonlinear model”. In: *Coupled systems mechanics* 1.3 (2012), pp. 247–268.
 - [29] Wang, Kai, Moan, Torgeir, and Hansen, Martin OL. “Stochastic dynamic response analysis of a floating vertical-axis wind turbine with a semi-submersible floater”. In: *Wind Energy* 19.10 (2016), pp. 1853–1870.
 - [30] Baishthakur, Shubham, Pakrashi, Vikram, Bhattacharya, Subhamoy, and Fitzgerald, Breiffni. “Impact of Sources of Damping On the Fragility Estimates of Wind Turbine Towers”. In: *ASCE-ASME J Risk and Uncert in Engrg Sys Part B Mech Engrg* (2024), pp. 1–47.
 - [31] Bathe, Klaus-Jürgen. *Finite element procedures*. Klaus-Jurgen Bathe, 2006.
 - [32] Sheibani, Mohammad and Akbari, Ali Akbar. “Finite element modeling of a wind turbine blade”. In: *Journal of Vibroengineering* 17.7 (2015), pp. 3774–3791.
 - [33] Navadeh, Navid, Goroshko, Ivan, Zhuk, Yaroslav, Etminan Moghadam, Farnoosh, and Soleiman Fallah, Arash. “Finite element analysis of wind turbine blade vibrations”. In: *Vibration* 4.2 (2021), pp. 310–322.
 - [34] Gentils, Theo, Wang, Lin, and Kolios, Athanasios. “Integrated structural optimisation of offshore wind turbine support structures based on finite element analysis and genetic algorithm”. In: *Applied energy* 199 (2017), pp. 187–204.
 - [35] Ali, Ahmer, De Risi, Raffaele, and Sextos, Anastasios. “Finite element modeling optimization of wind turbine blades from an earthquake engineering perspective”. In: *Engineering Structures* 222 (2020), p. 111105.
 - [36] Boudounit, Hicham, Tarfaoui, Mostapha, and Saifaoui, Dennoun. “Fatigue analysis of wind turbine composite blade using finite element method”. In: *Wind Engineering* 47.3 (2023), pp. 706–721.
 - [37] Jonkman, Bonnie et al. *OpenFAST/openfast: v3.5.2*. Version v3.5.2. 2024. DOI: 10.5281/zenodo.10530537. URL: <https://doi.org/10.5281/zenodo.10530537>.
 - [38] *Bladed Theory Manual*. Version 4.13. DNV Services UK Limited. Bristol, UK, 2022.
 - [39] Larsen, Torben Juul and Hansen, Anders M. “How 2 HAWC2, the user’s manual”. In: *RisøReport, Risø* 702 (2009).
 - [40] Sarkar, Saptarshi and Fitzgerald, Breiffni. “Vibration control of spar-type floating offshore wind turbine towers using a tuned mass-damper-inerter”. In: *Structural Control and Health Monitoring* 27.1 (2020), e2471.
 - [41] Abbas, Nikhar, Zalkind, Daniel, Pao, Lucy, and Wright, Alan. “A reference open-source controller for fixed and floating offshore wind turbines”. In: *Wind Energy Science Discussions* 2021 (2021), pp. 1–33.
 - [42] Johlas, Hannah M, Martínez-Tossas, Luis A, Churchfield, Matthew J, Lackner, Matthew A, and Schmidt, David P. “Floating platform effects on power generation in spar and semisubmersible wind turbines”. In: *Wind Energy* 24.8 (2021), pp. 901–916.
 - [43] Amaral, R et al. “A frequency-time domain method for annual energy production estimation in floating wind turbines”. In: *Journal of Physics: Conference Series*. Vol. 2265. 4. IOP Publishing. 2022, p. 042025.

- [44] IRENA, International Renewable Energy Agency. *Renewable power generation costs in 2020*. eBook Partnership, 2022.
- [45] Sarkar, Saptarshi, Fitzgerald, Breiffni, and Basu, Biswajit. “Individual blade pitch control of floating offshore wind turbines for load mitigation and power regulation”. In: *IEEE Transactions on Control Systems Technology* 29.1 (2020), pp. 305–315.
- [46] Sarkar, Saptarshi and Fitzgerald, Breiffni. “Fluid inerter for optimal vibration control of floating offshore wind turbine towers”. In: *Engineering Structures* 266 (2022), p. 114558.
- [47] Njiri, Jackson G and Söffker, Dirk. “State-of-the-art in wind turbine control: Trends and challenges”. In: *Renewable and Sustainable Energy Reviews* 60 (2016), pp. 377–393.
- [48] Tan, Jian Ding, Chang, Clifford Choe Wei, Bhuiyan, Mohammad Arif Sobhan, Nisa’Minhad, Khairun, and Ali, Kharudin. “Advancements of wind energy conversion systems for low-wind urban environments: A review”. In: *Energy Reports* 8 (2022), pp. 3406–3414.
- [49] Fitzgerald, Breiffni, McAuliffe, James, Baisthakur, Shubham, and Sarkar, Saptarshi. “Enhancing the reliability of floating offshore wind turbine towers subjected to misaligned wind-wave loading using tuned mass damper inerters (TMDIs)”. In: *Renewable Energy* 211 (2023), pp. 522–538.
- [50] Pinto, Paolo Emilio, Giannini, Renato, Franchin, Paolo, et al. *Seismic reliability analysis of structures*. IUSS press, 2004.
- [51] Melchers, Robert E and Beck, André T. *Structural reliability analysis and prediction*. John wiley & sons, 2018.
- [52] Haldar, Achintya and Mahadevan, Sankaran. *Probability, reliability and statistical methods in engineering design*. Wiley, 2000. URL: <https://cir.nii.ac.jp/crid/1130282272608504576>.
- [53] Straub, Daniel and Papaioannou, Iason. “Bayesian updating with structural reliability methods”. In: *Journal of Engineering Mechanics* 141.3 (2015), p. 04014134.
- [54] Goswami, Somdatta, Ghosh, Shyamal, and Chakraborty, Subrata. “Reliability analysis of structures by iterative improved response surface method”. In: *Structural Safety* 60 (2016), pp. 56–66.
- [55] Sørensen, John Dalsgaard, Kroon, IB, and Faber, Michael Havbro. “Optimal reliability-based code calibration”. In: *Structural Safety* 15.3 (1994), pp. 197–208.
- [56] Veers, PS. *Fatigue reliability of wind turbine components*. Tech. rep. Sandia National Labs., Albuquerque, NM (USA), 1990.
- [57] Lange, Clifford H. *Probabilistic fatigue methodology and wind turbine reliability*. Stanford University, 1996.
- [58] Winterstein, Steven R and Lange, Clifford H. “Moment-based probability models for wind engineering applications”. In: *Engineering Mechanics*. ASCE. 1995, pp. 159–162.
- [59] Lange, CH and Winterstein, SR. *Fatigue design of wind turbine blades: load and resistance factors from limited data*. Tech. rep. American Society of Mechanical Engineers, New York, NY (United States), 1996.

-
- [60] Veers, Paul S and Winterstein, Steven R. "Application of measured loads to wind turbine fatigue and reliability analysis". In: (1998).
 - [61] Winterstein, Steven R and Kashef, Tina. "Moment-based load and response models with wind engineering applications". In: *J. Sol. Energy Eng.* 122.3 (2000), pp. 122–128.
 - [62] Manuel, Lance, Veers, Paul S, and Winterstein, Steven R. "Parametric models for estimating wind turbine fatigue loads for design". In: *J. Sol. Energy Eng.* 123.4 (2001), pp. 346–355.
 - [63] Fitzwater, LeRoy M and Winterstein, Steven R. "Predicting design wind turbine loads from limited data: Comparing random process and random peak models". In: *J. Sol. Energy Eng.* 123.4 (2001), pp. 364–371.
 - [64] Seebregts, AJ, Rademakers, LWMM, and Van Den Horn, BA. "Reliability analysis in wind turbine engineering". In: *Microelectronics Reliability* 35.9-10 (1995), pp. 1285–1307.
 - [65] Rademakers, L WMM, Seebregts, AJ, Van den Horn, BA, Jehee, J NT, and Blok, BM. "Methodology for probabilistic safety assessment of wind turbines. Demonstrated by a case study of the Lagerwey LW 15/75 design". In: (1993).
 - [66] Rademakers, LWMM et al. "Assessment of Wind Turbine Safety". In: *Report, EUREC-Agency* (1995).
 - [67] Kongsø, Hans Erik, Kozin, Igor O, and Christensen, Palle. "The EFP-94 project on safety systems for wind turbines. Method for evaluation of failure modes and reliability". In: (1996).
 - [68] Ronold, Knut O, Wedel-Heinen, Jakob, and Christensen, Carl J. "Reliability-based fatigue design of wind-turbine rotor blades". In: *Engineering structures* 21.12 (1999), pp. 1101–1114.
 - [69] Michos, D, Dialynas, E, and Vionis, P. "Reliability and safety assessment of wind turbines control and protection systems". In: *Wind Engineering* 26.6 (2002), pp. 359–369.
 - [70] Khan, Muhammad Mohsin K. "Reliability analysis and condition monitoring of a horizontal axis wind turbine". PhD thesis. Memorial University of Newfoundland, 2005.
 - [71] Muskulus, Michael and Schafhirt, Sebastian. "Reliability-based design of wind turbine support structures". In: *Proceedings of the symposium on reliability of engineering system, hangzhou, china*. Vol. 1517. 2015.
 - [72] Sichani, Mahdi Teimouri, Nielsen, SR, and Naess, Arvid. "Failure probability estimation of wind turbines by enhanced Monte Carlo method". In: *Journal of engineering mechanics* 138.4 (2012), pp. 379–389.
 - [73] Naess, A and Gaidai, O. "Monte Carlo methods for estimating the extreme response of dynamical systems". In: *Journal of Engineering Mechanics* 134.8 (2008), pp. 628–636.
 - [74] Dong, Wenbin, Moan, Torgeir, and Gao, Zhen. "Fatigue reliability analysis of the jacket support structure for offshore wind turbine considering the effect of corrosion and inspection". In: *Reliability Engineering & System Safety* 106 (2012), pp. 11–27.

- [75] Dong, Wenbin, Moan, Torgeir, and Gao, Zhen. “Long-term fatigue analysis of multi-planar tubular joints for jacket-type offshore wind turbine in time domain”. In: *Engineering Structures* 33.6 (2011), pp. 2002–2014.
- [76] Tarp-Johansen, N J. *Examples of fatigue lifetime and reliability evaluation of larger wind turbine components*. Mar. 2003.
- [77] Hsu, Yao, Wu, Wen-Fang, and Chang, Yung-Chang. “Reliability analysis of wind turbine towers”. In: *Procedia Engineering* 79 (2014), pp. 218–224.
- [78] Bacharoudis, Konstantinos C and Philippidis, Theodore P. “A probabilistic approach for strength and stability evaluation of wind turbine rotor blades in ultimate loading”. In: *Structural Safety* 40 (2013), pp. 31–38.
- [79] IEC. *IEC 61400-1: Wind energy generation systems - Part 1: Design requirements*. 2019.
- [80] Mardfekri, Maryam and Gardoni, Paolo. “Multi-hazard reliability assessment of offshore wind turbines”. In: *Wind Energy* 18.8 (2015), pp. 1433–1450.
- [81] Hu, Weifei, Choi, KK, and Cho, Hyunkyoo. “Reliability-based design optimization of wind turbine blades for fatigue life under dynamic wind load uncertainty”. In: *Structural and multidisciplinary optimization* 54 (2016), pp. 953–970.
- [82] *Standard Practices for Cycle Counting in Fatigue Analysis*. West Conshohocken, PA: ASTM International, 2017.
- [83] Zwick, Daniel and Muskulus, Michael. “The simulation error caused by input loading variability in offshore wind turbine structural analysis”. In: *Wind energy* 18.8 (2015), pp. 1421–1432.
- [84] Häfele, Jan, Hübler, Clemens, Gebhardt, Cristian Guillermo, and Rolfes, Raimund. “A comprehensive fatigue load set reduction study for offshore wind turbines with jacket substructures”. In: *Renewable energy* 118 (2018), pp. 99–112.
- [85] Barone, Matthew, Paquette, Joshua, Resor, Brian, and Manuel, Lance. “Decades of Wind Turbine Load Simulation”. In: *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*. 2012, p. 1288.
- [86] Graf, Peter A, Stewart, Gordon, Lackner, Matthew, Dykes, Katherine, and Veers, Paul. “High-throughput computation and the applicability of Monte Carlo integration in fatigue load estimation of floating offshore wind turbines”. In: *Wind Energy* 19.5 (2016), pp. 861–872.
- [87] Liao, Ding et al. “Fatigue reliability of wind turbines: historical perspectives, recent developments and future prospects”. In: *Renewable Energy* 200 (2022), pp. 724–742.
- [88] Patel, K and Freathy, P. “A simplified method for assessing wind-induced fatigue damage”. In: *Engineering Structures* 6.4 (1984), pp. 268–273.
- [89] Kashef, Tina and Winterstein, Steven R. “Relating turbulence to wind turbine blade loads: Parametric study with multiple regression analysis”. In: (1999).
- [90] Veldkamp, Herman Frederik. “Chances in wind energy: a probalistic approach to wind turbine fatigue design”. In: (2006).

-
- [91] Benasciutti, Denis and Tovo, Roberto. “Spectral methods for lifetime prediction under wide-band stationary random processes”. In: *International Journal of fatigue* 27.8 (2005), pp. 867–877.
 - [92] Yeter, B, Garbatov, Y, and Guedes Soares, C. “Spectral fatigue assessment of an offshore wind turbine structure under wave and wind loading”. In: *Developments in Maritime Transportation and Exploitation of Sea Resources*. London, UK: Taylor & Francis Group (2014), pp. 425–433.
 - [93] Zwick, Daniel and Muskulus, Michael. “Simplified fatigue load assessment in offshore wind turbine structural analysis”. In: *Wind Energy* 19.2 (2016), pp. 265–278.
 - [94] Robertson, Amy N, Shaler, Kelsey, Sethuraman, Latha, and Jonkman, Jason. *Sensitivity of uncertainty in wind characteristics and wind turbine properties on wind turbine extreme and fatigue loads*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2019.
 - [95] Toft, Henrik Stensgaard, Svenningsen, Lasse, Moser, Wolfgang, Sørensen, John Dalsgaard, and Thøgersen, Morten Lybech. “Assessment of wind turbine structural integrity using response surface methodology”. In: *Engineering Structures* 106 (2016), pp. 471–483.
 - [96] Vera-Tudela, Luis and Kühn, Martin. “Analysing wind turbine fatigue load prediction: The impact of wind farm flow conditions”. In: *Renewable Energy* 107 (2017), pp. 352–360.
 - [97] Kulkarni, Pravin A, Dhoble, Ashwinkumar S, and Padole, Pramod M. “Deep neural network-based wind speed forecasting and fatigue analysis of a large composite wind turbine blade”. In: *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* 233.8 (2019), pp. 2794–2812.
 - [98] Luna, Julio, Falkenberg, Ole, Gros, Sébastien, and Schild, Axel. “Wind turbine fatigue reduction based on economic-tracking NMPC with direct ANN fatigue estimation”. In: *Renewable Energy* 147 (2020), pp. 1632–1641.
 - [99] Li, Xuan and Zhang, Wei. “Long-term fatigue damage assessment for a floating offshore wind turbine under realistic environmental conditions”. In: *Renewable energy* 159 (2020), pp. 570–584.
 - [100] Bai, Hao, Shi, Lujie, Aoues, Younes, Huang, Changwu, and Lemosse, Didier. “Estimation of probability distribution of long-term fatigue damage on wind turbine tower using residual neural network”. In: *Mechanical Systems and Signal Processing* 190 (2023), p. 110101.
 - [101] Papi, Francesco, Jonkman, Jason, Robertson, Amy, and Bianchini, Alessandro. “Going Beyond BEM with BEM: an Insight into Dynamic Inflow Effects on Floating Wind Turbines”. In: *Wind Energy Science Discussions* 2023 (2023), pp. 1–29.
 - [102] Dimitrov, Nikolay, Kelly, Mark C, Vignaroli, Andrea, and Berg, Jacob. “From wind to loads: wind turbine site-specific load estimation with surrogate models trained on high-fidelity load databases”. In: *Wind Energy Science* 3.2 (2018), pp. 767–790.
 - [103] Ren, Chao. “Reliability assessment of an offshore wind turbine jacket with active learning approaches”. PhD thesis. Normandie Université, 2022.

- [104] Müller, Kolja, Dazer, Martin, and Cheng, Po Wen. “Damage assessment of floating offshore wind turbines using response surface modeling”. In: *Energy Procedia* 137 (2017), pp. 119–133.
- [105] Rinker, Jennifer M. “Calculating the sensitivity of wind turbine loads to wind inputs using response surfaces”. In: *Journal of physics: Conference series*. Vol. 753. 3. IOP Publishing. 2016, p. 032057.
- [106] Lee, Sang-Lae and Shin, SangJoon. “Wind turbine blade optimal design considering multi-parameters and response surface method”. In: *Energies* 13.7 (2020), p. 1639.
- [107] Li, JY, Li, R, Gao, Y, and Huang, J. “Aerodynamic optimization of wind turbine airfoils using response surface techniques”. In: *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy* 224.6 (2010), pp. 827–838.
- [108] Meng, Debiao, Yang, Shiyuan, Jesus, Abílio MP de, and Zhu, Shun-Peng. “A novel Kriging-model-assisted reliability-based multidisciplinary design optimization strategy and its application in the offshore wind turbine tower”. In: *Renewable Energy* 203 (2023), pp. 407–420.
- [109] Teixeira, Rui, O’Connor, Alan, Nogal, Maria, Nichols, James, and Spring, Mark. “Structural probabilistic assessment of Offshore Wind Turbine operation fatigue based on Kriging interpolation”. In: *Proceedings of 27th European Safety and Reliability Conference (ESREL 2017)*. 2017, pp. 18–22.
- [110] Teixeira, Rui, O’Connor, Alan, Nogal, Maria, Krishnan, Nandakumar, and Nichols, James. “Analysis of the design of experiments of offshore wind turbine fatigue reliability design with Kriging surfaces”. In: *Procedia Structural Integrity* 5 (2017), pp. 951–958.
- [111] Teixeira, Rui, Nogal, Maria, O’Connor, Alan, Nichols, James, and Dumas, Antoine. “Stress-cycle fatigue design with Kriging applied to offshore wind turbines”. In: *International Journal of Fatigue* 125 (2019), pp. 454–467.
- [112] Morató, A, Sriramula, S, and Krishnan, N. “Kriging models for aero-elastic simulations and reliability analysis of offshore wind turbine support structures”. In: *Ships and Offshore Structures* 14.6 (2019), pp. 545–558.
- [113] Ling, Le, Li, Yan, and Fu, Sicheng. “A reliability analysis strategy for main shaft of wind turbine using importance sampling and Kriging model”. In: *International Journal of Structural Integrity* 13.2 (2022), pp. 297–308.
- [114] Keprate, Arvind, Bagalkot, Nikhil, Siddiqui, Muhammad Salman, and Sen, Subhamoy. “Reliability analysis of 15MW horizontal axis wind turbine rotor blades using fluid-structure interaction simulation and adaptive kriging model”. In: *Ocean Engineering* 288 (2023), p. 116138.
- [115] Zuhal, Lavi R, Faza, Ghifari A, Palar, Pramudita S, and Shimoyama, Koji. “Multi-objective kriging-based optimization for high-fidelity wind turbine design”. In: *AIAA Scitech 2019 Forum*. 2019, p. 0539.
- [116] Zheng, Yuqiao, Ma, Huidong, Wei, Jianfeng, and Zhu, Kai. “Robust optimization for composite blade of wind turbine based on kriging model”. In: *Advanced composites letters* 29 (2020), p. 2633366X20914631.

- [117] Pholdee, Nantiwat, Bureerat, Sujin, and Nuantong, Weerapon. “Kriging surrogate-based genetic algorithm optimization for blade design of a horizontal axis wind turbine”. In: *Computer Modeling in Engineering & Sciences* 126.1 (2021), pp. 261–273.
- [118] Beyaoui, Moez et al. “Dynamic behaviour of a wind turbine gear system with uncertainties”. In: *Comptes Rendus. Mécanique* 344.6 (2016), pp. 375–387.
- [119] Murcia, Juan Pablo et al. “Uncertainty propagation through an aeroelastic wind turbine model using polynomial surrogates”. In: *Renewable Energy* 119 (2018), pp. 910–922.
- [120] Mylonas, Charilaos, Abdallah, Imad, and Chatzi, Eleni. “Surrogate modelling for fatigue damage of wind-turbine blades using polynomial chaos expansions and non-negative matrix factorization”. In: *IABSE Symposium 2017: Engineering the Future*. Vol. 109. International Association for Bridge and Structural Engineering (IABSE). 2017, pp. 809–816.
- [121] Ganesh, N and Gupta, Sayan. “Estimating the Rain-Flow Fatigue Damage in Wind Turbine Blades Using Polynomial Chaos”. In: *Proceedings of the International Symposium on Engineering under Uncertainty: Safety Assessment and Management (ISEUSAM-2012)*. Springer. 2013, pp. 711–727.
- [122] Thapa, Mishal and Missoum, Samy. “Uncertainty quantification and global sensitivity analysis of composite wind turbine blades”. In: *Reliability Engineering & System Safety* 222 (2022), p. 108354.
- [123] Moura Paredes, Guilherme, Eskilsson, Claes, and P. Engsig-Karup, Allan. “Uncertainty quantification in mooring cable dynamics using polynomial chaos expansions”. In: *Journal of Marine Science and Engineering* 8.3 (2020), p. 162.
- [124] Paredes, Guilherme Moura, Thomsen, Jonas Bjerg, Ferri, Francesco, and Eskilsson, Claes. “Mooring system reliability analysis of an ORE device using general Polynomial Chaos”. In: *13th European Wave and Tidal Energy Conference*. European Tidal and Wave Energy Conference. 2019, pp. 1271–1.
- [125] Haghi, Rad and Crawford, Curran. “Surrogate models for the blade element momentum aerodynamic model using non-intrusive polynomial chaos expansions”. In: *Wind Energy Science Discussions* 2021 (2021), pp. 1–20.
- [126] Marugán, Alberto Pliego, Márquez, Fausto Pedro García, Perez, Jesus María Pinar, and Ruiz-Hernández, Diego. “A survey of artificial neural network in wind energy systems”. In: *Applied energy* 228 (2018), pp. 1822–1836.
- [127] Wang, Yun, Zou, Runmin, Liu, Fang, Zhang, Lingjun, and Liu, Qianyi. “A review of wind speed and wind power forecasting with deep neural networks”. In: *Applied Energy* 304 (2021), p. 117766.
- [128] Rai, Man Mohan and Madavan, Nateri K. “Aerodynamic design using neural networks”. In: *AIAA journal* 38.1 (2000), pp. 173–182.
- [129] Cinnella, Paola and Congedo, PM. “Optimal airfoil shapes for viscous transonic flows of Bethe–Zel’dovich–Thompson fluids”. In: *Computers & Fluids* 37.3 (2008), pp. 250–264.
- [130] Giannakoglou, KC. “Design of optimal aerodynamic shapes using stochastic optimization methods and computational intelligence”. In: *Progress in Aerospace sciences* 38.1 (2002), pp. 43–76.

- [131] Duvigneau, Régis and Visonneau, Michel. “Hybrid genetic algorithms and artificial neural networks for complex design optimization in CFD”. In: *International journal for numerical methods in fluids* 44.11 (2004), pp. 1257–1278.
- [132] Mortazavi, Seyed Mehdi, Soltani, Mohammad Reza, and Motieyan, Hamid. “A Pareto optimal multi-objective optimization for a horizontal axis wind turbine blade airfoil sections utilizing exergy analysis and neural networks”. In: *Journal of Wind Engineering and Industrial Aerodynamics* 136 (2015), pp. 62–72.
- [133] Ribeiro, AFP, Awruch, AM, and Gomes, HM. “An airfoil optimization technique for wind turbines”. In: *Applied Mathematical Modelling* 36.10 (2012), pp. 4898–4907.
- [134] Schröder, Laura, Dimitrov, Nikolay Krasimirov, Verelst, David Robert, and Sørensen, John Aasted. “Wind turbine site-specific load estimation using artificial neural networks calibrated by means of high-fidelity load simulations”. In: *Journal of Physics: Conference Series*. Vol. 1037. 6. IOP Publishing. 2018, p. 062027.
- [135] Schröder, Laura, Dimitrov, Nikolay Krasimirov, and Sørensen, John Aasted. “Uncertainty propagation and sensitivity analysis of an artificial neural network used as wind turbine load surrogate model”. In: *Journal of Physics: Conference Series*. Vol. 1618. 4. IOP Publishing. 2020, p. 042040.
- [136] Raissi, Maziar, Perdikaris, Paris, and Karniadakis, George E. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational physics* 378 (2019), pp. 686–707.
- [137] Sarkar, Saptarshi and Fitzgerald, Breiffni. “Use of kane’s method for multi-body dynamic modelling and control of spar-type floating offshore wind turbines”. In: *Energies* 14.20 (2021), p. 6635.
- [138] Jonkman, Jason Mark. *Modeling of the UAE Wind Turbine for Refinement of FAST { } AD*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2003.
- [139] Kane, Thomas R and Levinson, David A. *Dynamics, theory and applications*. McGraw Hill, 1985.
- [140] Jonkman, Bonnie J. *TurbSim user’s guide: Version 1.50*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2009.
- [141] Gaertner, Evan 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.
- [142] Fluck, Manuel and Crawford, Curran. “A fast stochastic solution method for the Blade Element Momentum equations for long-term load assessment”. In: *Wind Energy* 21.2 (2018), pp. 115–128.
- [143] Haghi, Rad and Crawford, Curran. “Surrogate models for the blade element momentum aerodynamic model using non-intrusive polynomial chaos expansions”. In: *Wind Energy Science* 7.3 (2022), pp. 1289–1304.
- [144] Manwell, James F, McGowan, Jon G, and Rogers, Anthony L. *Wind energy explained: theory, design and application*. John Wiley & Sons, 2010.

-
- [145] Glauert, H. *Aerodynamic Theory: A General Review of Progress, volume IV, chapter Division L, Airplane Propellers*. 1963.
 - [146] Buhl Jr, Marshall L. *New empirical relationship between thrust coefficient and induction factor for the turbulent windmill state*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2005.
 - [147] Ning, S Andrew. “A simple solution method for the blade element momentum equations with guaranteed convergence”. In: *Wind Energy* 17.9 (2014), pp. 1327–1345.
 - [148] Ning, Andrew, Hayman, Gregory, Damiani, Rick, and Jonkman, Jason M. “Development and validation of a new blade element momentum skewed-wake model within AeroDyn”. In: *33rd Wind Energy Symposium*. 2015, p. 0215.
 - [149] Brent, Richard P. “An algorithm with guaranteed convergence for finding a zero of a function”. In: *The computer journal* 14.4 (1971), pp. 422–425.
 - [150] Murphy, Kevin P. *Machine learning: a probabilistic perspective*. MIT press, 2012.
 - [151] Kingma, Diederik P and Ba, Jimmy. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
 - [152] Abbas, Nikhar J, Zalkind, Daniel S, Pao, Lucy, and Wright, Alan. “A reference open-source controller for fixed and floating offshore wind turbines”. In: *Wind Energy Science* 7.1 (2022), pp. 53–73.
 - [153] Wang, Sifan, Teng, Yujun, and Perdikaris, Paris. “Understanding and mitigating gradient flow pathologies in physics-informed neural networks”. In: *SIAM Journal on Scientific Computing* 43.5 (2021), A3055–A3081.
 - [154] Monaco, Simone and Apletti, Daniele. “Training physics-informed neural networks: One learning to rule them all?” In: *Results in Engineering* 18 (2023), p. 101023.
 - [155] Singh, Deepali, Dwight, Richard P, Laugesen, Kasper, Beaudet, Laurent, and Viré, Axelle. “Probabilistic surrogate modeling of offshore wind-turbine loads with chained Gaussian processes”. In: *Journal of Physics: Conference Series*. Vol. 2265. 3. IOP Publishing. 2022, p. 032070.
 - [156] Huang, Bin and Wang, Jianhui. “Applications of physics-informed neural networks in power systems-a review”. In: *IEEE Transactions on Power Systems* 38.1 (2022), pp. 572–588.
 - [157] Cai, Shengze, Mao, Zhiping, Wang, Zhicheng, Yin, Minglang, and Karniadakis, George Em. “Physics-informed neural networks (PINNs) for fluid mechanics: A review”. In: *Acta Mechanica Sinica* 37.12 (2021), pp. 1727–1738.
 - [158] Antonelo, Eric Aislan et al. “Physics-informed neural nets for control of dynamical systems”. In: *Neurocomputing* 579 (2024), p. 127419.
 - [159] Cai, Shengze, Wang, Zhicheng, Wang, Sifan, Perdikaris, Paris, and Karniadakis, George Em. “Physics-informed neural networks for heat transfer problems”. In: *Journal of Heat Transfer* 143.6 (2021), p. 060801.
 - [160] Behera, Ranjan Kumar, Jena, Monalisa, Rath, Santanu Kumar, and Misra, Sanjay. “Co-LSTM: Convolutional LSTM model for sentiment analysis in social big data”. In: *Information Processing & Management* 58.1 (2021), p. 102435.

- [161] Bappy, Jawadul H, Simons, Cody, Nataraj, Lakshmanan, Manjunath, BS, and Roy-Chowdhury, Amit K. “Hybrid lstm and encoder–decoder architecture for detection of image forgeries”. In: *IEEE transactions on image processing* 28.7 (2019), pp. 3286–3300.
- [162] Graves, Alex, Jaitly, Navdeep, and Mohamed, Abdel-rahman. “Hybrid speech recognition with deep bidirectional LSTM”. In: *2013 IEEE workshop on automatic speech recognition and understanding*. IEEE. 2013, pp. 273–278.
- [163] Banik, Abhishek, Behera, Chinmaya, Sarathkumar, Tirunagaru V, and Goswami, Arup Kumar. “Uncertain wind power forecasting using LSTM-based prediction interval”. In: *IET Renewable Power Generation* 14.14 (2020), pp. 2657–2667.
- [164] Yu, Ruiguo et al. “LSTM-EFG for wind power forecasting based on sequential correlation features”. In: *Future Generation Computer Systems* 93 (2019), pp. 33–42.
- [165] Woo, Seongcheol, Park, Junyoung, and Park, Jinkyoo. “Predicting wind turbine power and load outputs by multi-task convolutional LSTM model”. In: *2018 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE. 2018, pp. 1–5.
- [166] Choe, Do-Eun, Kim, Hyoung-Chul, and Kim, Moo-Hyun. “Sequence-based modeling of deep learning with LSTM and GRU networks for structural damage detection of floating offshore wind turbine blades”. In: *Renewable Energy* 174 (2021), pp. 218–235.
- [167] Xiang, Ling, Wang, Penghe, Yang, Xin, Hu, Aijun, and Su, Hao. “Fault detection of wind turbine based on SCADA data analysis using CNN and LSTM with attention mechanism”. In: *Measurement* 175 (2021), p. 109094.
- [168] Chen, Hansi, Liu, Hang, Chu, Xuening, Liu, Qingxiu, and Xue, Deyi. “Anomaly detection and critical SCADA parameters identification for wind turbines based on LSTM-AE neural network”. In: *Renewable Energy* 172 (2021), pp. 829–840.
- [169] Liew, Jaime, Lio, Wai Hou, Urbán, Albert Meseguer, Holierhoek, Jessica, and Kim, Taeseong. “Active tip deflection control for wind turbines”. In: *Renewable Energy* 149 (2020), pp. 445–454.
- [170] Dimitrov, Nikolay and Göçmen, Tuhfe. “Virtual sensors for wind turbines with machine learning-based time series models”. In: *Wind Energy* 25.9 (2022), pp. 1626–1645.
- [171] Rumelhart, David E, Hinton, Geoffrey E, and Williams, Ronald J. “Learning representations by back-propagating errors”. In: *nature* 323.6088 (1986), pp. 533–536.
- [172] Hu, Yuhuang, Huber, Adrian, Anumula, Jithendar, and Liu, Shih-Chii. “Overcoming the vanishing gradient problem in plain recurrent networks”. In: *arXiv preprint arXiv:1801.06105* (2018).
- [173] Hochreiter, Sepp and Schmidhuber, Jürgen. “Long short-term memory”. In: *Neural computation* 9.8 (1997), pp. 1735–1780.
- [174] Calazone, Ottavio. *An Intuitive Explanation of LSTM*. Accessed on February 20, 2024. 2022. URL: <https://medium.com/@ottaviocalzone/an-intuitive-explanation-of-lstm-a035eb6ab42c>.

-
- [175] Wold, Svante, Esbensen, Kim, and Geladi, Paul. "Principal component analysis". In: *Chemometrics and intelligent laboratory systems* 2.1-3 (1987), pp. 37–52.
 - [176] Fisher, Ronald A. "The use of multiple measurements in taxonomic problems". In: *Annals of eugenics* 7.2 (1936), pp. 179–188.
 - [177] Kramer, Mark A. "Nonlinear principal component analysis using autoassociative neural networks". In: *AIChE journal* 37.2 (1991), pp. 233–243.
 - [178] Simley, Eric, Pao, Lucy, Kelley, Neil, Jonkman, Bonnie, and Frehlich, Rod. "Lidar wind speed measurements of evolving wind fields". In: *50th AIAA aerospace sciences meeting including the new horizons forum and aerospace exposition*. 2012, p. 656.
 - [179] Markaki, Maria E and Stylianou, Yannis. "Dimensionality reduction of modulation frequency features for speech discrimination." In: *INTERSPEECH*. 2008, pp. 646–649.
 - [180] Van Der Maaten, Laurens, Postma, Eric O, Herik, H Jaap van den, et al. "Dimensionality reduction: A comparative review". In: *Journal of Machine Learning Research* 10.66-71 (2009), p. 13.
 - [181] Hamarneh, Ghassan, McIntosh, Chris, and Drew, Mark S. "Perception-based visualization of manifold-valued medical images using distance-preserving dimensionality reduction". In: *IEEE transactions on medical imaging* 30.7 (2011), pp. 1314–1327.
 - [182] Lataniotis, Christos. "Data-driven uncertainty quantification for high-dimensional engineering problems". PhD thesis. ETH Zurich, 2019.
 - [183] Skittides, Christina and Früh, Wolf-Gerrit. "Wind forecasting using principal component analysis". In: *Renewable Energy* 69 (2014), pp. 365–374.
 - [184] Geng, Dawei, Zhang, Haifeng, and Wu, Hongyu. "Short-term wind speed prediction based on principal component analysis and LSTM". In: *Applied sciences* 10.13 (2020), p. 4416.
 - [185] Zhang, Kai, Tang, Baoping, Deng, Lei, and Yu, Xiaoxia. "Fault detection of wind turbines by subspace reconstruction-based robust kernel principal component analysis". In: *IEEE Transactions on Instrumentation and Measurement* 70 (2021), pp. 1–11.
 - [186] Wang, Yifei, Ma, Xiandong, and Joyce, Malcolm J. "Reducing sensor complexity for monitoring wind turbine performance using principal component analysis". In: *Renewable energy* 97 (2016), pp. 444–456.
 - [187] Schär, Styfen, Marelli, Stefano, and Sudret, Bruno. "Emulating the dynamics of complex systems using autoregressive models on manifolds (mNARX)". In: *Mechanical Systems and Signal Processing* 208 (2024), p. 110956.
 - [188] Hübner, Clemens, Gebhardt, Cristian G, and Rolfes, Raimund. "Methodologies for fatigue assessment of offshore wind turbines considering scattering environmental conditions and the uncertainty due to finite sampling". In: *Wind Energy* 21.11 (2018), pp. 1092–1105.
 - [189] Thrun, Sebastian and Mitchell, Tom M. "Learning one more thing". In: *IJ-CAI*. Vol. 95. 1995, pp. 1217–1223.
 - [190] Caruana, Rich. "Multitask learning". In: *Machine learning* 28 (1997), pp. 41–75.