Lifetime performance of wind turbines and wind farms

Jack Kelly

St. Edmund Hall
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2022

Abstract

Wind energy is rapidly expanding as the demand for a low-emissions electricity grid increases. Acknowledging this, and the intermittent nature of the wind resource, both short-term (circa hours) and long-term (circa years) planning must be implemented. To do this accurately, the performance degradation of wind farms must be quantified and accounted for in electricity systems and financial models. This thesis is concerned with quantifying aerodynamic performance decline of wind turbines and wind farms to better understand changes in lifetime performance and more accurately predict wind farm power output.

Wind turbine aerodynamic performance losses are investigated by analysing the effect of roughness on airfoil performance (lift and drag). One of the challenges in quantifying wind turbine performance decline due to roughness is the lack of roughness airfoil data available. A 'roughness evolution parameter' is therefore parameterised based on publicly available roughened airfoil experimental tables, which enables roughened airfoil data to be estimated from clean airfoil data. The calculation of the parameter is proposed as a framework, with room for further refinement when more experimental data comes available.

Blade element momentum theory is used to analyse the impact of blade roughness on wind turbine performance. Exemplar wind turbine rotors are designed and the impact of roughness on annual energy production is quantified to be in the range of 2.9–8.6% for realistic levels of roughness. A novel control method is proposed to mitigate losses due to roughness and 0.1–1.0% of annual energy production is recovered.

To account for the long-term variability in wind speed, wind farm performance decline is quantified using long-term corrected capacity factor data from wind farms in the UK, US and Australia – 0.26, 0.72 and 0.43 capacity factor percentage points per year respectively. To do so, power output data is combined with reanalysis wind speed data to correct for variability and biases in the data, allowing the underlying performance trend to be realised. To address shortcomings in the data (e.g. undetected downtime or reanalysis errors) a novel long-term correction approach is used which reduces variability in the results.

Reanalysis wind speed data and wind farm datasets are utilised for a supervised machine learning performance prediction model. Due to the inclusion of additional metadata, this model is able to predict wind farm capacity factor output with a 56–90% improvement in mean absolute error value when compared to predicting power output using a wind farm power curve.

Lifetime performance of wind turbines and wind farms



Jack Kelly

St. Edmund Hall

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2022

Acknowledgements

I would like to thank Dr. Christopher Vogel and Prof. Richard Willden for their support, advice, wisdom and encouragement throughout my DPhil. They provided excellent academic guidance as well as providing a stable base to help me navigate all the hurdles that were thrown my way (part-time to full-time, funding shortfalls, studying internationally, COVID-19). Their flexibility through many uncertain times cannot go without praise, it was truly what allowed me to complete this degree. Beyond this, they allowed me to chase my academic interests and also balance other interests to ensure I enjoyed my time at Oxford. I would especially like to emphasise Dr. Vogel's influence and dedication, which even included a trip to Sydney! I would also like to thank my fellow students for all their advice and insightful conversations that helped along the way.

I am grateful to The Australia Day Foundation UK Trust for helping fund my studies, as well as all those who helped me fund raise and showed their support.

I would like to thank all my friends who supported me along the way and made my time in Oxford (and Sydney) so rewarding.

I would finally like to thank the never ending support of my family – I truly stand on the shoulders of giants. I could not have done it without you.

Abstract

Wind energy is rapidly expanding as the demand for a low-emissions electricity grid increases. Acknowledging this, and the intermittent nature of the wind resource, both short-term (circa hours) and long-term (circa years) planning must be implemented. To do this accurately, the performance degradation of wind farms must be quantified and accounted for in electricity systems and financial models. This thesis is concerned with quantifying aerodynamic performance decline of wind turbines and wind farms to better understand changes in lifetime performance and more accurately predict wind farm power output.

Wind turbine aerodynamic performance losses are investigated by analysing the effect of roughness on airfoil performance (lift and drag). One of the challenges in quantifying wind turbine performance decline due to roughness is the lack of roughness airfoil data available. A 'roughness evolution parameter' is therefore parameterised based on publicly available roughened airfoil experimental tables, which enables roughened airfoil data to be estimated from clean airfoil data. The calculation of the parameter is proposed as a framework, with room for further refinement when more experimental data comes available.

Blade element momentum theory is used to analyse the impact of blade roughness on wind turbine performance. Exemplar wind turbine rotors are designed and the impact of roughness on annual energy production is quantified to be in the range of 2.9–8.6% for realistic levels of roughness. A novel control method is proposed to mitigate losses due to roughness and 0.1–1.0% of annual energy production is recovered.

To account for the long-term variability in wind speed, wind farm performance decline is quantified using long-term corrected capacity factor data from wind farms in the UK, US and Australia – 0.26, 0.72 and 0.43 capacity factor percentage points per year respectively. To do so, power output data is combined with reanalysis wind speed data to correct for variability and biases in the data, allowing the underlying performance trend to be realised. To address shortcomings in the data (e.g. undetected downtime or reanalysis errors) a novel long-term correction approach is used which reduces variability in the results.

Reanalysis wind speed data and wind farm datasets are utilised for a supervised machine learning performance prediction model. Due to the inclusion of additional

metadata, this model is able to predict wind farm capacity factor output with a 56–90% improvement in mean absolute error value when compared to predicting power output using a wind farm power curve.

Contents

List of Figures **xi**

Nomenclature (Chapters 2–4) **xix**

Nomenclature (Chapters 5–8) **xxi**

List of Abbreviations **xxiii**

1 Introduction **1**

1.1 Motivation 5

1.2 Thesis Objectives 5

1.2.1 Objective 1: Optimising the Performance of a Rough-Bladed Turbine 7

1.2.2 Objective 2: Roughness Evolution Parameter 7

1.2.3 Objective 3: Wind Farm Aerodynamic Performance Loss . . 8

1.2.4 Objective 4: Wind Farm Performance Prediction 8

1.3 Thesis Outline 8

2 Performance Decline of Wind Turbines **11**

2.1 Wind Turbine Blade Roughness 12

2.1.1 Roughness Characterisation 12

2.1.2 Insect Damage 13

2.1.3 Ice Accretion 15

2.1.4 Rain 17

2.1.5 Airborne particles 18

2.1.6 Impact of Roughness on Performance 19

2.2 Analytical Methods for Modelling a Wind Turbine 20

2.2.1 Turbine Control (using BEM) 23

2.2.2 BEM implementation 24

3 Aerodynamic Performance Degradation and Mitigation of a Single Roughened Turbine 27
3.1 Introduction . . . 28
3.2 Method . . . 30
3.2.1 Computational Method and Available Airfoil Tables . . . 30
3.2.2 Rotor Design . . . 34
3.2.3 Turbine Control . . . 36
3.2.4 Performance Loss . . . 42
3.3 Results . . . 43
3.3.1 Rotor Performance . . . 43
3.3.2 Turbine Performance Comparisons . . . 45
3.3.3 Spanwise Effects of Roughness . . . 49
3.3.4 Roughness Effects on Airfoils . . . 52
3.3.5 Energy Loss . . . 55
3.4 Conclusions . . . 56

4 Parameterising the Evolution of Roughness on an Airfoil 59
4.1 Introduction . . . 60
4.2 Roughened Airfoil Data . . . 62
4.2.1 Comparing the Effect of Roughness . . . 62
4.2.2 Available Airfoil Datasets . . . 63
4.2.3 Data Comparison . . . 67
4.3 Roughness Effects on Airfoil Performance . . . 68
4.4 Roughness Impacts on Wind Turbine Performance . . . 70
4.4.1 Turbine Model . . . 71
4.4.2 Performance Change Metrics . . . 72
4.4.3 Impact of Roughness on Performance Over Time . . . 75
4.5 Discussion and Conclusions . . . 76

5 Wind Speed Variability and Wind Farm Degradation 79
5.1 Wind speed variability . . . 81
5.1.1 Comparing results to the Weibull distribution . . . 87
5.2 Reanalysis Modelling . . . 88
5.3 Performance Decline of Wind Farms . . . 90
5.4 Data Sources . . . 92
5.4.1 Ofgem . . . 93
5.4.2 EIA . . . 94
5.4.3 AEMO . . . 95
5.5 Data processing . . . 95
5.5.1 Raw data filtering . . . 97

5.5.2 Wind Speed Data . . . 98
5.5.3 Missing hub height data . . . 98
5.5.4 Scaling wind speed to hub height . . . 100
5.5.5 Converting wind speed to power . . . 101
5.5.6 Curtailment . . . 107
5.5.7 Removing outliers . . . 109
5.6 Conclusion . . . 110

6 Performance decline of wind farms in the UK, US and Australia 111
6.1 Introduction . . . 112
6.2 Previous long term correction work . . . 114
6.3 Raw Data Trends . . . 118
6.4 Long-Term Correction . . . 125
6.4.1 Existing LTC methods . . . 127
6.4.2 Proposed LTC method . . . 129
6.4.3 Testing the LTC assumptions . . . 132
6.5 Importance of Wind Farm Size . . . 137
6.6 Results . . . 139
6.6.1 Reanalysis wind speed data . . . 139
6.6.2 Wind speed trends . . . 141
6.6.3 Comparing estimated and observed power output . . . 144
6.6.4 Long-term correction and performance decline . . . 147
6.6.5 Capacity-weighted results . . . 160
6.7 Discussion . . . 164
6.7.1 Application . . . 167
6.8 Conclusions . . . 168

7 Wind Farm Performance Prediction 171
7.1 Machine learning model selection . . . 173
7.1.1 XGBoost . . . 177
7.1.2 Random Forest . . . 178
7.2 Model implementation . . . 178
7.3 Results . . . 181
7.3.1 Feature importance . . . 185
7.3.2 Data limitations . . . 188
7.3.3 Modelling results within a country of the dataset . . . 191
7.4 Discussion and conclusions . . . 193

8 Conclusion **199**
8.1 Conclusions 200
8.1.1 Wind turbine roughness 200
8.1.2 Long-term correction and wind farm performance decline . . 201
8.1.3 Wind farm power prediction 202
8.2 Future work 202
8.2.1 Roughness data 202
8.2.2 Wind farm data and reporting standards 202
8.3 Contributions 203

Appendices

A Wind Farm Metadata **207**

References **209**

List of Figures

1.1 An example of roughness on an airfoil section. Adapted from Gao et al. [10] 3

2.1 Insect damage on a wind turbine blade. Source: Doosttalab [36] . . 14
2.2 Streamwise icing accumulation on the leading edge of a wind turbine blade. Source: Doosttalab [36] 15
2.3 Leading edge erosion on a wind turbine blade. Source: Verma et al. [52] 18
2.4 The velocity triangle overlayed onto an airfoil. Adapted from Ning [71]. 22
2.5 A schematic showing a blade element sweeping out an annular ring. Source: Wind Energy Handbook [70]. 22
2.6 An example power curve, and operating regions for a 5 MW turbine. Adapted from Aho et al. [73] 24

3.1 Lift to drag ratio and aerodynamic polars for the three selected airfoils (NACA4415, S801, S810). Plots (a) and (b) present the NACA4415 airfoil data, (c) and (d) present the S801 airfoil, (e) and (f) present the S810 airfoil. 33
3.2 Spanwise distribution of rotor solidity σ_r (left) and blade twist angle γ_t (right) for rotor designs using the NACA4415 (red), S801 (blue) and S810 (green) airfoils. 36
3.3 Schematic of the different control strategies used for both clean and rough rotors. All control methods with a clean blade yield identical power curves and have been labelled "CMclean". 38
3.4 Wind speed probability distributions modelled with (a) the two parameter Weibull distribution and (b) the cumulative probability density for the Kentish Flats, Thanet, Greater Gabbard and Gwynt y Mor sites. Vertical lines have been added to indicate U_{ci}, U_r and U_{co}. Wind speed data is from 10 minute average data extrapolated through a log law profile to a common altitude of 80m. 44

3.5 Performance comparison between rotors with roughened and clean blades for the three rotor designs (NACA4415 left, S801 centre, S810 right) for all control methods. The x-axis for all plots is wind speed. Two inset graphs are present: one for the NACA4415 power vs wind speed, which demonstrates the improved performance of CM3; and one for the NACA4415 torque vs wind speed, which illustrates the change in gradient for the CM3 torque profile in this area. 45
3.6 Torque as a function of rotational speed for the (a) NACA4415 rotor, (b) S801 rotor and (c) S810 rotor, for both clean and rough airfoil conditions and all control methods. 47
3.7 Spanwise contribution to (a) AEP (evaluated for the Kentish Flats wind distribution) and (b) blade torque for the NACA4415, S801 and S810 rotors in clean and rough conditions using CM2/CM2b. . 50
3.8 Spanwise analysis of NACA4415, S801, S810 rotors using CM2/CM2b for a wind speed of 10m/s showing the spanwise variation of (a) angle of attack, α, (b) tangential force coefficient, $C_t(\phi)$, (c) axial induction factor, a and (d) normal force coefficient, $C_n(\phi)$. 51
3.9 Airfoil characteristics c_l/c_d in the angle of attack region most relevant to turbine operation, $0^\circ < \alpha < 10^\circ$ for the NACA4415 airfoil. The operational α values (40–100% span) for the different control strategies have been highlighted, with clean cases denoted CM_{clean} and rough cases denoted CM_{rough}. Arrows indicate the impact of roughness and the control recovery mechanism of CM3. 53
3.10 Airfoil characteristics c_l (a), c_d (b) and c_l/c_d (c) in the angle of attack region most relevant to turbine operation, $0^\circ < \alpha < 10^\circ$ for NACA4415, S801 and S810 airfoils. The operational α values (40–100% span) for the different control strategies have been highlighted, with all clean cases denoted CM_{clean}. 54
3.11 Energy capture comparison between clean and rough rotors operated with different control methods over a 1 year time period. AEP changes are measured relative to the clean operating rotor, and CM3 recovery is measured relative to CM2b. 55

4.1 Experimental ice shapes (a) and LEWICE ice shapes (b), sourced from Hann et al. [102]. 65
4.2 Experimental wind tunnel data comparison between OSU study at $Re = 1.5 \times 10^6$ and Sandia study at $Re = 1.6 \times 10^6$. 67
4.3 Percentage change in c_l is plotted against percentage change c_d for all datasets at the angle of attack for the maximum clean lift-to-drag ratio ($(c_l/c_d)^*$) point with linear (black) and power law (orange) best fit lines, plus the surrounding angles ($\pm 2^\circ$) indicated in grey. 69

4.4 Percentage change in lift is plotted against percentage change in drag (at α^* for clean airfoils) for the OSU, Sandia and Jasinski et al. datasets with the linear (black) and power law (orange) best fit lines superimposed. The legend indicates the roughness height. 71
4.5 Percentage change in lift is plotted against percentage change in drag for all datasets with linear (black) and power law (orange) best fit lines are shown in all plots, and contour plots of the modelled C_P values (a), percentage change AEP (b) and C_T values (c) for the DTU 10 MW RWT. 74
4.6 Wind farm age against AEP curves from Staffell and Green [18], Olauson et al. [19] and Germer and Kleidon [20] (a). γ against AEP as calculated using BEM analysis of the DTU 10 MW RWT (b). Wind turbine age against γ (c). 76
4.7 Wind turbine age against AEP performance change due to roughness based on wind farm performance decline rates from Staffell and Green [18], Olauson et al. [19] and Germer and Kleidon [20] 77

5.1 Location of the Edithburgh and Swan Island weather stations. The closest wind farm to these weather stations are also shown. This figure shows the south-east corner of Australia. 82
5.2 Wind distributions for average wind speed between 7–8 m/s over the averaging periods of 1 hour, 1 day and 1 month for the Swan Island weather station. 84
5.3 Wind distributions for varying average wind speeds (3–4 m/s, 6–7 m/s, and 9–10 m/s) over the averaging periods of 1 hour (left) and 1 day (right) for the Swan Island weather station. 85
5.4 Wind distributions for varying average wind speeds (3–4 m/s, 6–7 m/s, and 9–10 m/s) over the averaging period of 1 hour fit with the best fit distribution indicated and a Weibull minimum distribution. The analysis is repeated for the Swan Island weather station (left) and Edithburgh weather station (right). 89
5.5 Location, size and commencement year information for UK wind farms that had at least 5 years of operation to 2020 used in this thesis. 94
5.6 Location, size and commencement year information for US wind farms that had at least 5 years of operation by 2020 used in this thesis. 95
5.7 Location, size and commissioning year information for Australian wind farms that had at least 5 years of operation by 2020 used in this thesis. This figure shows the south-east corner of Australia. . . 96

5.8 Summary data for the wind farms used in this study. Figures in each column represent the number of wind farms, number of observations by year, number of observations by farm age and distribution of capacity factors. Each row represents the different countries – UK, US, AUS, and ALL, from top to bottom. UK and US observations are monthly, AUS observations are hourly and ALL is a blend of monthly and hourly. 97
5.9 Scatter plot of predicted hub heights from the validation dataset of the gradient-boosting algorithm used to predict unknown hub heights. 100
5.10 Capacity factor as a function of wind speed for an individual wind turbine (blue) and compared to a wind farm (orange) for an exemplar turbine CF curve. 103
5.11 (a) Reanalysis hourly wind speeds for Woolnorth Wind Farm (from 2009 to 2020) in Australia, with Weibull distribution fit. (b) The resulting wind speed distributions used for estimating CF when the hourly wind speed $U = 8.19$m/s (orange curve), decreases by 3m/s (blue curve) and increases by 3m/s (grey curve). 106
5.12 An overview of the CF production process. 106
5.13 Percentage of wind farm rated capacity, on average, that is curtailed during recorded curtailment events annually for applicable Australian wind farms. Applicable Australian wind farms are those who are required to produce curtailment data (semi-scheduled generators, 19 out of 33 wind farms in the dataset). 108
5.14 Wind speed against capacity factor with curtailment events shown in red, for the Musselroe Wind Farm. Note a CF of 100% is not achieved because the data represents exported power (to the electricity grid), not mechanical power. 109

6.1 Failure rate and downtime of onshore European wind turbines [149]. 113
6.2 Yearly averaged observed capacity factors against age for wind farms in (a) UK, (b) US, (c) AUS and (d) all three countries combined. The shaded regions indicate yearly percentile ranges of CF_{obs} and the yearly means are shown with squares. Each country has a linear regression line computed from age 1 onwards. 120
6.3 Mean lifetime CF for wind farms in the (a) UK, (b) US, (c) AUS and (d) ALL grouped by their commission year. The size of the markers represents the total capacity commissioned in the year, and the colour represents the average length of the data series. 122
6.4 A fictive example of CF against age, demonstrating the importance of including seasonal components to the regression equation to accurately capture the linear trend. 123

6.5 Distribution of lifetime degradation rates for individual wind farms for (a) UK, (b) US, (c) AUS and (d) ALL. 124
6.6 Mean degradation rate (regression slope from Equation 6.2) for (a) UK, (b) US, (c) AUS and (a) ALL wind farms grouped by their commission year. The size of the markers represents the total capacity commissioned in that year, and the colour represents the average length of the data series. 125
6.7 Monthly averaged observed capacity factors against age for (a) UK, (b) US, (c) AUS and (d) all three countries combined. The shaded areas indicate percentile ranges and the monthly means are showed. Each country has a regression line computed from age 1 onwards. . 126
6.8 LTC process flow chart. 129
6.9 Correlation between the error correction in the LTC process and age for each wind farm in the dataset. 135
6.10 Average age of each wind farm plotted against the LTC correction process ($\Delta CF \cdot D$) described in Equation 6.8 calculated using reanalysis wind speed data from (a) ERA5 and (b) MERRA-2. 136
6.11 Fictitious data showing examples of heteroscedasticity (left), linearity violated (middle), and compliant data (right). The upper row is plotting age against CF and the bottom row shows predicted CF (from the linear regression) against residuals. 137
6.12 The mean wind farm capacity by commission year for UK, US, and AUS. 138
6.13 The Pearson correlation between ERA5 and MERRA-2 wind speed data at each wind farm for (a) UK, (b) US, and (c) AUS. 140
6.14 The mean difference in m/s between ERA5 and MERRA-2 wind speed data for each wind farm for (a) UK, (b) US, and (c) AUS. . . 142
6.15 The absolute wind speed time series regression slopes at individual wind farm locations using both ERA5 and MERRA-2 data for (a) UK, (b) US, and (c) AUS. On average, at individual wind farm locations wind speed is shown to decrease for the UK and US, and increase for AUS. 143
6.16 Monthly averaged ERA5 wind speed trends for wind farm locations in (a) UK, (b) US, and (c) AUS. 145
6.17 Monthly averaged MERRA-2 wind speed trends for wind farm locations in (a) UK, (b) US, and (c) AUS. 146
6.18 Scatter plot comparing average estimated capacity factors calculated using ERA5 data against average actual capacity factors for each wind farm in the UK (blue), US (grey) and AUS (orange). 147

6.19 Rolling correlation between CF_{est} and CF for (a) the UK, (b) US, and (c) AUS. The correlation period for the UK and US is three months and is one week for AUS. 148
6.20 A comparison of LTC methods in across a 5 year period for the Ardrossan wind farm in Scotland. Shown is (a) the reanalysis wind speed, (b) predicted CF and (c) CF_{LTC}. 150
6.21 A comparison of LTC methods across a 1 day period for the Capital wind farm in Australia at age 7 (i.e., the time period between age 7 years and 0 days to age 7 years and 1 day). Shown is (a) the reanalysis wind speed, (b) predicted CF and (c) CF_{LTC}. 151
6.22 Wind farm performance plotted against age for the long term corrected capacity factors (using the Kelly method and ERA5 reanalysis wind speeds) for wind farms in (a) the UK, (b) US, (c) AUS and (d) ALL (as per Figure 7.2 for raw data only). A linear regression line was fitted to the yearly averaged CF values across all wind farms. . 153
6.23 Wind farm performance plotted against age for the long-term corrected capacity factors (using the Kelly method and ERA5 reanalysis wind speeds) for wind farms in the US separated by wind farms with a final age (in the dataset) of above or below age 13 years. A linear regression line was fitted to the yearly averaged CF values across all wind farms in each sample set. 154
6.24 Wind farm degradation rates averaged by commission year and calculated using the long-term corrected capacity factors (using the Kelly method and ERA5 reanalysis wind speeds) for wind farms in (a) the UK, (b) US, (c) AUS and (d) ALL. A linear regression line was fitted to the yearly averaged CF values across all wind farms. . 155
6.25 Wind farm performance plotted against age for the monthly averaged long-term corrected capacity factors (using Kelly LTC method and ERA5 reanalysis wind speeds) for wind farms in (a) the UK, (b) US, (c) AUS and (d) ALL. A linear regression line was fitted to the monthly averaged CF values for all wind farms. 157
6.26 Distribution of degradation rates for individual wind farms using the CF_{LTC} data produced using the ERA5 reanalysis wind speed data for (a) the UK, (b) US, (c) AUS and (d) ALL. 159
6.27 Wind farm performance plotted against age for the long-term corrected capacity factors (using the Kelly method and ERA5 reanalysis wind speeds) for wind farms in (a) the UK, (b) US, (c) AUS and (d) ALL. A linear regression line was fitted to the wind farm capacity weighted-average yearly CF values across all wind farms. 162

6.28 The distribution of wind farm rated capacities in 10 MW increments by country (UK, US, and AUS). 163
6.29 Wind farm performance plotted against age for (a) the monthly averaged long-term corrected capacity factors, and (b) monthly wind farm capacity weighted-averaged long term corrected capacity factors (using Kelly LTC method and ERA5 reanalysis wind speeds) for wind farms in the UK. Linear regression lines were fitted to the monthly averaged CF values for (a) all wind farms, and (b) to wind farm capacity weighted-averaged CF values for all wind farms. 164

7.1 An example decision tree for predicting CF produced using XGBoost and a reduced feature list. 177
7.2 The predicted CF from the XGBoost models plotted against the validation (observed) CF values. The models were trained with data from the UK (a), US (b), AUS (c) and ALL (d). All plots show models trained using the ERA5 wind speed features and the ALL plot is displaying an equal number of CF_{pred} data points for each country with the outputs colour coded for the UK (blue), US (red) and AUS (green). 183
7.3 The predicted CF from the XGBoost models plotted against the validation (observed) CF values. The models were trained with data from wind farms in AUS and show models trained on hourly data (a), daily data (b), weekly data (c) and monthly data (d). All plots show models trained using the ERA5 wind speed features. 184
7.4 Feature importance from the XGBoost models trained with data from wind farms in the UK (a), US (b), AUS (c) and ALL (d). All plots show models trained using the ERA5 wind speed features and feature importance values adding up to 0.9. 185
7.5 Mean ERA5 wind speed at each wind farm in the UK, US, AUS and ALL. 188
7.6 Mean CF for the entire dataset (ALL) by commissioning year, and manufacturers GE, Vestas and Siemens (a) and cumulative number of wind farms in the data set (b) for manufacturers GE, Vestas and Siemens. 189
7.7 The cumulative density functions for $|CF_{obs} - CF_{est}|$ from Chapter 6 (labelled 'Power Curve'), and $|CF_{obs} - CF_{pred}|$ using the XGBoost model for ERA5 data (a) and MERRA-2 data (b). 190
7.8 The MAE for predicting CF for each wind farm that was removed from the training dataset for (a) the UK, (b) US, (c) AUS and (d) ALL. Each wind farm is modelled using the ERA5 wind speed features. 192

7.9 The size of wind farm in MW plotted against the MAE of CF_{pred} for the LOOCV analysis split by models trained with wind farm data from the UK (a), US (b), AUS (c) and ALL (d). 193
7.10 The number of turbines in a wind farm plotted against the CF MAE for the LOOCV analysis split by models trained with wind farm data from the UK (a), US (b), AUS (c) and ALL (d). 194
7.11 The size of wind farm in MW plotted against the CF MAE (a), and the number of turbines in a wind farm plotted against the MAE (b) for the LOOCV analysis split by models trained with wind farm data from ALL. The colour of the scatter points indicate the start year of the wind farm. 195

Nomenclature (Chapters 2–4)

α	Angle of attack
$\tilde{\alpha}$	Representative angle of attack across a rotor blade
α_{sep}	Angle of attack where separation starts to occur
β	Root pitch angle of the wind turbine blades
β_g	Initial geometric twist angle along a wind turbine blade
β_{max}	Maximum pitch of a wind turbine rotor
γ_t	Blade twist angle distribution, accounting for root pitch
κ	von Kármán constant
λ	Tip speed ratio
μ	Air viscosity
ω	Rotational speed
ω_r	Rated rotational speed
ϕ	Inflow angle
ρ	Air density
σ	Rotor solidity
τ	Torque
τ_{aero}	Aerodynamic torque
τ_{gen}	Generator torque
τ_r	Rated torque
A	Rotor swept area
a	Axial induction factor
a'	Tangential induction factor
B	Number of turbine blades
c	Blade chord length
c_s	Weibull scale parameter

c_d	Sectional drag coefficient
c_l	Sectional lift coefficient
c_l/c_d	Lift to drag ratio
C_n	Normal (thrust-generating) force component
C_P	Power coefficient
C_T	Thrust coefficient
C_t	Tangential (torque-generating) force component
C_x	Local thrust coefficient
d	Sectional drag force
F	Tip-loss factor
F_x	Streamwise force
k_s	Weibull shape parameter
K	Turbine control parameter
l	Sectional lift force
P	Power
P_{ci}	Cut-in power
P_r	Rated power
r	Radius
R	Blade tip radius
Re	Reynolds number
t	Time
T	Turbine thrust
U	Wind speed
U_∞	Free stream wind speed
U_f	Friction velocity
U_x	Steamwise wind velocity at the rotor plane
V	Wind speed relative to the airfoil section
W	Relative wind speed
z	Height
z_0	Wind profile roughness height

Nomenclature (Chapters 5–8)

ΔCF	The difference between observed CF values and outputs from LTC methods
Γ	Gamma function
γ	Roughness evolution parameter
κ	Von Karman's constant
μ	Mean
c_s	Scale parameter
CF	Capacity factor
CF_{est}	Estimated capacity factor (from wind farm power curve model)
CF_{obs}	Observed capacity factor
CF_{pred}	Predicted capacity factor (from machine learning model)
D_t	Dampening correction
M	Median
k	Roughness height
k_s	Shape parameter
P	Power
R^2	R-squared value
U	Wind speed
U_f	Friction velocity
w	Weighting factor
z	Height above ground
z_0	Wind profile roughness height

List of Abbreviations

2-D	Two dimensional
3-D	Three dimensional
AARL	Aeronautical and Astronautical Research Laboratory
AEMO	Australian Energy Market Operator
AEP	Annual energy production
AUS	Australia
BEM	Blade element momentum theory
BOM	Bureau of Meteorology
CCWE	China Creative Wind Energy
CDF	Cumulative density function
CFD	Computational fluid dynamics
CHP	Combined heat and power
CM	Control method
COSMO	Consortium for Small-scale Modeling
CV	Cross validation
DTU	Technical University of Denmark
ECMWF	European Centre for Medium-Range Weather Forecasts
EIA	Energy Information Administration
EMD	Europe Mesoscale Data
ERA	ECMWF reanalysis
ERCOT	Electric Reliability Council of Texas
HCS	Hill-climb search
JRA55	Japanese 55-year Reanalysis
LASSO	Least absolute shrinkage and selection operator
LCOE	Levelised cost of energy

LEGR	Leading edge grit roughness
LOOCV	Leave-one-out cross-validation
LTC	Long-term correction
LWK	*Landwirtschaftskammer Schleswig-Holstein*
MAD	Mean absolute deviations
MAE	Mean absolute error
MAPE	Mean absolute percentage error
MARS	Multivariate adaptive regression splines
MERRA-2	Modern Era Analysis for Research and Applications-2
ML	Machine learning
MLP	Multilayer perceptron
MPPT	Maximum power point tracking
MSE	Mean squared error
NASA	National Aeronautics and Space Administration (United States)
NCAR	National Center for Atmospheric Research (United States)
NCEP	National Centers for Environmental Prediction (United States)
NEG	Nordtank Energy Group
NEM	National electricity market
NREL	National Renewable Energy Laboratory (United States)
Ofgem	Office of Gas and Electricity Markets
OLS	Ordinary least squares
OSU	Ohio State University
OTC	Optimal torque control
R1	Reanalysis 1
RANS	Reynolds-Averaged Navier-Stokes
REA	Regional analysis
RE	Relative error
REF	Renewable Energy Foundation
RF	Random Forest
RWT	Reference wind turbine
Sandia	Sandia National Laboratories

SCADA	Supervisory control and data acquisition
SMC	Sliding mode control
SVM	Support Vector Machine
UK	United Kingdom
US	United States of America
VRE	Variable renewable energy
VTT	Technical Research Centre of Finland
WI	Wind index
WMEP	*Wissenschaftliches Mess- und Evaluierungsprogramm*
WRF	Weather Research and Forecast
XGBoost	Extreme gradient boosting

We are the first generation to feel the impact of climate change and the last to be able to do anything about it.

— US President Barack Obama

1 Introduction

Contents

1.1 Motivation 5
1.2 Thesis Objectives 5
1.2.1 Objective 1: Optimising the Performance of a Rough-Bladed Turbine 7
1.2.2 Objective 2: Roughness Evolution Parameter 7
1.2.3 Objective 3: Wind Farm Aerodynamic Performance Loss 8
1.2.4 Objective 4: Wind Farm Performance Prediction 8
1.3 Thesis Outline 8

Interest in wind energy has led to a number of studies aimed at quantifying and eventually harnessing the available resource. Hoogwijk et al. [1] estimate that global onshore wind energy potential is 96 PWh per year, which corresponds to approximately 6 to 7 times the 2001 world electricity consumption. Extending this analysis to include offshore sites, Archer and Jacobson [2] used wind data from meteorological sites to show that in the year 2000, approximately 628 PWh of electricity could have been produced globally using 1.5 MW wind turbines. If only ≈20% of this electricity was captured, it could satisfy 100% of the world's energy demand for all purposes and over seven times the world's electricity needs (in the year 2000) [2]. Lu et al. [3] built on Archer and Jacobson's methodology to show that if a network of land-based 2.5 MW wind turbines were used and operating

at 20% of their rated capacity, then they could supply > 40 times the worldwide consumption of electricity and > 5 times total global use of energy in all forms. While these studies vary in their modelling assumptions and outcomes, it is clear that wind energy provides an abundant energy source globally. Needless to say, there is vast potential in wind energy and the interest in large scale installation of wind turbines is therefore justified.

Several global wind potential studies assume a 20% capacity factor (CF) as a conservative input assumption, however many studies expect the CF to be greater than 30% [4–6]. However, energy production of wind turbines deployed in the field has sometimes fallen below these high-level resource estimates. Boccard [7] used recorded production data of onshore wind power in Europe in years 2003–2007 to show that the capacity factors for wind farms in this region were below 21%. Boccard claimed the expected capacity factor was in the 30–35% range by summarising previous estimates from the literature. This under-performance of wind farms could be due to many factors, and is an important aspect to understand in order to improve yield forecasts and financial models. In summarising previous literature, Boccard shows the wide range of CF estimates, from 17% [8] to 60% [9], which highlights the difficulty in estimating the variable wind resource.

Wind farm under-performance is not attributed directly to one source, however variability of the wind resource is often a key driver. Two aerodynamic factors that directly lead to performance decline are wake effects and degradation of the turbine blades. Wake effects will impact a wind farm performance immediately, whereas degradation (often referred to as roughness) of the blades will accumulate over time, which contributes to the ageing affect on wind farms. Roughness can be understood as altering the shape of the airfoil, and usually occurs on the leading edge (due to exposure to incoming airflow). Figure 1.1 is adapted from Gao et al. [10] and shows an example of an airfoil section with altered shape on the leading edge. By understanding these effects, wind farms could be operated more optimally, reducing costs and allowing for further implementation globally. It will also lead to improved yield prediction, which is increasingly important as the

penetration of wind power continues to rise. This is important to adequately fulfill the expected future reliance on wind energy.

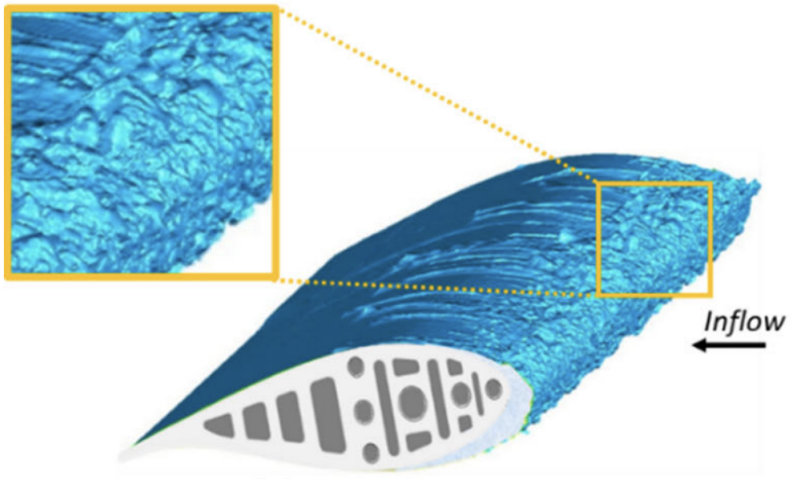


Figure 1.1: An example of roughness on an airfoil section. Adapted from Gao et al. [10]

In recent years, large scale wind energy has become competitive with fossil fuel sources, with offshore wind farms winning bids to produce energy more cheaply than gas-fired power plants in the UK [11]. This has been aided by the increasing size of wind turbines (up to 12 MW machines [12]) to maximise energy extraction per turbine. Blanco [13] estimates that up to 80% of the cost of implementing a wind farm comes from the wind turbine materials, foundations, road construction and grid connection, so producing more power from a similar-sized footing and in the same location makes it substantially cheaper. The same principle applies for offshore wind; Heptonstall et al. [14] assume 88% of capital costs come from the turbine and tower (47%), foundations (22%) and electrical infrastructure (19%). Another way to reduce costs is by understanding performance decline to reduce uncertainty around wind farm lifespans.

Different methods have been investigated to assess the performance reduction impact for roughness-affected wind turbines [15–17]. The majority of studies have used Blade Element Momentum (BEM) theory to model the impact of blade roughness on a wind turbine. The aerodynamic affects of roughness are included in the model by modifying input airfoil data, sourced from experimental or simulations

of (usually) 2-D airfoils. However, it is currently unclear what effect operational control methods could have on mitigating the performance reduction.

Wind farm performance decline is separate from BEM-based roughness studies. The impact of blade roughness is very difficult to monitor over time because academic studies have only investigated static roughness at one or multiple heights and not how it changes over time. Furthermore, roughness is location-dependent and not monitored on operational wind farms. Due to this, when investigating wind farm performance, roughness is a part of many reasons for performance decline.

For ageing wind farms, performance declines have been quantified by Staffell and Green [18], Olauson et al. [19], Germer and Kleidon [20], and Hamilton et al. [21] using mesoscale wind models and electricity production data. Staffell and Green show that UK farms between the years 2002–2012 degrade at 1.6% $\pm$ 0.2% per year, Olauson et al. show a 0.5% performance decline for Swedish wind farms, Germer and Kleidon show a 0.63% $\pm$ 0.01% decline for German wind farms, and Hamilton et al. show a 0.17% performance decline for wind farms aged under 10 years and a 0.53% decline thereafter. While these studies provide insight into general performance decline trends, these trends have not yet been connected to causes of performance decline.

The research presented hereafter focuses on the effects of blade roughness on wind turbine performance, as well as modelling wind farm performance decline over time. Blade roughness effects on a single wind turbine have been studied using an analytical model. This model is used to quantify performance decline of a single wind turbine and explore control strategies that mitigate performance losses of turbines with roughened rotor blades. This model is then extended to broader wind farm effects by incorporating recorded power production data and comparing it with simulated wind speeds from a reanalysis model to quantify performance trends of wind farms. One of the main challenges with this process is accounting for the yearly and seasonal variation of wind speed, and thus a novel long-term correction (LTC) application has been implemented to find the underlying performance trends. Wind

farm data used to quantify performance trends are then combined with reanalysis wind speed data to predict wind farm power output.

1.1 Motivation

Improved understanding of lifetime performance of wind farms is crucial for future grid planning and levelised cost of energy (LCOE) models. This is especially pivotal as the world transitions to variable renewable energy (VRE) sources for power production and grid operators have to ensure the required electricity generation capacity is available. LCOE is calculated for the lifetime of a wind farm and energy generation is a key input [22], which makes understanding how performance changes over time vitally important. Understanding this also mitigates an investment risk of the wind farm generating less electricity over time than expected [23]. Furthermore, understanding how to operate turbines more efficiently as they age will significantly increase their value over time. To utilise improved models and turbine operation, performance declines of wind turbines and farms need to be accurately quantified and novel control methodologies implemented.

Another important motivator for this work is using accessible wind speed datasets to understand and improve the wind energy sector. In this thesis, quantifying the impact of blade roughness on turbine performance and wind farm performance prediction will be improved using publicly accessible datasets. Wind speed data from reanalysis models are used to both quantify performance losses and predict operational or potential wind farm power output. Using publicly accessible datasets accurately and effectively can unlock further research into a currently data-limited area.

1.2 Thesis Objectives

Significantly utilising the global wind resource will require wind turbines and wind farms to operate efficiently throughout their operational life, as well as extend their life spans. To do this, quantifying and understanding performance loss of

operational wind turbines and farms is necessary. Performance loss can be split into two main categories: aerodynamic and mechanical losses. The objective of this thesis is to create frameworks to consider aerodynamic performance degradation for wind turbines and performance decline for wind farms.

Performance degradation of individual wind turbines has become established in the literature, however, applying aerodynamic losses to new rotors and optimising performance of aerodynamically degraded – or roughened – turbines is lacking. Furthermore, quantifying entire wind farm performance decline with age is not well understood. These challenges are addressed through four primary objectives: the first concerns finding a 'roughness evolution parameter' which can be used to simulate the effect of blade roughness on a wind turbine's performance; the second involves optimising the performance of an individual wind turbine with roughened blades using an analytical model under different control algorithms; the third incorporates observed wind farm data to quantify and investigate real-world performance loss in wind farms; and the fourth uses wind farm output data combined with modelled wind speed data to predict wind farm power output.

This research aims to further the understanding of wind farm performance decline for an individual wind turbine and a wind farm. In particular, this work focuses on:

1. Quantifying the aerodynamic performance loss on airfoils (reduction in lift, increase in drag) and a wind turbine due to blade roughness

2. Developing a wind turbine control strategy to mitigate the performance loss from blade roughness

3. Defining a 'roughness evolution parameter' that can convert clean airfoil data into roughened data, which can be used to simulate the effect of roughness on a wind turbine without the need for experimental data

4. Producing a framework to more accurately predict wind farm power using reanalysis-derived wind speeds compared to existing methods

5. Producing a novel long-term correction method that is used to quantify performance decline of wind farms

6. Utilising a supervised machine learning model to predict wind farm power output for a given location and wind farm size (incorporating performance decline)

1.2.1 Objective 1: Optimising the Performance of a Rough-Bladed Turbine

Modelling rough-bladed wind turbines analytically requires airfoil data produced under roughened conditions as inputs. Previous BEM methods in this area have explored performance losses without optimising for aerodynamic control. The objective of this section is to mitigate performance losses due to blade roughness by changing the control of the wind turbine to optimise pitch for a roughened blade with modified aerodynamic characteristics.

1.2.2 Objective 2: Roughness Evolution Parameter

A challenge of analysing roughened turbines is the lack of suitable airfoil data available to input into models. This is limiting because of difficulties around producing these data, both numerically and experimentally. Numerically producing the data through computational fluid dynamics (CFD) methods is especially difficult due to the inability to capture the complex flow phenomena on this scale. Instead, carefully produced experimental data has generally been preferred.

The objective of this research is to determine a method which allows the impact of airfoil surface roughness to be estimated based on a parameter that modifies the clean airfoil data. This enables the effect of roughness on wind turbines to be estimated without necessitating the production of roughened airfoil data.

1.2.3 Objective 3: Wind Farm Aerodynamic Performance Loss

Analysing a single turbine is useful to understand what is occurring aerodynamically, but it neglects other ageing effects such as generator performance decline and mechanical wear. It also ignores interactions between turbines, such as wake effects and blockage, that occur on a wind farm scale. Realistically, performance decline of wind farms will occur from a combination of sources and may change with respect to aspects such as location, turbine size and farm size.

The third objective of this thesis is to create a framework to analyse aerodynamic performance decline for wind farms in the United Kingdom, the United States and Australia.

1.2.4 Objective 4: Wind Farm Performance Prediction

The final objective of this thesis is to utilise wind farm power output data, wind farm metadata (number of turbines, turbine size, etc.), and reanalysis wind speed data to predict power output of a wind farm for a target period based on wind farm site data only.

1.3 Thesis Outline

The thesis has been broadly divided into three sections – performance decline of a single turbine, performance decline of wind farms, and wind farm power prediction.

The first section investigates aerodynamic impacts of roughness on airfoils and a turbine. Chapter 2 begins with a review of the contributing factors of blade roughness and the use of semi-analytic models in modelling a single turbine (with and without roughness applied). A guaranteed convergence blade element momentum model is introduced, which is used throughout the first section. Chapter 3 utilises roughened airfoil data to model the impacts on turbine performance and investigate the potential for control strategies to mitigate performance losses, and Chapter 4 draws trends from existing experimental airfoil roughness studies to generalise the effect of roughness on airfoil lift and drag characteristics.

The second section quantifies the performance decline of wind farms in the United Kingdom, United States, Australia and all three countries combined. This is investigated using reanalysis wind speed data and long-term correction analysis. Chapter 5 introduces the importance of wind variability for wind farms and details the data sources used in this section, and Chapter 6 displays the LTC process and results.

The final section uses wind farm metadata and reanalysis wind speed data from the second section to predict capacity factor. Chapter 7 investigates two supervised machine learning models that can be trained to predict wind farm power output.

Chapter 8 presents the main conclusions of the thesis and suggestions for future work.

There is perhaps no better demonstration of the folly of human conceits than this distant image of our tiny world. To me, it underscores our responsibility to deal more kindly with one another, and to preserve and cherish the pale blue dot, the only home we've ever known.

— Dr. Carl Sagan [24]

2

Performance Decline of Wind Turbines

Contents

2.1 Wind Turbine Blade Roughness 12
2.1.1 Roughness Characterisation 12
2.1.2 Insect Damage 13
2.1.3 Ice Accretion 15
2.1.4 Rain 17
2.1.5 Airborne particles 18
2.1.6 Impact of Roughness on Performance 19
2.2 Analytical Methods for Modelling a Wind Turbine . . 20
2.2.1 Turbine Control (using BEM) 23
2.2.2 BEM implementation 24

This chapter reviews methods developed to analyse aerodynamic performance decline of wind turbines. The different types of roughness that occur are reviewed, as well as analytical methods for modelling a wind turbine. Thereafter, the implementation of BEM in the context of roughness is reviewed.

Analytical modelling techniques reveal that BEM methods provide a simplified way to model aerodynamical performance effects of blade roughness. Effects of blade roughness vary according to their causes, and reviewing the different sources of roughness in previous studies reveals a range of performance effects. Moreover, analysing turbine control methods through BEM exposes a lack of understanding for the control of aerodynamically degraded turbines.

2.1 Wind Turbine Blade Roughness

Blade roughness has been extensively studied, originating in the context of aeroplane wings. One difference between aeroplane wings and wind turbine blades is that aircraft with damaged wings can be inspected and repaired. Another key difference is that aircraft wings operate at a narrowly defined angle of attack specified by cruise flight, whereas wind turbines experience a much larger range of flow conditions because the blades pitch. This leads to a different design approach and the use of different types of airfoils.

A lot of studies on aeroplane wings are focussed on the effect of rain [25–27] and ice accretion on the airfoil [28–30] because these impact performance in a very short period of time. These studies give insights into the characteristics of airfoils when their shape is changed because of accretion and their lift and drag characteristics are altered as a consequence. However, Blasco et al. [15] note that although many papers have investigated icing penalties for aircraft airfoils, few investigations have focussed on thicker wind turbine airfoils, which has resulted in less availability of validation data for wind turbine icing studies. The change in aerodynamic performance of wind turbine airfoils is important for understanding performance losses as it provides the lift and drag values which are inputs to a BEM model.

Airborne particles may damage or alter the aerodynamic surfaces of the rotor blades and thus change the blade's aerodynamic performance. Roughness may occur on wind turbine blades due to accretion of ice or insects, damage caused by rain, or erosion due to airborne particles. These causes are discussed further below.

2.1.1 Roughness Characterisation

Some attempts have been made to characterise roughness on a wind turbine blade in terms of its height, concentration and location [31]. Roughness height refers to the height of a single roughness element (tangentially from the surface plane), or an average of elements on a surface. Roughness size can be defined by its height from the original blade surface for randomly shaped elements, or by its diameter for spherical elements. To compare multiple cases, the roughness height, k, is

non-dimensionalised by the chord length, c, as $^{k}/_{c}$. Note that this differs from roughness height or length, z_0, which is used in log law models for the change in wind speed with altitude [32]. The collective study from Ohio State University (OSU) researchers in collaboration with the National Renewable Energy Laboratory (NREL) investigated the effect of grit roughness across 13 airfoils using the same $^{k}/_{c}$ value for all tests [33]. The $^{k}/_{c}$ value was informed from measurements of an insect pattern taken from a wind turbine in the field. Roughness density refers to how closely the roughness particles are distributed on the surface, and roughness location refers to where on the airfoil the roughness elements are located. As the contamination particles follow the streamlines of the flow past the airfoil, they tend to build up near the stagnation point, which is usually at the leading edge [31] and is called 'leading edge roughness' (see Figure 1.1).

2.1.2 Insect Damage

Corten and Veldkamp [34] claimed that insects can halve wind-turbine power. Having found that performance was decreasing mainly at high wind speeds (greater than 10 m/s), they postulated that insect damage occurs on the leading edge of the rotor blade, in low wind periods with temperatures above 10° C. These environmental factors are important because this is a time period when the wind turbine is still operating but it is calm enough for insect swarms to fly. In addition to atmospheric conditions, insect presence also depends on altitude, with a rapidly decreasing density from ground level to 150 m [35]. These insect swarms collide with the exposed leading edge of the turbine blade causing roughness. An example of insect damage on an operational wind turbine blade is shown in Figure 2.1 [36].

In Corten and Veldkamp performance drops were noticed after low wind speed weather conditions occurred. A time series study was conducted for four different turbines and found that the power at high speeds decreased markedly after every period of low wind speed, rising again after the blades were cleaned either manually or by rain [34]. This led Corten and Veldkamp to suggest that insect damage was affecting these wind turbines. Insect damage is a region-specific phenomenon



Figure 2.1: Insect damage on a wind turbine blade. Source: Doosttalab [36]

because large insect swarms are required to infiltrate the zone through which the turbine's blades sweep in specific wind and temperature conditions.

Corten and Veldkamp also conducted field experiments to simulate insect damage using two identical 44 metre diameter Nordtank Energy Group (NEG) Micon wind turbines situated within 50 metres of each other to ensure similar flow conditions. One turbine was artificially roughened by installing zigzag tape of maximum thickness 1.15mm on the blades. The roughness was replicated on the leading edge of the blades, as that is where impacts with insects are assumed to occur. They also used stall flags [37] as a flow-separation detector to try to identify why performance drops were only occurring at high wind speeds. At high wind speeds roughness on the leading edge caused flow separation around the blade, leading to higher drag and lower lift (hence lower lift-to-drag ratio). Their results replicated those recorded for an operating wind turbine before and after a low wind period, with the power output being close to equal between the two turbines at low wind speeds, but much lower for the roughened blades at high wind speeds (greater than 10 m/s). The magnitude of the performance drop off also resembled the change seen in after the insect damage (approximately 35% reduction in power).

2.1.3 Ice Accretion

Ice accretion can alter the shape of the blade and thus change its lift and drag characteristics. Cold-climate regions, such as Northern Europe and North America, typically contain a high wind resource, and the higher air density results in greater available kinetic energy flux (more power available) from the wind [15]; however these regions are also subject to atmospheric icing events. Many studies show that wind turbine icing events lead to severe power degradation [38–41], which results in significant losses of revenue for wind turbine operators during the winter season and an annual energy production (AEP) loss of up to 20% [40]. The driving force behind power loss is degradation of blade performance due to the accretion of ice on the blade surface, which increases surface roughness and disrupts the airflow around the blade. An example of icing roughness on an operational wind turbine blade is shown in Figure 2.2 [36].



Figure 2.2: Streamwise icing accumulation on the leading edge of a wind turbine blade. Source: Doosttalab [36]

In an icing event, super-cooled water droplets accrete and solidify on wind turbine blades, causing an alteration in both the shape and roughness of local blade airfoil sections [42, 43]. Switchenko et al. [44] suggest that the impact of increased surface roughness can, at times, be more significant than the size and shape of

the accreted ice. However, they advise that further research is needed on surface roughness effects for particular icing conditions. Bragg et al. [45] show that the general trend for iced airfoils is to have a decrease in lift and increase in drag for various types of ice accretion (e.g. streamwise, horn, or spanwise ridge). They note that the aerodynamic penalty varies between different airfoil geometries for the same icing conditions, making it crucial to not only change the icing shape/roughness but also test different airfoil sections as well.

Blasco et al. [15] extensively studied the effect of ice on a single airfoil (DU 93-W-210), testing different icing events on airfoil aerodynamic characteristics and wind turbine performance experimentally. They explored the effect of icing conditions by first creating ice shapes on an airfoil section using an icing facility in order to obtain aerodynamic data for various icing conditions on a representative utility-scale wind turbine. Ice shapes were then moulded and cast for preservation and the surface roughness attributes of these shapes recorded. Wind tunnel experiments were used to measure airfoil performance, and these data were used to compute the performance of a 1.5 MW wind turbine subject to these various icing conditions.

Six separate 45-minute icing events were simulated to make six different alterations of the DU 93-W-210 airfoil sections. Results showed a 16% loss of airfoil lift at operational angle of attack for freezing fog conditions. Airfoil drag increased by 190% at temperatures near 0° C, 145% near 10° C and 80% near 20° C. For freezing drizzle icing conditions, lift loss and drag rise are more severe at 25% and 220% respectively.

Wind turbine performance was modelled using an external BEM tool, XTurb-PSU [46]. An analysis of the wind turbine aerodynamic loads showed power losses ranging from 16 to 22% for freezing fog conditions and 26% for a freezing drizzle condition. Blasco et al. [15] also explored altering the control mechanism of the XTurb-PSU BEM model by altering when the blade pitches, as will be discussed in Section 2.2.1.

2.1.4 Rain

Corrigan and DeMiglio [47] described rain as having three distinct effects on the performance of a moving airfoil. One effect is a loss in the momentum of the moving airfoil as a result of the mass of raindrops adhering to the airfoil surface which must then be accelerated to the speed of the rotor. The second effect is the additional weight of the water adhering to the airfoil surface. The third effect is the change in the condition of the airfoil surface as a result of the adherence and impact of raindrops. Corrigan and DeMiglio went on to show that the first two effects were insignificant and that the driver of performance loss from rain and in raining conditions is disturbances to the airfoil surface [47]. They conducted experimental tests on the two-bladed Mod-0 horizontal-axis wind turbine with three different rotor configurations to find that light rainfall ($<$ 20 mm/hour) degraded turbine performance by as much as 20% at higher wind speeds ($>$ 6 m/s), while heavy rainfall ($>$ 30 mm/hour) degraded performance by as much as 30%. This was an early analysis that was heavily based on research of the aerodynamic penalties of heavy rain on a landing aircraft [25]. Consequently, performance loss was measured during rain events and did not investigate the ongoing impact of rain-induced damage on wind turbines.

Currently, the predominant consideration as to the impact of rain on wind turbines is its ability to erode surfaces on the outermost section of the blade due to high impact speeds [48]. Typical blade tip region velocities are in the 90 to 110 metres per second (m/s) range [49] which is fast enough to cause erosion through impact damage of rain droplets in this region [49, 50]. Multiple studies have concentrated on rotor design and materials to mitigate this effect [31, 50, 51]. Chen et al. [50] suggest that a comprehensive protective coating can be obtained by combining and modifying resins with different properties. In their review, Dalili et al. [51] suggest that a single surface engineered coating designed to reduce accumulation of ice and insect damage, and to protect from erosive deterioration, is a desirable solution. Sagol et al. [31] also mention the importance of surface protection and finishes, such as nonstick coatings or plastic films.

2.1.5 Airborne particles

Small particles (dust, salt, sand, etc.) that can be transported by the wind to the height of the wind turbine rotor can erode the blade surface. As these particles hit the rotor blade, the smoothness of the surface is perturbed, mainly along the leading edge. An example of an eroded wind turbine blade is shown in Figure 2.3 [52].

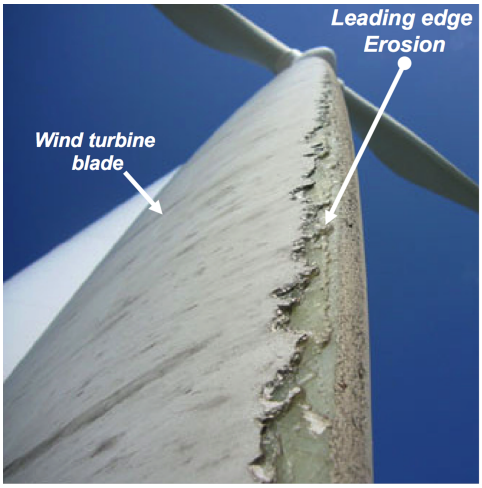


Figure 2.3: Leading edge erosion on a wind turbine blade. Source: Verma et al. [52]

Khalfallah and Koliub [53] show the effects of dust accumulation on a 300-kW pitch-regulated wind turbine over a 9-month recording period (without cleaning). The dust accumulation pattern saw a concentration of particles on the leading edge and covering between 10–20% of the blade surface. A higher concentration of particles were also observed on the tip of the blade. These areas are the most contaminated due to the exposure to the incoming flow as well as the speed increasing with radius causing a bigger impact. In Khalfallah and Koliub's study, the dust roughness area on the blades accumulated on the leading edge and extended up to 20% of the chord line, and the mean power loss varied from 6.5% for 0.05 mm diameter dust grains, to 38% when the diameter of dust particles was 0.3 mm [53].

With the expansion of wind turbines offshore, a increasingly important form of erosion comes from salinity in the air. The salty air and water can be corrosive and can erode the leading edge of the blades [54–56].

Erosion is often included in the more general topic of wind turbine blade roughness, as it can have various causes and is hard to quantify.

2.1.6 Impact of Roughness on Performance

A large number of experimental tests have been completed detailing the effects of roughness on wind turbine airfoil performance. Janiszewska et al. [57] show the S814 airfoil has a 25% decrease in maximum lift and a 60% increase in minimum drag when exposed to roughness, as part of a larger NREL study that did the same tests on 13 separate airfoils. Sareen et al. [58] showed the effect of erosion on airfoils in wind tunnel experiments, producing a decrease in lift of 17% and an increase in drag up to 500% for a heavily eroded case.

The previously mentioned research conducted at OSU in conjunction with NREL, which used wind tunnel experiments to show lift and drag profiles of specific airfoils when clean and roughened [33], has a publicly available dataset accompanying its results. These data were used by Munduate and Ferrer [59] to examine the modelling capability of Xfoil [60] by comparing it to CFD analysis. Xfoil is an analysis and design system for subsonic isolated airfoils. By comparing experimental data to simulated results, Munduate and Ferrer showed that Xfoil was an insufficient modelling tool for predicting wind turbine blade roughness effects. Munduate and Ferrer also show that airfoil inputs can be found computationally; however, transition between laminar and turbulent flow, as well as roughness effects, are difficult to calculate this way. Therefore, experimental data is often used.

Schramm et al. [17] used numerical methods to simulate roughness effects on a wind turbine over all operational wind speeds by using CFD methods to develop aerodynamic polars which were then utilised to model a turbine through BEM methods (FAST v8 [61]). However, they only analysed the NREL 5 MW reference turbine [62] with FAST v8. Pechlivanoglou et al. [16] used an in-house BEM simulation method based on Hansen [63] to produce a power distribution against wind speed for clean and roughened airfoils on a 1.5 MW turbine. Blasco et. al [15] modified the XTurb-PSU BEM model to only pitch a roughened (ice accretion)

blade once the maximum torque is reached (4m/s after rated flow speed), instead of when the rated flow speed occurs. This is because a rough-bladed turbine will not be at rated power at the rated flow speed, whereas for a clean-bladed turbine, these points coincide. The modified BEM model ensured that maximum turbine power was maintained after the rough-bladed turbine reached rated power despite the presence of aerodynamic penalties from roughness, although the performance improvement from this control method (when compared to the unmodified XTurb-PSU BEM model) was not quantified. This was the only investigation of a turbine control strategy that accounted for roughened blades in the reviewed literature.

2.2 Analytical Methods for Modelling a Wind Turbine

To understand the impact of roughness on wind turbine performance, the roughened airfoil lift and drag data must be modelled on a theoretical rotor. To do this, semi-analytical methods can be used.

Theoretical modelling of energy extraction from fluid flows originated with Rankine [64] and Froude [65] who described the flow field around a ship propeller. This led to actuator disk theory by replacing the propeller with a uniformly porous disk. This simplified the problem by converting it to a linear momentum problem with no rotating elements. When this theory was combined with incompressible fluid flow assumptions, the steady state flow around the disk could be found. It was later shown by both Lanchester [66] and Betz [67] that the same idea could be applied to an inverse propeller in an infinite flow field. The key result was that it is not possible to extract all of the undisturbed kinetic energy flux upstream of the turbine; to attempt to do so would cause the flow to divert around the turbine as it becomes an increasingly large resistance to the flow, thereby limiting the overall energy capture. A consequence of this is the Lanchester-Betz limit, $^{16}/_{27}$, which is the maximum fraction of the undisturbed kinetic energy flux that can be extracted from an unbounded flow stream.

These insights formed the basis of blade element momentum theory, which was further developed by Glauert in 1926 [68]. BEM theory is still extensively used in wind turbine conceptual design and aerodynamic analysis. The rotor area is divided into multiple concentric rings, which allows the combination of a control volume conservation of momentum analysis with a blade sectional analysis [63, 69, 70]. Various additions have been used to augment BEM theory and account for neglected physics, such as the effect of a finite number of blades and skewed wakes [71].

The Lanchester-Betz limit is a theoretical upper limit to the power extraction of a turbine in an infinite flow field, but in reality there are several factors preventing real turbines from achieving it. BEM theory is a steady flow approximation to the flow conditions through a turbine, and it assumes that only the axial and tangential velocities are affected by the energy extraction, neglecting the radial velocity changes that occur. Furthermore, there are aerodynamic, mechanical and electrical losses which limit turbine power. However, BEM theory and the Lanchester-Betz limit provide a simple benchmark against which real turbines can be compared.

BEM theory requires a numerical solution as characteristics – chord, pitch, airfoil section – change along the blade radius, for each blade element. The most common approach is to form the BEM equations as a function of two variables, the axial induction factor (a), which represents the change in wind speed from the inflow velocity to the velocity through the turbine, and a tangential induction factor (a'), which represents the added tangential wind speed due to the rotating blades. The axial and tangential velocities are perpendicular to each other and thus create the velocity triangle seen in Figure 2.4, where Ω is the rotational speed of the rotor, r is the radius of the blade element from the hub, U_∞ is the free stream wind speed, and V is the relative wind speed. The induction factors can therefore be defined through the relative wind speed:

$$W = \sqrt{U^2(1-a)^2 + \Omega^2 r^2 (1+a')^2} \tag{2.1}$$

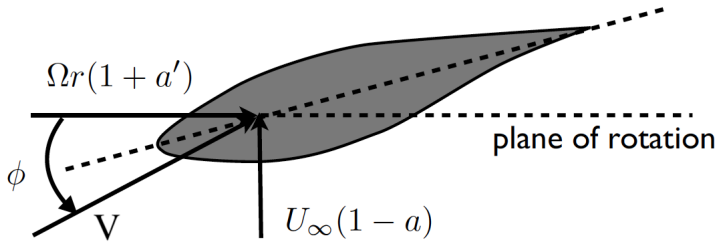


Figure 2.4: The velocity triangle overlayed onto an airfoil. Adapted from Ning [71].

The velocity triangle occurs across all blade elements, and is contextualised in Figure 2.5, which shows a single blade element sweeping out an annular ring, with the definitions of the induction factors indicated.

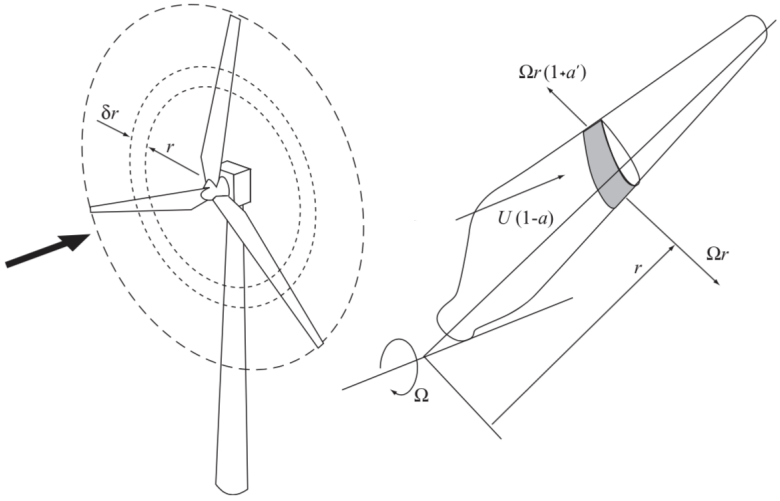


Figure 2.5: A schematic showing a blade element sweeping out an annular ring. Source: Wind Energy Handbook [70].

The induction factors are found through a fixed-point problem for each radial rotor annulus and are then used to calculate turbine performance. Most textbooks [69, 70] solve this equation through an iterative approach to converge on final values for each blade element. However, convergence is not guaranteed in all areas of the design space [71]. Even if the final solution is solvable, the iteration from initial guess to converged value may still pose a problem. This limitation may not always hinder solving the model (depending on the rotor and airfoil inputs, as well as the

operating conditions) but it is not as computationally robust as desired.

Ning [71] presented a guaranteed convergence method to solve the BEM equations. This is accomplished by converting the two-variable, two-equation fixed-point problem into an equivalent single variable, root-finding problem for which robust optimisation/solution algorithms exist [71]. Ning does this by parametrising the induction factors in terms of the inflow angle (ϕ), which is a function of both the axial and tangential induction factors, as shown in Figure 2.4. Ning uses Brent's method [72] as the root-finding algorithm of choice. This algorithm retains the guaranteed convergence properties of bisection (assuming that an initial bracket encompassing the solution is given), but converges much faster (superlinearly in most cases) [72]. This methodology was adopted as the BEM model implementation for the current research.

BEM theory requires both rotor and airfoil data inputs to work. Rotor solidity ($\sigma_r = Bc/2\pi r$, where B is the number of blades) and twist (γ_t) values along the blade span are used for the rotor, and lift and drag values as functions of angle of attack are required for the airfoil. Blade roughness can be modelled in BEM theory by using roughened airfoil data as the airfoil inputs and observing performance differences against the case with clean airfoil data.

2.2.1 Turbine Control (using BEM)

To analyse the change in power output of roughened blades on a wind turbine, simulations must be conducted across all typical operating points (tip speed ratios) of interest. A conventional variable-speed, variable-blade, pitch-to-feather configuration wind turbine is the commonly investigated turbine type. The NREL 5 MW reference turbine [62] operates by assuming the turbine operator only has access to: the rotational speed of the blades, the root pitch angle of the blades, and the torque on the generator. This turbine operates using two basic control systems: a generator-torque controller, and a full-span collective blade-pitch controller. The two control systems are designed to work largely independently, in the below-rated and above-rated wind speed ranges respectively – Region 2 and 3 in Figure 2.6. The

goal of the generator-torque controller is to maximise power capture below the rated operation point and the goal of the blade-pitch controller is to regulate generator power above this point. This is the underlying method for most BEM models that model turbine power over a range of wind speeds. Modelling turbine control using BEM simplifies the transition between the regions shown in Figure 2.6 with abrupt changes to the power curve, however in reality there is a smoothed conversion period between regions. As mentioned, Blasco et al. [15] modified the blade-pitch controller such that pitching occurs once rated torque is reached, in order to maximise power.

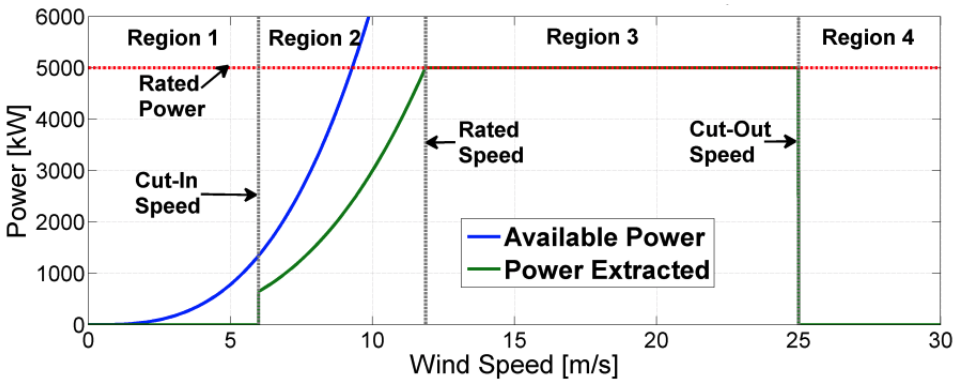


Figure 2.6: An example power curve, and operating regions for a 5 MW turbine. Adapted from Aho et al. [73]

2.2.2 BEM implementation

In this thesis, the BEM method will be used to examine the effects of airfoil roughness on wind turbine performance in Chapters 3 and 4. The BEM method used for this simulation is a single-parameter, guaranteed-convergence formulation proposed by Ning [71]. The model will use airfoil lift and drag data (both clean and roughened).

Prandtl's tip-loss model [74] is included to account for the influence of discrete and 3-D blade effects that cause a spanwise variation in blade loads and are not directly modelled in the 2-D momentum and blade element equations used in BEM theory. Buhl's thrust correction [75] for the turbulent wake state that couples momentum theory with Glauert's empirical model for heavily loaded turbines with hub and tip losses is also included.

The model has been validated with reference to BEM-based studies, see Kelly et al. [76], Kelly et al. [77], and Vogel et al. [78] for further details.

Climate change is the greatest threat to our existence in our short history on this planet. Nobody's going to buy their way out of its effects.

— Mark Ruffalo

3

Aerodynamic Performance Degradation and Mitigation of a Single Roughened Turbine

Contents

3.1 Introduction 28
3.2 Method 30
3.2.1 Computational Method and Available Airfoil Tables . . 30
3.2.2 Rotor Design 34
3.2.3 Turbine Control 36
3.2.4 Performance Loss 42
3.3 Results 43
3.3.1 Rotor Performance 43
3.3.2 Turbine Performance Comparisons 45
3.3.3 Spanwise Effects of Roughness 49
3.3.4 Roughness Effects on Airfoils 52
3.3.5 Energy Loss 55
3.4 Conclusions 56

The work presented in this chapter has been previously published in *Wind Energy* titled "Impact and mitigation of blade surface roughness effects on wind turbine performance" [76].

This chapter presents a numerical study of the effects of blade roughness on wind turbine performance and annual energy production, and how these effects may

be partially mitigated through improved control. Three rotors are designed using the NACA4415, S801 and S810 airfoils, and blade element momentum theory is used to model wind turbine behaviour. The aim is to explore how rotors designed with different lift and drag characteristics are affected by roughness. The aerodynamic lift and drag data for the clean and roughened airfoils are taken from previous experimental work. In Chapter 4, these data show that surface roughness leads to a decreased airfoil lift coefficient and an increased drag coefficient across an angle of attack range typical for large wind turbines, which will clearly lead to decreased turbine performance. Three separate control methods are considered for wind turbines with roughened blades operating at each of four candidate wind sites with different wind speed distributions. Results show that, compared to clean rotor blades, the roughened blades lead to a performance (energy yield) drop in the range of 2.9–8.6% for a torque-based control strategy. A control-based performance recovery strategy, in which the controller gain coefficient is re-optimised, increased roughened rotor annual energy production by 0.1–1.0%.

3.1 Introduction

Roughness impacts on turbine performance can be significant. Corten and Veldkamp [34] showed in field tests that roughness due to insect contamination can result in power losses of over 25%. Bragg et al. [43] show that turbine performance can decrease by up to 20%, while Corrigan and DeMiglio [47] concluded that heavy rain can degrade performance by as much as 30%. Khalfallah and Koliub [53] also show power losses between 6.5%–38% for different diameter dust particles on turbine blades.

Many studies provide insight into the aerodynamic effects of airfoil roughness [57, 58] but few investigate the overall effect on wind turbine power performance. Blasco et al. [15] analysed the effect of icing on an airfoil using wind tunnel experimental testing. They used BEM software XTurb-PSU [46] to conclude that the changed aerodynamic performance of the airfoil implied a 16–22% reduction in power for a wind turbine in freezing fog conditions and 26% for a freezing

drizzle condition. They also showed that modifying the blade pitch controller to start blade pitching at a higher wind speed resulted in greater power production when the blades were roughened.

The literature on the effect of roughness on wind turbine performance predominantly focuses on the impact on C_P and not AEP. However, Schramm et al. [17] use numerical methods to simulate roughness effects on the NREL 5 MW reference wind turbine across a range of wind speeds using CFD methods to evaluate airfoil aerodynamic coefficients which were then used to model the turbine through a BEM model (FAST v8 [61]). Roughness was represented by deforming the leading edge of the airfoil (to simulate erosion), and the airfoil performance coefficients were determined through RANS CFD simulation with the k-ω SST model, where a wall roughness function was used. While all operational wind speeds were considered, they did not investigate how turbine control might be used to mitigate performance losses.

Similarly, Pechlivanoglou et al. [16] used an in-house BEM simulation method based on Hansen [79] to produce a power distribution as a function of wind speed for a 1.5 MW turbine. However, their study was primarily focussed on the effects of roughness on airfoil performance through wind tunnel experiments, and they only provided a limited BEM analysis of the consequences for wind turbine performance. Bak et al. [80] investigated the effect of 5 different roughness types (bumps, erosion, overhang, etc) using CFD across three different wind climates and different maximum tip speeds. They show AEP losses of 0.5–4%, and also that increasing the maximum tip speed (for both clean and rough cases) decreased the magnitude of losses. This is an example of a potential design strategy for a clean rotor to mitigate the effect of roughness when it occurs. A different approach is used in this chapter, which develops a control strategy to mitigate the effect of roughness regardless of the design tip speed of the clean rotor.

Turbine control has the potential to reduce the negative impacts of turbine roughness, either by reducing the likelihood of blade damage occurring in the first place, or by adjusting turbine performance once damage has been incurred. Impact

damage roughness events such as insect damage and rain can result in permanent damage. However, these events can also be mitigated through the implementation of erosion-safe operation (reducing the tip speed during rain events or insect swarms to minimise impact damage), see Bech et al. [81], Tilg et al. [82] and Hasager et al. [83]. Erosion-safe operation has a smaller impact on AEP than permanently damaged blades because the performance losses are only temporary. The focus of this chapter is on permanent roughness, and a novel control strategy to mitigate the impact of roughness is proposed.

The aim of this chapter is to investigate the potential for turbine control methods to mitigate losses in annual energy production due to the development of roughness on the blades. The performance of roughened airfoils, tested in research conducted at OSU for the NREL [84], is used to model a single wind turbine simulated using BEM methods. The effects of roughness on turbine performance have been explored for three turbine designs, each designed using a different airfoil profile and exhibiting different aerodynamic characteristics. In all cases roughness was applied uniformly along the entire span of each rotor blade. Three turbine control strategies have been investigated with an in-house analytic BEM model; two widely adopted in the literature and a novel third control strategy that seeks to mitigate some of the effects of roughness on turbine performance and annual energy production.

3.2 Method

3.2.1 Computational Method and Available Airfoil Tables

The BEM method used to examine the effects of roughness on wind turbine performance is detailed in Section 2.2.2.

Although steady wind turbine operation generally corresponds to modest angles of attack, α, along most of the blade span (0° to 10°), good numerical convergence of the BEM model requires smooth and continuous lift and drag data over a much wider range of α. Thus, some processing of experimental data is inevitably required.

Airfoil aerodynamic characteristics (lift, drag) must be specified in order to use a BEM simulation for wind turbine performance. Airfoil data for both aerodynamically

smooth and rough conditions, published by OSU researchers under contract from the NREL [84], are used in this chapter and were described in Section 4.2.2. The data were collected using wind tunnel experiments to determine the lift and drag profiles of specific airfoils [33]. Lift and drag forces are converted to non-dimensional parameters, the sectional lift and drag coefficients, c_l and c_d respectively.

The advantage of using airfoil data from one study is that an equivalent level of roughness is applied in each test. As discussed in Section 4.2.2, a standard roughness pattern was developed from a molded insect pattern taken from a wind turbine in the field [57] and the non-dimensional roughness height was $k/c = 0.0019$ [57].

All airfoils were tested in the OSU Aeronautical and Astronautical Research Laboratory (AARL) 3 x 5 foot subsonic wind tunnel under steady flow conditions [57]. Data were acquired and processed from 60 surface pressure taps, four individual tunnel pressure transducers, an angle of attack potentiometer, a wake probe position potentiometer, and a tunnel thermocouple [57].

The OSU/NREL wind tunnel experiments were conducted across a range of chord-based Re (0.75 to 1.5 million) for each airfoil. Reynolds number is defined as the ratio of inertial forces to viscous forces within a fluid. Only the data from the highest Re experiment was used in the present study (1.25 or 1.5 million depending on the airfoil) as the airfoil boundary layer flow is expected to be post-transitional and therefore most representative of a utility-scale wind turbine.

The range of drag data needed to be extrapolated because drag measurements only occurred at low angle of attack values (approximately $-10°$ to $10°$). This is because the drag value was calculated from pressure measurements in the wake and could not be accurately calculated when there was a lot of turbulence, as arises in separated flow at high angles of attack [85]. Consequently, the lowest and highest recorded angles of attack were uniformly extended beyond this angle of attack range. This simplistic approximation, required to ensure convergence of the BEM model, did not affect the final predicted rotor performance in the conditions simulated in this study.

Airfoil Selection

Table 3.1: Summary of maximum lift to drag ratio and corresponding angle of attack and chord-based Reynolds number for three selected airfoils from the OSU study [33]. $(c_l/c_d)^*$ indicates the maximum value and α^* is the corresponding α.

Airfoil	**Condition**	$(\mathbf{c_l/c_d})^*$	α^* **[degrees]**	$\mathbf{c_{l,max}}$	α **range [degrees]**	**Re** $[\times\mathbf{10^6}]$
NACA4415	Clean	102.45	5.8	1.29	-10.2–12.2	1.5
NACA4415	Rough	52.64	4.2	1.09	-10.2–8	1.5
S801	Clean	129.75	3.3	1.44	-8.2–11.2	1.5
S801	Rough	60.50	4.5	1.27	-10.2–10.2	1.5
S810	Clean	92.62	4.1	1.14	-10.2–10.2	1.5
S810	Rough	36.84	5.8	1.01	-10.1–10.1	1.5

The three airfoils listed in Table 3.1 were chosen for analysis based on three factors: $(c_l/c_d)^*$ for both the clean and rough airfoils as well as the difference in magnitude of these values; the corresponding change in α^* between the clean and rough airfoils; and the shape of the c_l/c_d vs α graph. A large lift to drag ratio is desired to maximise torque production and thus achieve efficient turbine operation. The optimal angle of attack value, α^*, and the c_l/c_d vs α curve are important because a rotor designed to operate at α^*_{clean} could operate significantly away from $(c_l/c_d)^*_{rough}$ when the airfoil transitions from smooth to rough conditions. This has the potential to exacerbate the performance decline in roughened conditions due to the combined effects of roughness reducing the c_l/c_d ratio and operating away from the maximum c_l/c_d condition. Figure 3.1 shows c_l/c_d vs α as well as the aerodynamic polars for the NACA4415, S801 and S810 airfoils under clean and rough conditions. Note that the vertical bars (constant c_d) at the ends of the c_l/c_d curves are present due to the extrapolation of the final drag measurement to higher magnitude angle of attack values.

The chosen airfoils all had high clean airfoil $(c_l/c_d)^*$ values with differing α^* positions and differing c_l/c_d vs α characteristics (broadness of lift-to-drag peak). It was postulated that these differences would give insight into the performance

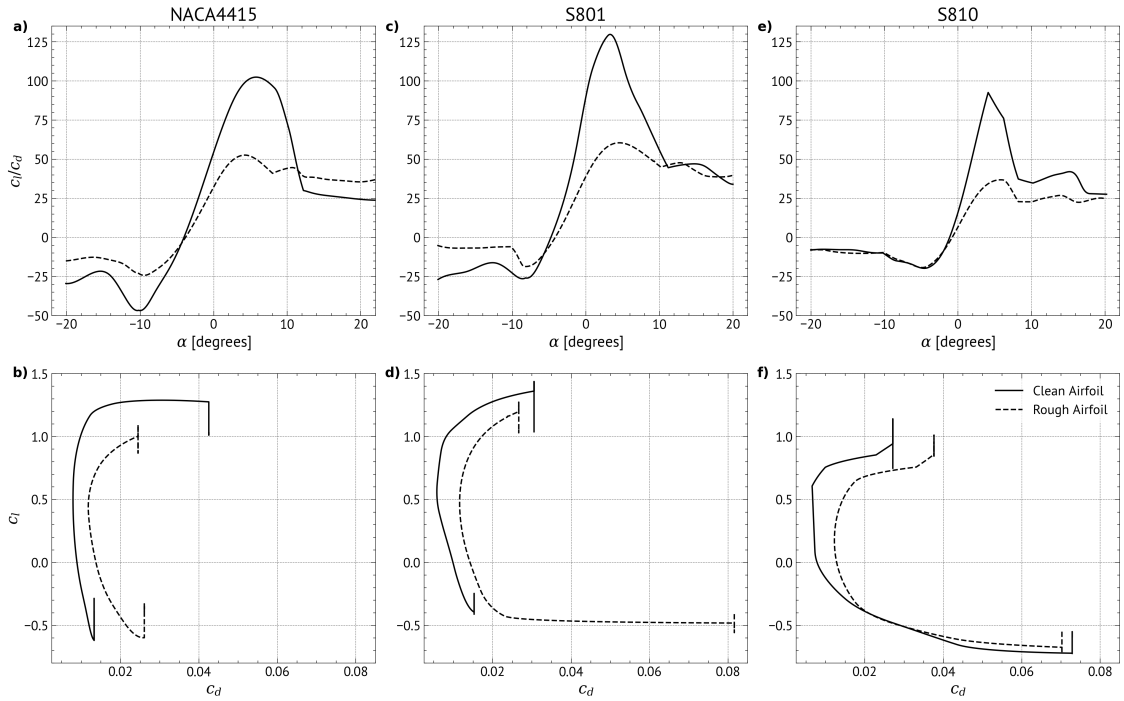


Figure 3.1: Lift to drag ratio and aerodynamic polars for the three selected airfoils (NACA4415, S801, S810). Plots (a) and (b) present the NACA4415 airfoil data, (c) and (d) present the S801 airfoil, (e) and (f) present the S810 airfoil.

differences experienced by different types of rotor as their airfoils transition from clean to rough.

The S801 and S810 airfoils were designed for use on small wind turbines (20–400kW) and to be resistant to leading edge erosion [86]. They were designed to have a $c_{l,max}$ which is largely insensitive to roughness (as shown in Table 3.1), however c_d may still increase significantly. These airfoils were designed to have laminar to turbulent transition on the suction surface occur close to the leading edge just prior to reaching $c_{l,max}$. This is not expected to significantly change at higher chord-based Reynolds number. Using these airfoils on larger rotors may not be as representative of airfoils used on utility scale turbines, but they do allow for a data-driven approach to analyse the effect of roughness. For more details on the original experimental studies completed on these airfoils, see Hoffman et al. [85], Ramsay et al. [87] and Ramsay et al. [88].

3.2.2 Rotor Design

Three rotors were designed to investigate the effects of roughness on turbine performance given the different characteristics of the selected airfoils. For simplicity in evaluating the impact of airfoil roughness on turbine performance, a constant airfoil profile was used throughout the span of each rotor's blades. A conventional variable-speed, variable blade pitch-to-feather configuration wind turbine was assumed, and the rotors were designed to meet the following parameters:

- Design Tip Speed Ratio $\lambda^* = 10; \lambda = R\omega/U$, where ω is the rotational speed of the turbine, R is blade radius, and U is wind speed.
- Maximum rotational speed $\omega_r = 2\text{rad/s}$
- Diameter = 100m
- Hub Radius = 5m
- Cut-in wind speed $U_{ci} = 3\text{m/s}$
- Rated wind speed $U_r = 10\text{m/s}$
- Cut-out wind speed $U_{co} = 20\text{m/s}$
- Approximate rated power $P_r \approx 2.4$ MW
- Design coefficient of performance $C_P^* \approx 0.5$; C_P is defined in Equation 4.6.

C_P is the power coefficient and is the ratio of the shaft power generated to the kinetic energy flux of the undisturbed flow through the rotor plane, and A is the turbine swept area. These parameters are chosen to be representative of the size and performance of wind turbines that are currently in the field. Current similar sized commercial offerings include the Vestas V120-2.2 MW [89][1] and Siemens SG 2.6-114 [90][2].

[1]The Vestas V120-2.2 is a 2.2 MW wind turbine with a 120 metre rotor diameter, and is available at hub heights of 80, 92 and 122 metres.

[2]The Siemens SG 2.6-114 is a 2.625 MW wind turbine with a 114 metre rotor diameter, and is available at hub heights of 63, 68, 75, 80, 88, 93 and 125 metres.

Two planform parameters were used to characterise the rotor blade design: blade twist distribution and chord length distribution. These were set via an iterative Blade Element Momentum solution in which the rotor is divided into a number of concentric annular elements. The blade twist distribution was set such that the angle of attack was equal to α^* for the clean airfoil along the blade, thereby approximately maximising torque production with the airfoils in the clean state. The chord distribution was determined by setting a target local thrust coefficient of $C_x = 2.0$ along the entire length of each blade, which corresponds to the optimal local thrust coefficient from Betz theory for the optimal rotor in an unbounded flow. The local thrust coefficient is defined as:

$$C_x = \frac{\delta F_x}{(1/2)\rho U_x^2 \delta A} \tag{3.1}$$

where δF_x is the streamwise force on a corresponding annular element of area δA, and U_x is the streamwise wind velocity at the rotor plane for that annular element.

The rotor solidity (σ_r) and twist (γ_t) of the rotor designs is shown in Figure 3.2. Because this thesis focuses on the aerodynamics, rather than on the practical structural requirements of blade design, the inboard taper of the blade was not modified to transition it to a cylindrical section. Consequently, the sectional Reynolds number is approximately constant along the blade span due to the contrasting effects of increasing incidence velocity and decreasing chord length as functions of radius. It should be noted that near the hub, the solidity is relatively large and would usually be replaced by thicker airfoils and a cylindrical section to provide the necessary structure to support the rotor blades. However, the focus of this chapter is on the impact of airfoil roughness on rotor performance, and hence a single airfoil cross-section is used for the entire blade span to simplify the following analysis. Relatively little torque and therefore power is generated near the blade root so this simplification is considered unlikely to significantly alter the conclusions of this roughness study.

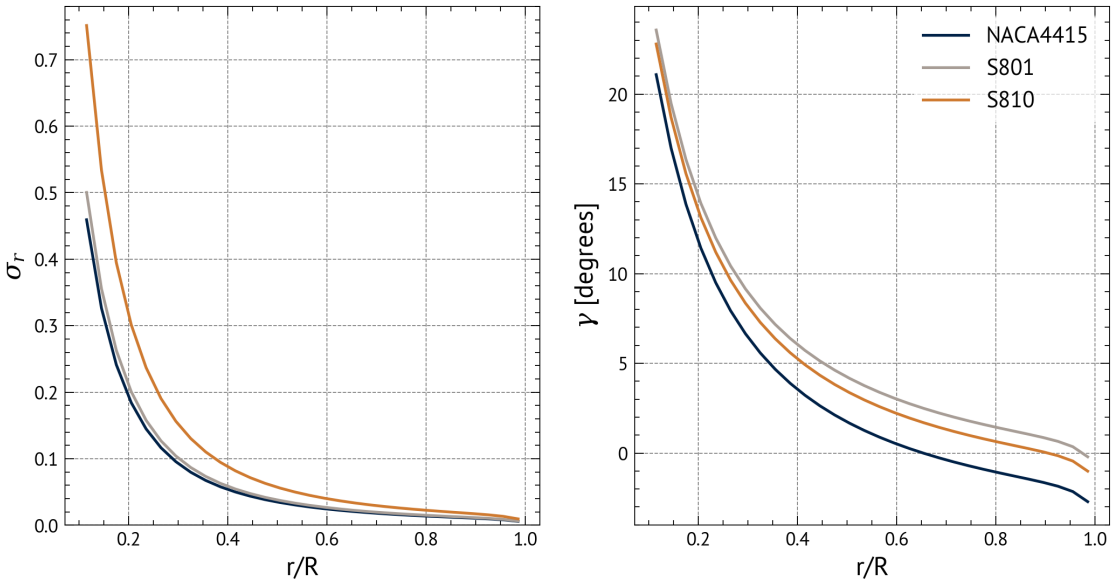


Figure 3.2: Spanwise distribution of rotor solidity σ_r (left) and blade twist angle γ_t (right) for rotor designs using the NACA4415 (red), S801 (blue) and S810 (green) airfoils.

3.2.3 Turbine Control

Rotor performance loss was evaluated by comparing turbine performance with the clean and roughened airfoils. For the roughened case, the roughened airfoil data were applied along the entire blade span, implying a uniform roughness distribution. This is clearly an approximation and it is known that blades tend to roughen more rapidly towards the blade tips where they experience higher impact velocities. Alternative roughness distributions may be more representative than this first-order approximation, but are beyond the scope of the project work. The completely roughened blade assumption is however useful as it provides an upper bound on the detrimental effects of blade roughness.

The reduced lift and increased drag that arises as roughness develops results in a reduction in rotor torque and thus power at a given wind speed. Consequently, higher wind speeds or higher turbine rotational speeds are required to recover the torque to the levels observed with clean airfoil conditions. The reduction in torque interacts non-linearly with rotor control mechanisms, which may act to modify the behaviour of the rotor and therefore impact on turbine performance.

Turbine performance was evaluated across the full range of wind speeds from cut-in to cut-out to assess the impact of roughness on overall energy production. It is assumed that turbine control is based only on the following data:

- Rotational speed of the rotor (ω)
- Root pitch setting angle of the blades (β)
- Generator torque (τ_{gen})

The generator torque, τ_{gen}, can be different to the expected aerodynamic torque on the rotor blades (τ_{aero}), which, in the present analysis, is determined through a BEM solution. This is an important distinction as generator torque can be directly measured, whereas the aerodynamic torque is inferred [91].

These control parameters are the same as those chosen to control the NREL 5 MW reference turbine [62]. The NREL turbine operates using two control systems: a generator-torque controller and a full-span collective blade-pitch controller. The two control systems work independently, for the most part, in the below-rated and above-rated wind speed ranges, respectively. The goal of the generator-torque controller is to maximise power capture below the rated operation point and the goal of the blade-pitch controller is to regulate generator speed and rotor torque above this point. The transition from below- to above-rated flow speed control systems typically starts from around 1 m/s below the rated flow speed. For simplicity, the transition between rotor control below and above rated flow speed is not considered in this chapter, but this simplification is not expected to significantly alter the conclusions of this study.

Two control paradigms were considered in the present work: referencing turbine control to the rotational speed ω; and referencing turbine control to generator torque τ_{gen}. In a BEM model, U is known and is used as a reference for the ω control paradigm. The blade-pitch controller was employed above the rated flow speed in both cases. An optimised generator torque method was also investigated, and the methodology for each is explained below. For all simulations, it was assumed

that the 'cut-in' torque was equal to the torque generated by the rotor with clean airfoil blades at the cut-in wind speed (3 m/s).

Figure 3.3 shows a schematic power-wind speed plot for all control methods and is presented to aid the explanation of the control methods. A clean and rough operation curve are shown for each control method, with exaggerated differences to highlight changes. All control methods with a clean blade yield identical power curves and are simply labelled "CMclean". The control methods are outlined below, and an overview is provided in Table 3.2.

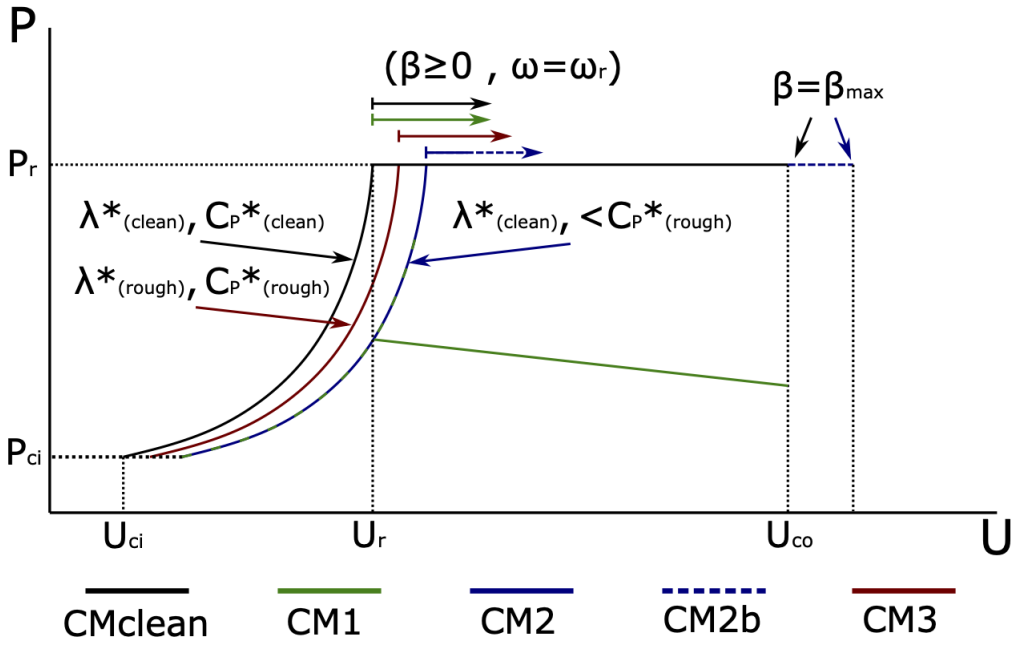


Figure 3.3: Schematic of the different control strategies used for both clean and rough rotors. All control methods with a clean blade yield identical power curves and have been labelled "CMclean".

Control Method 1 (CM1): Wind Speed Referenced

A reference case was established in which the turbine rotational speed and blade pitch angle are varied at the same rate with the flow speed equal to that of the clean turbine case. In reality, this would require the wind speed to be measured and provide the sole basis of the turbine control, without any influence of observed generator torque. This is consistent with existing power-wind speed control paradigms [70].

In CM1, the tip speed ratio λ is set to $\lambda^* = 10$ between the cut-in and rated flow speeds, and the rotor performs sub-optimally and below the performance level attainable for the roughened rotor, i.e. $C_P < C^*_{P,\ rough}$. Once the rated flow speed is reached $\omega = \omega_r$, the blade pitch angle β increases with flow speed at the same rate as in the clean turbine case. The turbine shuts down at the cut-out flow speed U_{co}.

As the rotor torque is reduced by the influence of roughness from the clean case, the maximum torque achieved with this method is less than that at rated conditions for the clean rotor ($\tau_{max} < \tau_r$), and thus maximum rotor power is necessarily less than the rated power ($P_{max} < P_r$). As the rough rotor follows the pitch profile of the clean rotor, it is likely that $\partial P/\partial U < 0$ beyond rated.

Control Method 2 (CM2): Generator Torque Referenced

CM2 is an optimal torque control (OTC) technique that is the standard Maximum Power Point Tracking (MPPT) algorithm that is implemented in commercial wind turbines [92]. OTC is favoured in this study over other MPPT techniques such as hill-climb search (HCS) and sliding mode control (SMC) because of its popularity in commercial wind turbines and the steady-state nature of BEM which does not allow for control responses to unsteady events such as turbulent flow.

In CM2, at below-rated flow speed, the generator torque is brought into balance with the aerodynamic torque in steady-state to achieve:

$$\tau_{gen} = \tau_{aero} \tag{3.2}$$

Note that this is not the case during transient or flow speed-varying operation, and it is the equalisation of these torques that is sought through adjustment of rotor speed or collective pitch. The above equilibrium condition yields the following relationship:

$$\tau_{gen} = \left(\frac{1}{2}\rho\pi R^5 \frac{C_P^*}{\lambda^{*3}}\right)\omega^2 = K\omega^2 \tag{3.3}$$

where C_P^* and λ^* are defined by the optimal operation point of the turbine. K is the generator torque constant that is pre-determined for the case of clean blades (K_{clean}). If the turbine performs such that it generates a power coefficient C_P^* at

λ^* then the controller (Equation 3.3) works to adjust the rotor's speed in response to changes in the generator torque which is matched to the aerodynamic torque experienced by the rotor (or computed via BEM as in the current analysis).

In CM2, the rotor under performs relative to the clean condition with $C_P < C^*_{P\ clean}$ and so the aerodynamic torque is less than expected at any given flow speed below rated, and the rotational speed, set via Equation 3.3, climbs more slowly. Consequently, in CM2, rated power is reached at a higher flow speed than in the clean condition. When a high enough flow speed is reached so that $\omega = \omega_r$, blade pitching commences to limit generator torque so that $\tau_{gen} = \tau_r$. This occurs at a higher flow speed than in the clean case, but once $\omega = \omega_r$, power yield is unaffected and $P = P_r$ until the cut-out flow speed is reached, at which point the rotor shuts down.

Control Method 2b (CM2b): Generator Torque Referenced with Full Pitch

CM2b is similar to CM2, except that the turbine shuts down when, under conditions of above rated operation and maximum torque, i.e. $\tau_{gen} = \tau_r$, maximum blade pitch angle, $\beta = \beta_{max}$, is reached. This means that the rotor will continue operating to wind speeds slightly higher than U_{co} due to the lower blade pitch angle required at any given wind speed above rated in the roughened case.

Control Method 3 (CM3): K-Optimised Torque

The third control method uses the generator constant K as an additional degree of freedom to improve turbine performance. Adjusting K to account for roughness effects enables the controller to modify the rotor's operation to maximise rotor efficiency. A roughness optimised K value, K_{rough}, was found by seeking to maximise AEP, as described below.

By adjusting K for roughness, the turbine will seek to operate at the maximum power coefficient achievable with roughened blades $C^*_{P\ (rough)}$. This will exceed the power coefficient that would be achieved without adjusting the controller

Table 3.2: Overview of control methods.

	$\mathbf{U < U_r}$		$\mathbf{U > U_r}$		
Control Method	**Target**	**Control Parameter**	**Target**	**Control Parameter**	**Notes**
CM1	$\lambda = \lambda^*_{clean}$	$\omega(U)$	$\omega = \omega_r$	$\beta(U)$	Pitching starts at ω_r. Clean pitch followed for both cases.
CM2	$\lambda = \lambda^*_{clean}$	τ_{gen}	$\tau_{gen} = \tau_r$	β	Pitching starts at $\tau_{gen} = \tau_r$. Cut out speed enforced at 20m/s.
CM2b	$\lambda = \lambda^*_{clean}$	τ_{gen}	$\tau_{gen} = \tau_r$	β	Pitching starts at $\tau_{gen} = \tau_r$. Cut out speed enforced at β_{max}.
CM3	$\lambda = \lambda^*_{rough}$	τ_{gen}, K	$\tau_{gen} = \tau_r$	β	K optimised for $\tau_{gen} < \tau_r$. Pitching starts at $\tau_{gen} = \tau_r$. Cut out speed enforced at β_{max}.

for roughness, and hence, the power curve rises faster with flow speed for CM3 than for CM2 (unadjusted controller with rough blades) but slower than for CMclean (clean blades).

Rated power is reached at a higher flow speed than with clean blades but at a lower flow speed than if there were no adjustments to the controller. This control method should therefore present a net energy yield benefit, relative to CM2, both because more time is spent at rated flow speed and because higher power is delivered below rated than would be generated if there were no controller adjustment.

Once rated power is achieved the rotor blades are again pitched to limit the generator torque to $\tau_{gen} = \tau_r$. Blade pitching continues until the maximum permitted pitch is reached at $\beta = \beta_{max}$, resulting in a cut-out flow speed slightly in excess of that seen in clean conditions, U_{co}.

3.2.4 Performance Loss

The impact of airfoil roughness on turbine Annual Energy Production depends not only on its impact on turbine performance (efficiency), but also its impact on the proportion of time that the turbine operates at wind speeds where performance loss is significant. AEP was assessed through numerical integration of Equation 4.8. The wind speed distribution used in this calculation $f(U)$ is evaluated using a Weibull distribution [93], which is defined in Equation 4.9 in Section 4.4.2.

Throughout this thesis, Weibull distributions are used to calculate energy yields of wind turbines and wind farms. A Weibull distribution has a zero probability of wind speeds of 0, which is not accurate as very low wind speed periods occur. However, it provides a good approximation of long-term wind speed distributions and wind speeds below the cut-in wind speed of a wind turbine will not impact modelled results in this thesis.

Wind speed data from four offshore wind farm sites in the UK with different flow speed distribution characteristics were investigated: the Kentish Flats (see Section 4.4.2), Thanet, Greater Gabbard and Gwynt y Mor. All data were retrieved from the Crown Estate's marine data exchange [94] and were extrapolated to a common height of 80m using the log law in a neutrally-stable condition [32]:

$$U(z) = \frac{U_f}{\kappa} \ln\left(\frac{z}{z_0}\right) \tag{3.4}$$

where U_f is the friction velocity, κ is von Kármán's constant (taken to be 0.4), z is the height at which the velocity is evaluated (here 80m) and z_0 is the roughness height, which is the height at which the flow speed in the log law profile approaches zero (assumed to be 0.0002m based off the Davenport terrain roughness classification [95]). Important data for each of these sites can be found in Table 3.3.

The wind speed data, available as 10 minute average wind speeds through the durations noted in Table 3.3, were used to determine best fit Weibull wind speed distributions. The determined Weibull shape parameters and scale factors were evaluated and are presented in Table 3.3. The proportions of the distributions

relative to U_{ci}, U_r and U_{co} are given in Table 3.3 and shown in Figure 3.4. Differences in these proportions will be important as different wind distributions will affect AEP by changing how much time the turbine spends in pre-rated, post-rated or non-operational regimes. However, the underlying aerodynamic mechanisms causing performance drop will not be considered to be affected by site location.

Table 3.3: Summary of the four wind sites considered for performance analysis. Flow speeds have been extrapolated from data collected at the "original altitude" through a log law profile to a common altitude of 80m.

Site name	**Years of data**	**Original altitude**	$\mathbf{k_s}$	$\mathbf{c_s}$ **[m/s]**	$\mathbf{U_{ci} < U < U_r}$	$\mathbf{U_r < U < U_{co}}$	$\mathbf{U < U_{ci}}$ **or** $\mathbf{U > U_{co}}$
Kentish Flats	2002 to 2005	80m	2.1073	8.8928	64.6%	27.6%	7.8%
Thanet	2005 to 2006	10m	2.3471	9.9750	57.8%	35.7%	6.5%
Greater Gabbard	2005 to 2014	86m	2.3028	10.5542	52.3%	40.7%	7.0%
Gwynt y Mor	2005 to 2007	85m	2.0306	10.3791	51.0%	39.2%	9.8%

3.3 Results

3.3.1 Rotor Performance

The rotor designs were simulated with their respective clean and rough airfoil data to evaluate the effects of roughness on overall performance. A summary of the power and thrust coefficients, C_P and C_T, for each rotor operating under rotor design conditions ($\lambda^* = 10$, $U_r = 10$m/s) for both clean and rough airfoil data is provided in Table 3.4. C_T is defined analogously to the local thrust coefficient, as the ratio of rotor thrust, T, to the force generated by the upstream dynamic pressure acting on the frontal area of the turbine:

$$C_T = \frac{T}{(1/2)\rho U^2 A}. \tag{3.5}$$

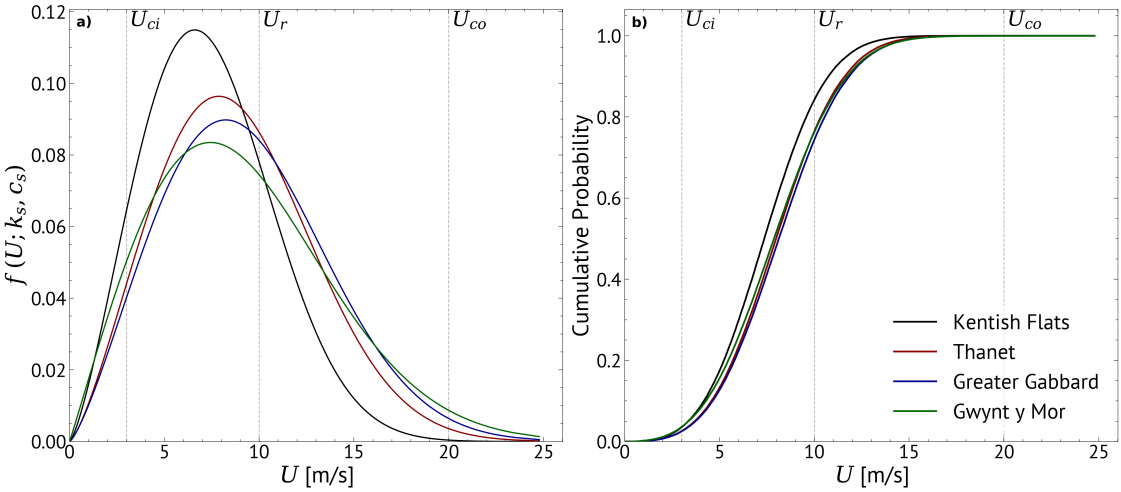


Figure 3.4: Wind speed probability distributions modelled with (a) the two parameter Weibull distribution and (b) the cumulative probability density for the Kentish Flats, Thanet, Greater Gabbard and Gwynt y Mor sites. Vertical lines have been added to indicate U_{ci}, U_r and U_{co}. Wind speed data is from 10 minute average data extrapolated through a log law profile to a common altitude of 80m.

In all cases, at the design point, the turbine power is adversely affected by blade roughness. This performance reduction is accompanied by a reduction in observed turbine thrust at this specific tip-speed-ratio ($\approx 10\%$ of C_T is shown in Table 3.4). Whilst this does not affect individual turbine analysis of this chapter such change of momentum extraction may be expected to alter the interactions between wind turbines and hence alter the wind farm power curve.

Table 3.4: Coefficient of performance and thrust coefficient for each rotor operating under design conditions ($\lambda^* = 10$, $U_r = 10$m/s) for clean and rough airfoil data.

Rotor	Airfoil Condition	C_P	C_T
NACA4415	Clean	0.4991	0.8540
NACA4415	Rough	0.4382	0.7784
S801	Clean	0.5107	0.8540
S801	Rough	0.4571	0.7647
S810	Clean	0.4932	0.8540
S810	Rough	0.3995	0.7567

3.3.2 Turbine Performance Comparisons

Figure 3.5 illustrates the performance of the turbines with wind speed for clean and rough rotors using each of the control methods. On initial inspection, the torque-referenced control strategies (CM2 and 3) are more effective than the wind speed-controlled strategy (CM1), as indicated by the higher power they achieve relative to CM1 across all flow speeds. All roughened rotor cases achieve lower power than the clean rotor examples in below rated conditions highlighting the effect that airfoil roughness has on rotor torque.

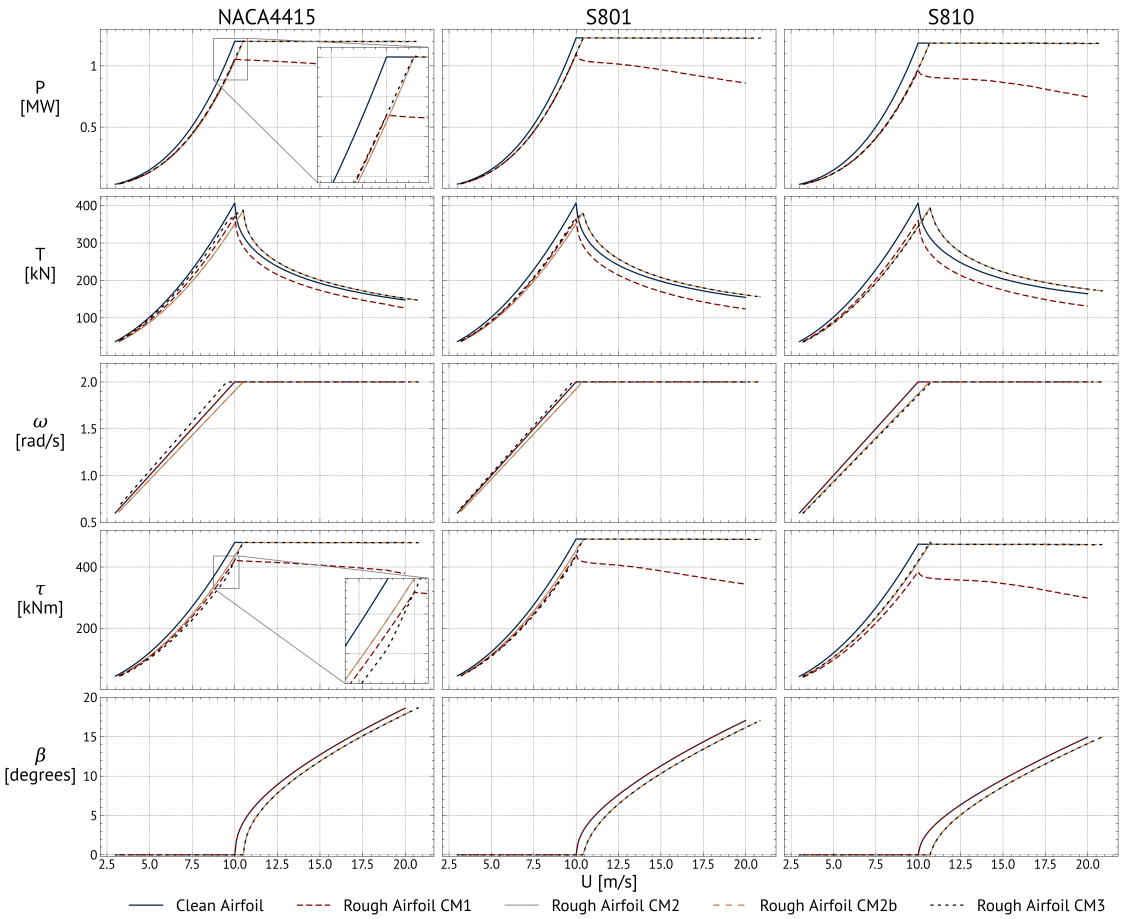


Figure 3.5: Performance comparison between rotors with roughened and clean blades for the three rotor designs (NACA4415 left, S801 centre, S810 right) for all control methods. The x-axis for all plots is wind speed. Two inset graphs are present: one for the NACA4415 power vs wind speed, which demonstrates the improved performance of CM3; and one for the NACA4415 torque vs wind speed, which illustrates the change in gradient for the CM3 torque profile in this area.

For all control strategies, the clean airfoil turbines perform identically because

the rotor is designed to operate with this airfoil, meaning that C_P and λ take optimal values for the clean rotor, and ω_r and τ_r occur at the same wind speed. The roughened airfoil rotors have cut-in speeds at higher wind speeds (3.2, 3.2 and 3.3 m/s for the NACA4415, S801 and S810 rotors respectively) than the clean rotor, as a high wind speed is required to attain the same cut-in torque as in clean operating conditions.

Analysing CM1 between cut-in wind speed and rated flow speed, the roughened blade produces less power due to the reduction in torque (at a given rotational speed) that arises due to the reduced lift and increased drag of the roughened airfoils. The NACA4415 rotor has a maximum power (P^*) difference of 0.29 MW with the clean rotor producing $P^*_{clean} = 2.40$ MW. Similarly, the S801 rotor maximum power drops from $P^*_{clean} = 2.46$ MW to $P^*_{rough} = 2.20$MW and the S810 rotor maximum power drops from $P^*_{clean} = 2.37$ MW to $P^*_{rough} = 1.92$ MW. For all CM1 cases, both smooth and rough rotors reach maximum rotational speed at the rated wind speed, as stipulated by the control method. Consequently, both clean and roughened rotors start to pitch at a wind speed of 10.0 m/s, which also corresponds to the point of maximum thrust (T^*) for all rotors. For all rotors $T^*_{clean} = 407$ kN, as they are designed using the same design condition as discussed in Section 3.2.2. The roughened NACA4415 rotor has a peak thrust $T^*_{rough} = 371$ kN, whilst for the S801 rotor $T^*_{rough} = 364$ kN, and for the S810 rotor $T^*_{rough} = 360$ kN.

Above rated speed, different trends between the clean and rough airfoil rotors emerge. The clean airfoil turbines adjust their blade root pitch angle to maintain rated power, giving rise to a 'typical' power curve. However, the roughened blade rotor has $\partial P_{CM1}/\partial U < 0$ above rated flow speed. This occurs because the pitching mechanism is designed to maintain generator torque for the clean airfoil, and is thus sub-optimal when the rotor aerodynamics have changed. The changed aerodynamics can be understood through Equations 3.6 and 3.7, where $C_t(\phi)$ is the tangential (torque-generating) component, and $C_n(\phi)$ is the normal (thrust-generating) component of force on the rotor respectively, where the flow incidence angle $\phi = \alpha + \gamma_t$, where $\gamma_t = \beta_g + \beta$ where β_g is the initial geometric twist of the blade,

and β is the additional blade pitch through control. The general trend of airfoil roughness is for lift to decrease and drag to increase, which results in $C_t(\phi)$ decreasing in all cases, and $C_n(\phi)$ remaining approximately constant. Therefore, at a given α, the torque is reduced and thrust is similar between clean and roughened cases.

$$C_t(\phi) = c_l(\alpha)\sin(\phi) - c_d(\alpha)\cos(\phi) \tag{3.6}$$

$$C_n(\phi) = c_l(\alpha)\cos(\phi) + c_d(\alpha)\sin(\phi) \tag{3.7}$$

Above rated flow speed the blade pitch is adjusted to reduce α and thus reduce lift in order to limit rotor torque. Under CM1, this adjustment is performed at the same rate, as a function of wind speed, as for the clean rotor which, all other flow conditions being equal, produces more torque than a roughened rotor. Consequently, rotor torque, and thus power, slowly decrease with flow speed for the roughened rotor under CM1. This is highlighted in Figure 3.6 where torque decreases once the maximum rotational speed is achieved.

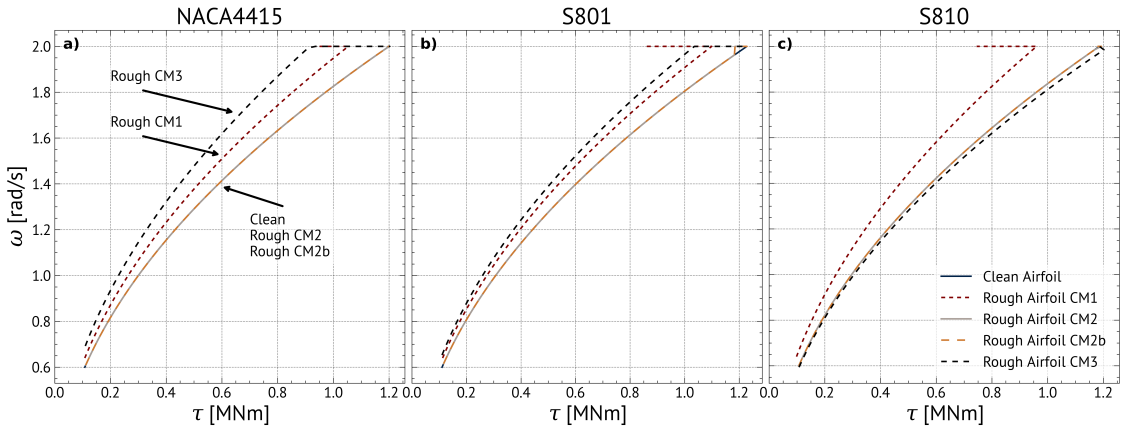


Figure 3.6: Torque as a function of rotational speed for the (a) NACA4415 rotor, (b) S801 rotor and (c) S810 rotor, for both clean and rough airfoil conditions and all control methods.

For the rough airfoil rotors, CM1 and CM2 perform similarly from cut-in wind speed to rated wind speed, as both control methods CM1 and CM2 utilise the same K_{clean} below rated. However, CM2 enables τ_{gen} to continue to increase beyond U_r until τ_r and P_r are reached. The rough airfoil turbine then pitches so as to maintain τ_r. For each of the clean rotors, rated power was reached at a wind speed of 10.0

m/s, beyond which the blade pitch angle was adjusted to limit rotor power. Blade pitch adjustment doesn't occur until slightly higher flow speeds are reached for the rough rotors, beginning once rated power is obtained at 10.5, 10.4 and 10.7 m/s for the NACA4415, S801 and S810 rotors respectively. At the point that the roughened rotor with CM2 control reaches the cut-out speed, it has not obtained the maximum pitch setting β_{max}. The operating range of the roughened rotors can be extended using CM2b so that cut-out occurs at a higher wind speed once β_{max} is reached. The rotors cut out at wind speeds of 20.8, 20.9 and 20.9 m/s for the roughened NACA4415, S801 and S810 rotors respectively, compared to a cut-out speed of 20 m/s for the clean rotors.

As may be inferred from Equations 3.2, 3.3 and 3.6, adjusting the turbine control to account for roughness may help increase rotor torque relative to control based on clean blade characteristics. This results in the different control curves in Figure 3.6, where increasing K reduces the curve gradient. CM3 adjusts this balance by determining a K value that maximises annual energy production based on the roughened airfoil characteristics. The power inset plot in Figure 3.5 shows the improved performance (higher power) for a given wind speed when using CM3 compared to CM2. Each rotor changes differently, as the different aerodynamic properties of the rotor airfoils result in the roughened rotor $C_P - \lambda$ power curve varying such that the required change to K to re-optimize for roughened performance can change by different magnitudes and in different directions. The purpose of changing K is to recover some of the power that is lost below rated torque due to erroneous control of the rotor as if it were still clean. Whether K_{rough} is greater or less than K_{clean} depends on the α required to achieve maximum rough c_l/c_d compared to that required to achieve the maximum in the clean condition.

It is useful to consider a representative angle of attack, $\tilde{\alpha}$, to understand the direction K will move for a given airfoil rotor. When operating using K_{clean} (i.e. CM2/CM2b) with roughened blades, a range of α occurs along the span of the rotor (see Figure 3.8 a) and $\tilde{\alpha}$ can be thought of as being representative of these values. If $\alpha^*_{rough} > \tilde{\alpha}$ then by increasing K, a lower rotational speed is prescribed,

via Equation 3.3, resulting in an increase in the flow angle ϕ and therefore α, until $\tilde{\alpha} \rightarrow \alpha^*_{rough}$ and a roughened rotor optimum is achieved. If initially $\alpha^*_{rough} < \tilde{\alpha}$ then K must be decreased so as to recover the optimum. However, the exact required changes in K will be a function of blade twist distribution and are therefore too complex to define precisely.

K is observed to increase for the S810 airfoil and decrease for the other two when optimised for roughness. In turn, these changes in K decrease the rotational speed and therefore tip-speed ratio for S810, and increase these for the other rotors. By changing K, operational $C_{P,rough}$ is maximised from cut-in wind speed to rated torque. The cut-in flow speed also occurs approximately 0.1–0.2m/s earlier with K optimised than with CM2/CM2b, as the required cut-in torque, τ_{ci}, is achieved at a lower wind speed with the optimised settings.

For CM3, as the flow speed is increased, the rotor spins more quickly than it would under CM2 and reaches maximum rotational speed, ω_r, ahead of obtaining maximum torque τ_r, so that between the point of maximum ω and τ, the torque must rise rapidly with flow speed, and hence $\partial\tau/\partial U$ increases as seen in the torque inset plot of Figure 3.5. This can be identified in the torque vs rotational speed plot for this case (Figure 3.6), which shows that once $\omega = \omega_r$, $\partial\omega/\partial\tau = 0$ as torque increases to its maximum value.

3.3.3 Spanwise Effects of Roughness

Figure 3.7 (a) shows the spanwise contribution to AEP along the rotor calculated using CM2/CM2b and evaluated for the turbines using the Kentish Flats wind distribution. The shape of these distributions differ from the optimal performance torque profiles which increase monotonically to maxima at $r/R = 0.865$ for all rotors (both clean and rough) before tip loss effects cause the blade loads to reduce (see Figure 3.7 (b)). Maximum torque occurs at the same radius for clean and rough rotors, whereas the radial location for maximum AEP contribution may shift inboard due to the influence of roughness. The change in peak torque per unit span produced between clean and rough blades is 10.9%–18.6%, whereas the change in

total AEP is 4.9%–8.6% (based on the Kentish Flats wind distribution). Hence, we note that a combination of both roughness and wind speed characteristics are important for calculating performance losses. Figure 3.7 (a) identifies that there is a fundamental performance difference between the S810 rotor and the NACA4415 and S801 rotors, especially over the inboard 50% of span. The differences in AEP distributions between the rotors arise as a result of the differences in torque generated by the different rotors and control methods.

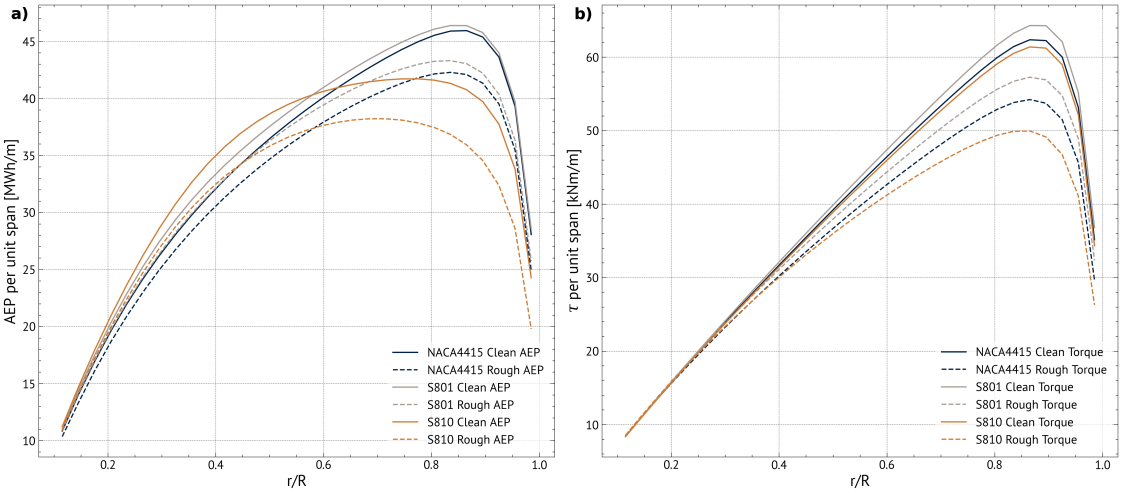


Figure 3.7: Spanwise contribution to (a) AEP (evaluated for the Kentish Flats wind distribution) and (b) blade torque for the NACA4415, S801 and S810 rotors in clean and rough conditions using CM2/CM2b.

Figure 3.8 shows the differences in performance characteristics due to roughness for all investigated blades in their optimal operating state (below rated torque) for CM2/CM2b. The clean rotors all have a spanwise-constant angle of attack, α, and axial induction factor, a, as these are specified through the rotor design process, see Section 3.2.2 ($\alpha = \alpha^*$, $C_x = 2$). A consequence of rotor roughness is a reduction in rotational speed at any given pre-rated flow speed (see Figure 3.5), such that the roughened rotors operate at a lower tip-speed-ratio than λ^*_{clean}. This, together with the reduction in airfoil forces due to roughness, results in a reduction in thrust along the length of each blade (see Figure 3.8 (d)), which results in an increase in axial flow speed through each turbine (reduction in induction factor a). Consequently, we observe a small increase in the angle of attack on the

roughened blades. The increase in α, however, does not typically have a dramatic effect on the average operational lift-to-drag ratio value of each airfoil, shown in Table 3.5, other than the rough NACA4415 rotor. The greatest changes due to roughness can be seen in the normal and tangential force component plots, calculated using Equations 3.6 and 3.7 respectively.

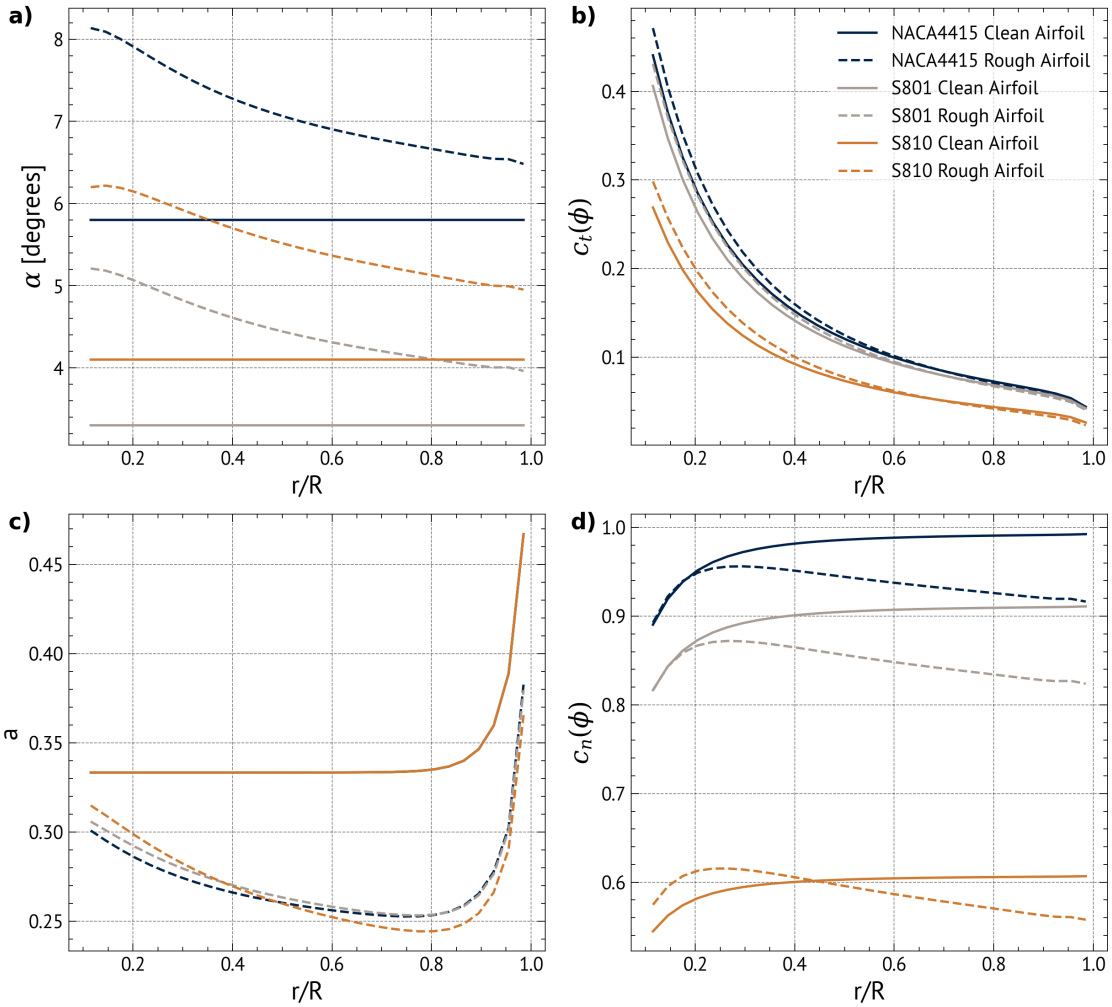


Figure 3.8: Spanwise analysis of NACA4415, S801, S810 rotors using CM2/CM2b for a wind speed of 10m/s showing the spanwise variation of (a) angle of attack, α, (b) tangential force coefficient, $C_t(\phi)$, (c) axial induction factor, a and (d) normal force coefficient, $C_n(\phi)$.

Furthermore, a simple control strategy for a roughened turbine to optimise α at the angle of attack to achieve $(c_l/c_d)^*_{rough}$ will not produce optimal performance,

Table 3.5: The maximum lift-to-drag ratios for NACA4415, S801 and S810 airfoils, together with the spanwise-averaged rotor lift-to-drag ratios for operation using CM2/CM2b.

Rotor	**Airfoil condition**	**Maximum c_l/c_d at α^***	**Spanwise average c_l/c_d**
NACA4415	Clean	102.45	102.45
NACA4415	Rough	52.64	44.86
S801	Clean	129.75	129.75
S801	Rough	60.50	60.33
S810	Clean	92.62	92.62
S810	Rough	36.84	36.59

because α^*_{rough} cannot be simultaneously achieved at all points along the rotor span. This is a consequence of the rotor being designed to operate with a clean airfoil at α^*_{clean}, and thus, when roughness alters the aerodynamics, even if these changes are known, α^* cannot be produced along the entire span of the blade as $\alpha^*_{clean} \neq \alpha^*_{rough}$.

The performance differences in $C_t(\phi)$ and $C_n(\phi)$ between clean and rough conditions are a result of the changes in lift and drag airfoil characteristics and in the operating point of the rotor (λ) under CM2/CM2b. It might be expected that roughness should lead to a reduction in tangential force at all spanwise locations. However, together with the change in tip-speed-ratio, the tangential force is observed to decrease across outboard sections but increase through inboard sections with a net resultant decrease in rotor torque due to roughness. Focusing on the operational airfoil characteristics in Table 3.5, the lift-to-drag ratio 'ranking' (from highest to lowest values) for the airfoils corresponds to the performance (AEP) order of the corresponding wind turbines for all wind sites.

3.3.4 Roughness Effects on Airfoils

Figure 3.9 shows the impact of blade roughness on the lift-to-drag ratio curve in the angle of attack range encountered in turbine operation $0° < \alpha < 10°$. Arrows are included to indicate the impact of blade degradation (c_l/c_d decreased by c. 56%) and the performance recovery impact of changing from CM2 to CM3 (average c_l/c_d increased by c. 18%).

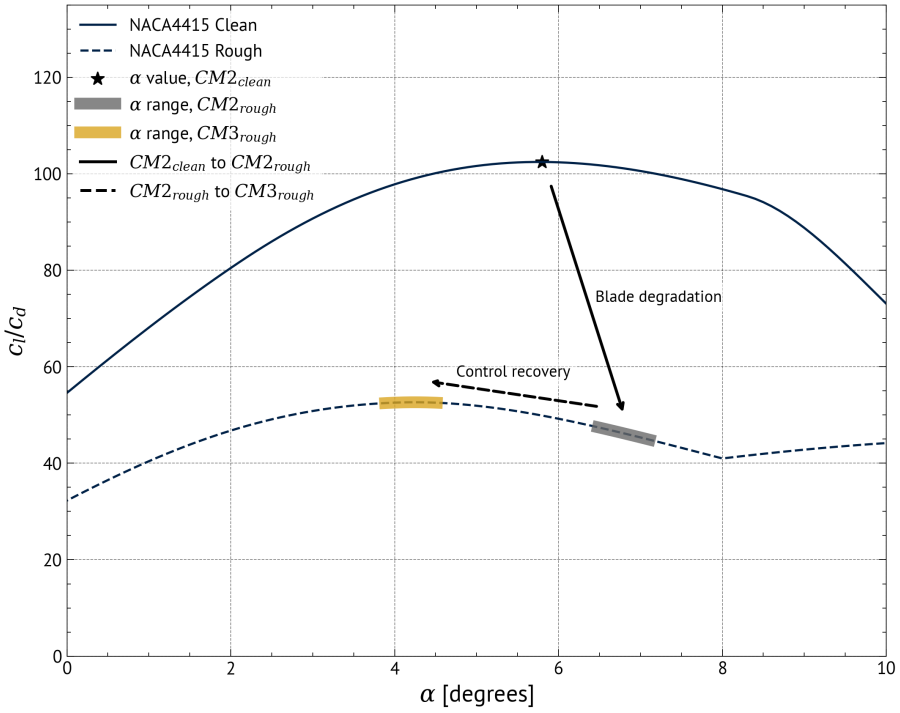


Figure 3.9: Airfoil characteristics c_l/c_d in the angle of attack region most relevant to turbine operation, $0° < \alpha < 10°$ for the NACA4415 airfoil. The operational α values (40–100% span) for the different control strategies have been highlighted, with clean cases denoted CM_{clean} and rough cases denoted CM_{rough}. Arrows indicate the impact of roughness and the control recovery mechanism of CM3.

Figure 3.10 shows the lift and drag profiles in the same angle of attack range as Figure 3.9, giving an indication of the effect of roughness on airfoil performance. Drag is approximately linear in this range and roughly doubles due to the effect of roughness for all airfoils. The increase in angle of attack that occurs when the airfoil is roughened (Figure 3.8 a) results in a similar magnitude of the lift force in the pre-stall region between the clean and rough case (approximate reduction of 5%). Consequently, the change in drag is the primary factor in determining the reduction in operational lift-to-drag ratio, and the magnitude of performance loss.

Figure 3.10 (c) highlights the range of operational α and c_l/c_d values that occur between 40-100% span in clean and rough conditions below rated power. When the clean rotor operates in its design conditions, it achieves a constant α that maximises c_l/c_d along the blade span. However, when blade roughness develops, the rotor no longer operates in design conditions and consequently the blade forces and the angle

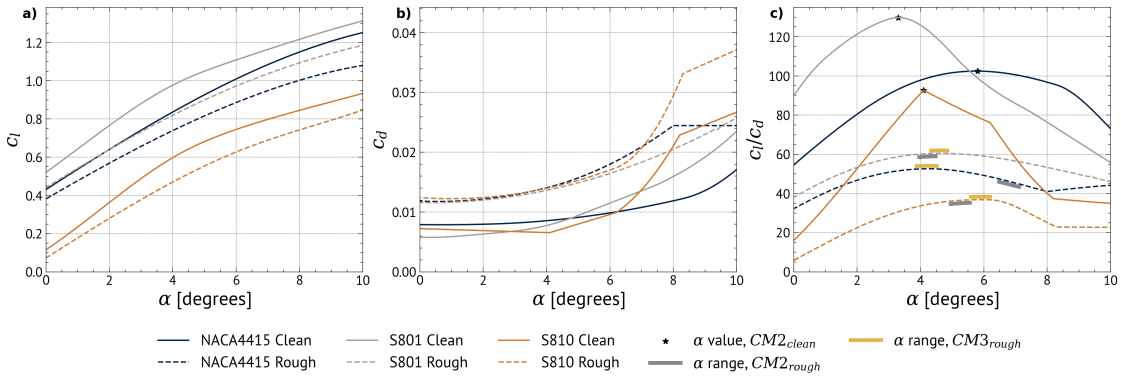


Figure 3.10: Airfoil characteristics c_l (a), c_d (b) and c_l/c_d (c) in the angle of attack region most relevant to turbine operation, $0° < \alpha < 10°$ for NACA4415, S801 and S810 airfoils. The operational α values (40–100% span) for the different control strategies have been highlighted, with all clean cases denoted CM_{clean}.

of attack will vary along the blade (see Figure 3.8 a). The operational angle of attack increases when airfoil roughness develops, as highlighted in Figure 3.10 (c). The figure also shows that operating the roughened rotor as if it were clean (CM2) tends to miss the c_l/c_d peak of the roughened airfoil. Adjusting turbine control to account for the degraded aerodynamic characteristics of the airfoil (CM3) allows some performance to be recovered by adjusting rotor speed to ensure that the operating α range of the rotor under CM3 encompasses the roughened airfoil c_l/c_d peak.

The increase in drag results in a decrease of the rough airfoil lift-to-drag ratio and an attendant reduction in the slope of the lift-to-drag curve $d\left(c_l/c_d\right)/d\alpha$. Therefore, even if the rotor does not achieve the optimal roughened angle of attack along the blades, $d\left(c_l/c_d\right)/d\alpha$ is small so that the rotor is found to operate close to, but not at, its maximum possible efficiency. The roughened performance can therefore only be marginally improved by further changing α. Consequently, while it is possible to alter K in CM3, this will result in only small improvements in power for the examined rotors and for the given levels of roughness. It is therefore of interest to ensure that airfoils used in wind turbine applications maintain a high lift coefficient and limited increase in drag due to roughness.

3.3.5 Energy Loss

The impact of roughness on energy production based on the wind distributions identified in Table 3.3 are shown in Figure 3.11. For all rotors, performance under CM1, where the turbine control is effectively referenced to a measured wind speed, is worst and results in very undesirable reductions in energy yield of up to around 20%. For the torque based CM2 control method in which the pre-rating control (K value) remains unaltered, the performance reduction ranges between 2-8% depending on the rotor and site, with the S801 rotor showing the least sensitivity in AEP to roughness, and the S810 rotor showing the greatest losses. For all rotors and sites some AEP may be recovered through control method CM2b, which extends the cut-out speed to higher flow speeds so that the maximum pitch angle β_{max} reached under rough airfoil conditions is the same as that reached at cut-out for the clean rotor. By altering the pre-rated control coefficient to K_{rough}, CM3 is able to achieve a partial yield recovery relative to CM2b, with a smaller performance drop relative to the clean rotors. The recovery in AEP from CM2b to CM3 across the sites is between 0.73–0.98%, 0.22–0.31% and 0.04–0.05% for the NACA4415, S801 and S810 airfoil rotors respectively (Figure 3.11).

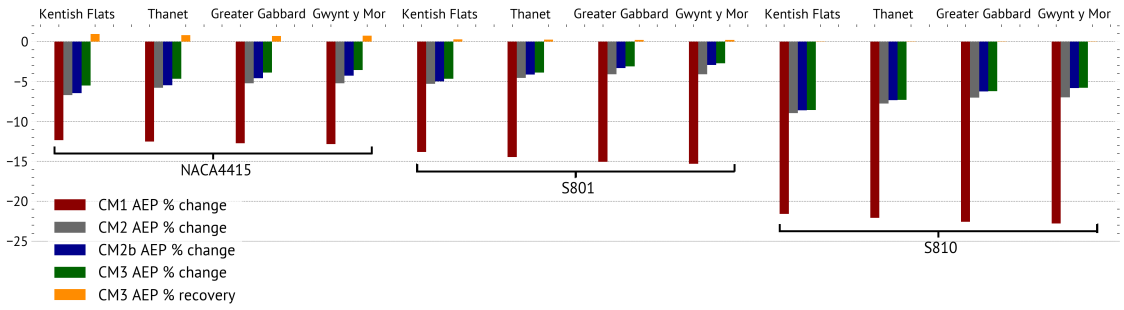


Figure 3.11: Energy capture comparison between clean and rough rotors operated with different control methods over a 1 year time period. AEP changes are measured relative to the clean operating rotor, and CM3 recovery is measured relative to CM2b.

The possibility for partial power recovery through this simple controller-based mechanism is demonstrated to be achievable but limited in scope due to the underlying performance reduction of the airfoils due to roughness; the approximate doubling of drag imposes an unrecoverable negative impact on rotor torque. However,

this is a relatively simple adjustment of rotor operation and therefore may still be an attractive method for mitigating performance loss, thus recovering some energy yield previously lost due to roughness. We note that the impact of roughness on turbine lifetime due to increased unsteady loading may also be significant but cannot presently be estimated with a steady-state model such as BEM theory and the available airfoil data.

The S810 rotor performs notably worse than the other two rotors presented herein. The S810 airfoil produces significantly less lift (clean or rough) than the other airfoils considered in this study, whereas its drag is similar. This is a function of the lower design max c_l for this airfoil, as highlighted in Table 3.1. Consequently, a large blade chord was required (see Figure 3.2) to meet the design requirements (thrust level), which also results in a greater contribution towards overall rotor AEP from the inboard portions of the blade (see Figure 3.7 a).

The varying ranges of performance losses shown in Figure 3.11 indicate the effect of wind distribution on the impact of roughness. The greater the likelihood of wind speeds above the rated speed (shown in Figure 3.4 and Table 3.3), the more likely the turbine is to operate at the rated power, whether clean or rough, thus reducing the negative effect of roughness on overall energy production. Figure 3.4 shows that the Kentish Flats site has the greatest probability of wind speeds below U_r, which gives rise to the largest performance decline from roughness. In effect, roughness mainly affects turbine performance when operating between P_{ci} and P_r, noting that roughness can cause a change in the wind speed at which P_r is achieved (roughness also impacts wind speeds above the cut-out point of the clean turbine). Similarly, CM3 is only able to improve power production in this range, so as the percentage change in performance due to roughness decreases, so does the percentage improvement from using CM3.

3.4 Conclusions

This chapter has explored the effect on wind turbine performance and annual energy production when operating with full-spanwise roughened rotor blades.

Experimental data from three NREL-tested airfoils were used to provide clean and rough aerodynamic lift and drag data. The clean airfoil data were used with an in-house BEM design tool to develop clean rotor blade designs that sought to operate at a design angle of attack along the blade and a prescribed local thrust coefficient.

The in-house BEM tool was used for clean operation and for roughened rotor operation considering three control methods: turbine operation referenced to wind speed (CM1); turbine operation referenced to generator torque (CM2); and a K-optimised torque control method for energy yield recovery (CM3). CM1 provided a baseline for comparison by adjusting the rotor speed and blade pitch angle with wind speed in the same manner as for each of the clean wind turbines. The rough airfoil rotors were unable to achieve rated torque, and energy losses ranged from 12.4 to 22.8% across the four different candidate wind sites considered.

In CM2, the rotor seeks to balance aerodynamic and generator torque below rated torque by adjusting rotational speed and above rated torque by pitch-to-feather control. Roughness leads to a slower rotational speed before rated torque is reached, with rated torque itself delayed to a higher flow speed than for the clean rotors. CM2 was found to be more effective than CM1, with energy losses ranging from 3.3% to 8.6% across the three rotors and four sites considered. A slight increase in energy yield was found using CM2b which extended the cut-out speed to higher flow speeds such that the maximum blade pitching of the roughened rotor blades matched that of the clean rotors at cut-out. CM3 improved on CM2 by optimising the controller gain parameter K for operating the rough rotors. The required change in K from clean to rough conditions is found to be dependent on rotor geometry and on the change in airfoil performance characteristics between clean and rough conditions. The rough rotors are found to produce the highest energy yield using CM3, which reduces energy losses by a further 0.1 to 1.0% relative to CM2b.

A significant performance drop was found for all simulated control strategies across the representative sites considered in this study. Using four separate wind farm sites for wind speed distributions, yearly energy losses between 2.9 and 8.6% were found for the three rotors using CM2/CM2b. CM3 was able to 'recover'

some portion, up to about 1%, of the power lost through re-optimising K. This re-optimisation adjusts the rotational speed of the rotor below rated torque to account for roughness effects. This is a simple control based energy yield recovery strategy. The extent to which energy yield recovery can be achieved is limited due to the high increase in blade drag through roughness, and the resulting reduction in the rate of change of the lift-to-drag ratio with angle of attack. Further research is required to identify airfoils and rotor designs that are less sensitive to the development of blade roughness.

There's always a story. It's all stories, really. The sun coming up every day is a story. Everything's got a story in it. Change the story, change the world.

— Terry Pratchett [96]

4

Parameterising the Evolution of Roughness on an Airfoil

Contents

4.1 Introduction 60
4.2 Roughened Airfoil Data 62
4.2.1 Comparing the Effect of Roughness 62
4.2.2 Available Airfoil Datasets 63
4.2.3 Data Comparison 67
4.3 Roughness Effects on Airfoil Performance 68
4.4 Roughness Impacts on Wind Turbine Performance . . 70
4.4.1 Turbine Model 71
4.4.2 Performance Change Metrics 72
4.4.3 Impact of Roughness on Performance Over Time 75
4.5 Discussion and Conclusions 76

The results presented in this chapter have been previously published in *Wind* titled "Parameterising the Impact of Roughness Evolution on Wind Turbine Performance" [77].

This chapter investigates the effects of surface roughness on airfoil performance and its consequences for wind turbine energy yield. This chapter addresses the main challenge related to understanding the impact of roughness on wind turbine performance: the lack of available roughened airfoil data that can be used in turbine simulations. 51 sets of experimental data across 16 airfoils designed for

wind turbine application are used to identify trends in roughened wind turbine airfoil performance. The trends are used to formulate a novel 'roughness evolution parameter' that can be applied to airfoils with no roughened data available to predict the impact of roughness on performance. Blade element momentum theory is used to model the performance of the Technical University of Denmark (DTU) 10 MW reference wind turbine (RWT) [97], with uniformly roughened blades estimated using the roughness evolution parameter. An annual energy production loss between 0.6–9.6% is found for the DTU 10 MW turbine when considering a plausible range of values for the roughness evolution parameter derived from the experimental data. A framework is developed in this chapter to evaluate how the roughness evolution parameter changes over time, informed by observed changes in wind farm performance from previous studies.

4.1 Introduction

Wind turbine performance declines over time because all systems are subject to deterioration as they age [98]. While there have been some advances in the diagnostics of wind turbine faults, relatively little research has been completed on the fault prognostics of wind turbines [99]. Research on individual components has indicated that the gearbox, generator, main bearing, blades and tower are critical components for turbine reliability [100].

Despite the increasing knowledge around the failure rates of, and downtime due to, individual components, there is less understanding of performance reduction caused by degradation of the aerodynamic surfaces, or 'roughness', of the turbine blades. Roughness that occurs on rotor blades may arise due to causes such as erosion, ice accretion, and insect contamination, as discussed in Chapter 2. The effect on airfoil and turbine performance from each of these causes can vary depending on operating conditions and environmental factors, making it difficult to quantify.

The impact of ice accretion on the aerodynamic performance of an airfoil has been tested experimentally in multiple studies, showing comparable changes in lift and drag [15, 101, 102]. Krøgenes and Brandrud [101] and Hann et al. [102] both

investigate the impacts on the NREL S826 airfoil under different icing conditions. Their results were synthesised by Vinnes and Hearst [103] and showed a 10% decrease in lift and 80% increase in drag. They also showed that a heavily ice affected airfoil reduces the angle of attack stall first occurs, which agrees with results published by DeGregorio et al. [104] who investigated the NACA0012 airfoil under similar ice accretion geometries. Furthermore, Jasinski et al. [43] conducted wind tunnel testing of the S809 airfoil and used the data to simulate a 20% performance decrease on a 450 kW wind turbine, Hudecz et al. [105] tested the NACA 64-418 airfoil at a constant angle of attack (α) and showed a decrease in lift coefficient of 22–34% for differing icing conditions, and Blasco et al. [15] tested the DU 93-W-210 airfoil showing a loss in lift of 16–25% and an increase of drag of 80–220% for different conditions.

While experimental studies, such as those reviewed above, provide valuable information about the effects of airfoil roughness, one of the key challenges is the limited number of airfoils for which there are sufficient data publicly available. Some studies have no or limited performance data [106, 107] or only provide data for a single angle of attack [105, 108]. This presents a challenge when trying to incorporate these data with wind turbine models, which typically require lift and drag data across a range of angles of attack. A further issue is that airfoil roughness is computationally challenging to simulate because the non-homogeneous and three-dimensional airfoil roughness geometry is difficult to define and computationally prohibitive to resolve [59, 76]. Additionally, wall-modelling turbulence closures in CFD models are not validated for roughness applications [76]. To overcome the difficulties in producing robust roughened airfoil data, the present chapter aggregates all available experimental airfoil roughness data in order to provide insights into the generalisable performance trends of roughened airfoils that can be used to estimate key characteristics of the impact of roughness on other airfoils.

Experimental airfoil data is sourced from research conducted at OSU under contract from NREL [84], Sandia National Laboratories (Sandia) [109], Hann et al. [102], Blasco et al. [15], Jasinski et al. [43], and Sareen et al. [58].

Using the existing experimental data, a novel 'roughness evolution parameter' is proposed in this chapter that can be applied to clean (non-roughened) airfoil data to synthesise roughened airfoil data. This parameterisation can be updated as more experimental roughened airfoil data become available in future studies. The impacts of airfoil roughness are then evaluated for the DTU 10 MW RWT.

This chapter is structured as follows: Section 4.2 describes the data sources used, Section 4.3 details the effect of roughness on airfoil performance and defines the roughness evolution parameter, and in Section 4.4 the application of the roughness evolution parameter to the DTU 10 MW RWT is demonstrated to analyse the impacts of airfoil roughness on wind turbine performance.

4.2 Roughened Airfoil Data

This section summarises the available airfoil data and their sources. All airfoil data were collected from publicly available sources, as detailed below. For all datasets, only the data from the highest chord-based Reynolds number ($Re = \rho V c/\mu$) experiment with a sufficiently wide range of angle of attack, α (at least $-5^\circ \leq \alpha \leq 10^\circ$) was used in this thesis, where ρ is air density, V is wind speed relative to the airfoil section, and μ is the viscosity of the air. The non-dimensional roughness height is defined relative to the chord length, k/c. While the magnitude of the experimental chord-based Reynolds numbers is in the range 10^5 to 10^6, compared to a 10^7 for a utility-scale turbine, the presence of roughness is expected to trip the boundary layer into turbulent flow near the leading edge of the airfoil. Consequently, the roughened airfoil experimental data are expected to provide a reasonable comparison to the conditions experienced by a utility-scale wind turbine.

4.2.1 Comparing the Effect of Roughness

Lift (l) and drag (d) forces per unit span, as functions of angle of attack, for all airfoils are converted to non-dimensional lift and drag coefficients, $c_l = l/((1/2)\rho V^2 c)$ and $c_d = d/((1/2)\rho V^2 c)$ respectively. The effects of roughness are evaluated using the percentage change in the lift and drag coefficients between clean and rough

conditions, shown in Equations 4.1 and 4.2 respectively, where Δ represents the percentage change of these metrics.

$$\Delta c_l(\alpha) = \frac{c_{l,clean}(\alpha) - c_{l,rough}(\alpha)}{c_{l,clean}(\alpha)} \cdot 100\% \tag{4.1}$$

$$\Delta c_d(\alpha) = \frac{c_{d,rough}(\alpha) - c_{d,clean}(\alpha)}{c_{d,rough}(\alpha)} \cdot 100\% \tag{4.2}$$

Taking the percentage change value allows for aerodynamic differences in lift and drag magnitudes to be shown, as well as allowing comparisons with other experimental datasets that might not be produced in directly comparable experimental conditions. The mean, median and variance of the relative lift and drag changes are then calculated.

4.2.2 Available Airfoil Datasets

Airfoil data for both aerodynamically clean and rough conditions are used in this chapter. Experimental data were available from six sources; data published from OSU research under contract from the NREL [84], Sandia National Laboratories (Sandia) [109], Hann et al. [102], Jasinski et al. [43], Sareen et al. [58], and Blasco et al. [15]. In total, 51 datasets and 16 airfoils are used in this study, and a summary of the data used in this paper is presented in Table 4.1.

OSU

The OSU study has publicly available data [84] for 13 airfoils (the L303 airfoil has not been included in the present chapter due to its largely cylindrical shape). A standard roughness pattern was developed from a molded insect pattern taken from a wind turbine in the field [57] and applied to all airfoils. To simulate leading edge accumulation of roughness, the density of roughness particles ranged from 5 particles per cm^2 in the centre of the pattern to 1.25 particles per cm^2 at the

Table 4.1: Summary of the datasets, and their sources, used in this study. "Airfoil(s)" lists the airfoils used in the dataset, Re is the Reynolds number of the data retrieved for the present work, k/c indicates the roughness height(s) applied to each airfoil in a given study and "Cases" indicates the number of roughness configurations applied to each airfoil in that study. For some studies, the roughness height k/c was unable to be defined due to the manner of leading edge geometry modification.

Source	**Airfoil(s)**	***Re*** **($\times 10^6$)**	**k/c**	**Cases**
OSU [84]	LS417, LS421, S801 NACA4415, S809, S810 S812, S813, S814 S815, S824, S825	1.25 to 1.5	1.9×10^{-3}	1
Sandia [109]	NACA63$_3$418, S814	3.2	1.23×10^{-4} 1.72×10^{-4} 2.46×10^{-4}	5
Hann et al. [102]	S826	0.2 to 0.6	N/A	6
Jasinski et al. [43]	S809	1.5 to 2	9×10^{-4} 1.9×10^{-3}	5
Sareen et al. [58]	DU 96-W-180	1.85	N/A	13
Blasco et al. [15]	DU 93-W-210	1.5	N/A	6

edge. Based on average particle size from the field specimen, the non-dimensional roughness height was $k/c = 1.9 \times 10^{-3}$ [57].

The OSU wind tunnel experiments were conducted across a range of chord-based $Re = 0.75 \times 10^6$ to 1.5×10^6 for each airfoil.

A minor alteration to this dataset is performed in order to effectively use the data in the present chapter. Repeated lift/drag data points at a common α value are averaged. The repeated values always occurred at the point of the lowest drag.

These data have also been used for other airfoil roughness studies by Munduate and Ferrer [59], Mendez et al. [110] and Kelly et al. [76].

Sandia

Sandia has conducted studies on leading edge erosion [109] that include wind tunnel testing for two airfoils: the NACA63$_3$418 and S814. These tests were completed between $Re = 1.6 \times 10^6$ to 4.0×10^6, at three different roughness heights (100, 140 and 200μm) and three different roughness densities (3, 9 and 15% airfoil coverage). These data are publicly available from the US Department of Energy [111]. For the

purpose of this paper, the $Re = 3.2 \times 10^6$ datasets were used for both airfoils as it was the highest Re data that had an adequate α range (at least $-5° < \alpha < 10°$).

The k/c values for the three roughness heights in the Sandia study are 1.23×10^{-4}, 1.72×10^{-4} and 2.46×10^{-4} respectively.

Hann et al.

Hann et al. recently explored the performance penalties associated with leading-edge icing on the S826 airfoil using both experimental and numerical techniques [102]. They completed experimental tests on six geometries at Re ranging between 2×10^5 to 6×10^5. Three ice geometries were produced from experimental tests at the icing wind tunnel of the Technical Research Centre of Finland (VTT) [112]. The ice geometries were then digitised by manually tracing their outlines [113]. The other three geometries were generated using a simulation tool, LEWICE [114]. The experimental ice shapes are reproduced from Hann et al. [102] in Figure 4.1 (a), and the LEWICE ice shapes in Figure 4.1 (b).

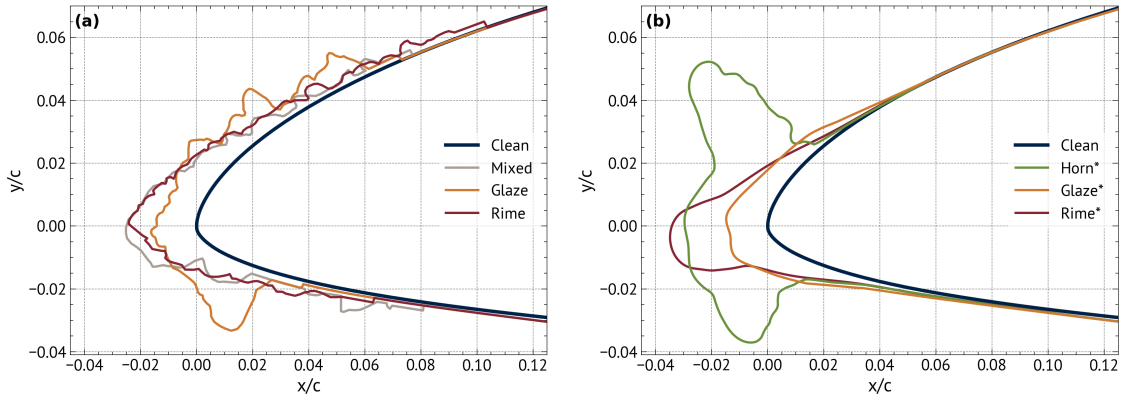


Figure 4.1: Experimental ice shapes (a) and LEWICE ice shapes (b), sourced from Hann et al. [102].

Data have been obtained for all six geometries at the highest available Re (6×10^5). Given the method used to produce the roughness seen in Figure 4.1, it is difficult to define a roughness height for these shapes, due to their irregularity and a change in airfoil shape.

Jasinski et al.

Jasinski et al. conducted experimental tests for four different supercooled fog conditions on the S809 airfoil [43]. The rime ice accretions were predicted using the LEWICE code, and the Reynolds number for the wind tunnel testing was in the range $Re = 1.0 \times 10^6$ to 2.0×10^6.

Data tables for lift and drag testing for the experimental tests shown in the Jasinski et al. paper[1] were made available by the corresponding author and have been used in this chapter. In these figures, experimental results from roughness heights of $k/c = 9 \times 10^{-4}$ and $k/c = 1.9 \times 10^{-3}$ (matching the OSU roughness height) are presented. Consistent with the other datasets, only the highest Re tests were used in this chapter[2].

Blasco et al.

As part of a study to quantify the power loss of a representative 1.5 MW wind turbine under icing conditions, Blasco et al. conducted wind tunnel testing to evaluate airfoil performance under various conditions [15]. Experimental tests were completed on the DU 93-W-210 airfoil, and six icing configurations were tested at $Re = 1.5 \times 10^6$. Icing conditions for these tests were suggested by collaborators at the National Center for Atmospheric Research, to replicate conditions experienced in the northern USA.

Roughness heights and shapes were recorded for each of the experiments, but due to the non-homogeneous and three-dimensional shapes, a k/c value was not assigned for each case. The experimental data tables are available in Blasco [115].

Sareen et al.

Sareen et al. conducted wind tunnel testing to investigate the effect of leading edge erosion on the aerodynamic performance of the DU 96-W-180 airfoil [58]. Experimental tests were performed between $Re = 1.0 \times 10^6$ to 1.85×10^6 under varying erosion conditions. Bug damage was also simulated on the airfoil to assess the impact of insect accretion on airfoil performance.

[1] Figures 4 and 5 from the Jasinski et al. paper have been used in this chapter.
[2] 2×10^6 for Figure 4 and 1.5×10^6 for Figure 5 in Jasinski et al.

Roughness height was not provided; however, it can be inferred from the depth of the erosion on the blade which is given in Table 1 of [58]. These roughness heights vary from other experimental tests because the roughness applied here is subtractive rather than additive to the surface. Only data from the $Re = 1.85 \times 10^6$ experimental tests are used in this thesis.

4.2.3 Data Comparison

The purpose of this section is to summarise the direct comparisons that can be made across the datasets used in this chapter. The S814 airfoil is common across both the OSU and Sandia datasets and can therefore be used for comparison. Figure 4.2 plots the lift and drag comparison between the two wind tunnel datasets (OSU $Re = 1.5 \times 10^6$ and Sandia $Re = 1.6 \times 10^6$), showing good agreement in lift and a systematic difference in drag in the range $-10° \leq \alpha \leq 10°$. This chapter adopts relative change metrics to minimise the potential impacts of systematic differences between different experimental datasets.

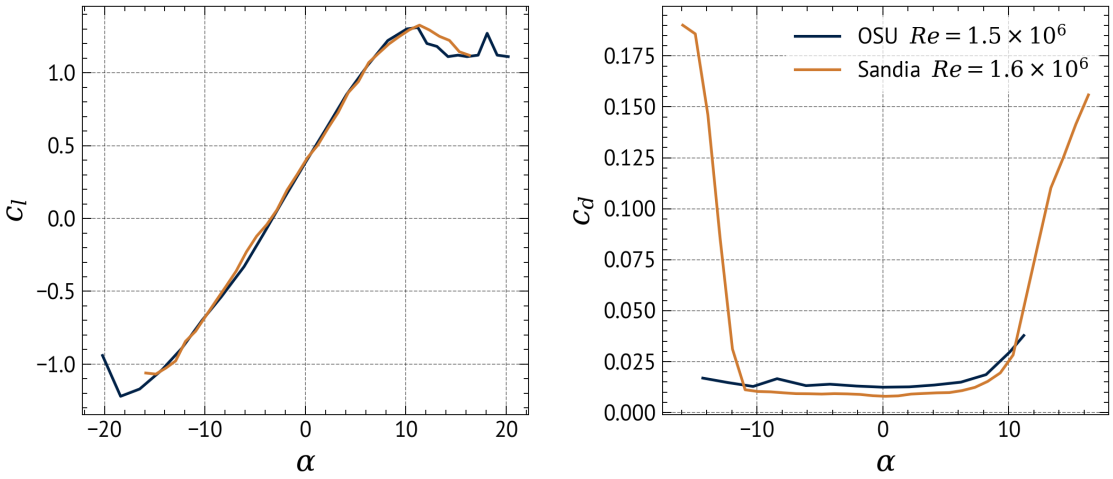


Figure 4.2: Experimental wind tunnel data comparison between OSU study at $Re = 1.5 \times 10^6$ and Sandia study at $Re = 1.6 \times 10^6$.

The S809 airfoil is common amongst this study and the OSU studies, and experimental results are validated in Figure 2 of the Jasinski et al. paper [43].

4.3 Roughness Effects on Airfoil Performance

The percentage changes in c_l and c_d for all datasets described in Section 4.2 are shown in Figure 4.3. The scatter points indicate the changes in lift and drag at the angle of attack α corresponding to the maximum lift-to-drag ratio ($(c_l/c_d)^*$) for the clean airfoil. Different symbols are used to represent the different studies included herein, and $(c_l/c_d)^*$ for all the airfoils in this chapter occurred at angles of attack between $3.7° \leq \alpha \leq 8.0°$. It is observed that there is a greater relative change in drag than in lift as the value of roughness increases.

One effect of roughness can be to change the angle of attack at which the maximum lift-to-drag ratio occurs from the clean case, and there can additionally be spanwise variation in angle of attack along the blades [76]. Consequently, Figure 4.3 also shows the changes in lift and drag in a range of $\alpha = \pm 2°$ around the angle of attack which maximises the clean lift-to-drag ratio. These ranges are shown as grey lines and intersect with the corresponding scatter point. This reinforces the observation that the relative changes in drag are greater than those in lift.

The general relationship between the relative change in lift and relative change in drag at α^* for clean airfoils can be described with a linear or a power-law model, as shown in Figure 4.3. These models have been evaluated for the data at the optimal angle of attack of the clean airfoil (shown as scatter points in the figure), with the constraint that the fits must pass through the origin. The equation fits were chosen based on minimising the mean squared error (MSE) and the mean absolute error (MAE). The two equations and the quality of their fit to the data are described in Table 4.2.

Table 4.2: Summary of the equation fits seen in Figure 4.3. Note, for the equations Δ represents percentage change in the corresponding metric.

Curve Fit	Equation	Eq. Number	MSE	MAE
Power law	$\Delta c_d = 13.12 \cdot \Delta c_l^{0.493}$	(4.3)	118.42	8.62
Linear	$\Delta c_d = 3.19 \cdot \Delta c_l$	(4.4)	166.03	10.91

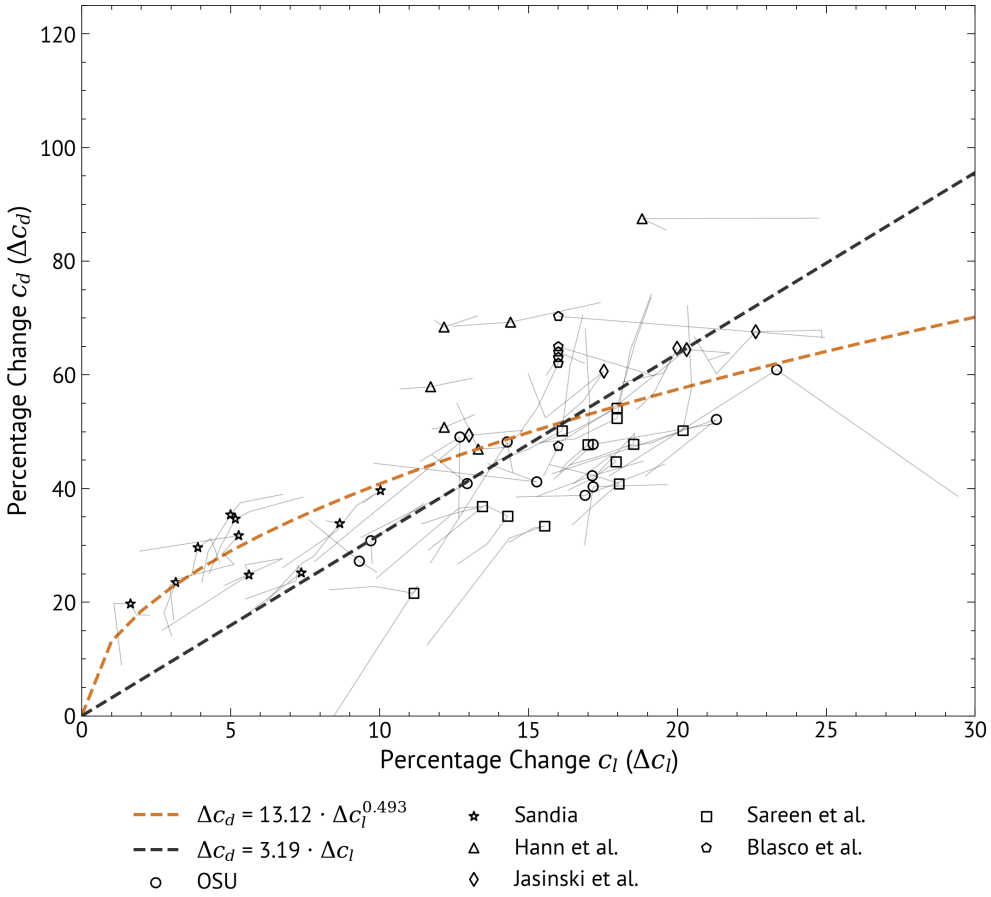


Figure 4.3: Percentage change in c_l is plotted against percentage change c_d for all datasets at the angle of attack for the maximum clean lift-to-drag ratio ($(c_l/c_d)^*$) point with linear (black) and power law (orange) best fit lines, plus the surrounding angles ($\pm 2°$) indicated in grey.

Both Equations 4.3 and 4.4 demonstrate that roughness has a larger effect on the relative increase in drag than the relative decrease in lift. Roughness has a larger impact on the skin friction component of drag than the pressure component which also affects the airfoil lift, as demonstrated by the rapid increase in drag at low levels of roughness, which is better captured by the power-law fit. While there are differences for a given roughness height, the curve-fits show that as an airfoil becomes increasingly roughened, it will tend to experience a reduction in lift and increase in drag at α^* for clean airfoils. Furthermore, the changes in lift and drag tend to be greater for airfoils with larger roughness heights and greater roughness extent. The rate at which airfoil performance declines due to roughness is not expected to be constant in time, but rather to decrease over time, that is

$\partial^2(\Delta c_l)/\partial t^2 < 0$ and $\partial^2(\Delta c_d)/\partial t^2 < 0$.

While it is intuitive that changes in lift and drag due to roughness will vary as a function of time, these changes are difficult to quantify and likely depend on environmental factors. However, by using the existing roughness datasets, a general 'roughness evolution parameter' can be created which can be used to convert clean lift and drag data to roughened data. This enables the effect that roughness could have on an airfoil or wind turbine to be evaluated without any roughened data being available. An equation for modifying airfoil data using a roughness evolution parameter, γ, can then be defined based on the power law model (Equation 4.5). γ is a parametric quantity that allows change in airfoil performance to be linked to the change in roughness over time. γ is a percentage value defined to be equal to the Δc_l (e.g. 20% decrease in lift is represented as 20) and changes over time as a turbine becomes more roughened and operating conditions change. The available data can be used as a guide to help choose a value of γ based on an assigned or assumed roughness height; this is discussed further in Section 4.4.3.

$$\begin{aligned} c_{d,rough} &= \left(1 + \frac{13.12 \cdot \gamma^{0.493}}{100}\right) \cdot c_{d,clean} \\ c_{l,rough} &= \left(1 - \frac{\gamma}{100}\right) \cdot c_{l,clean} \end{aligned} \tag{4.5}$$

Figure 4.4 shows the percentage change in lift and drag for different roughness heights aggregated across the datasets where k/c was defined. This shows that as roughness height progressively increases, lift decreases (i.e. Δc_l increases) and the data move up the best fit line as the drag increases. The development of airfoil roughness over time corresponds to moving towards the right of the figure.

4.4 Roughness Impacts on Wind Turbine Performance

This section details the methods and outcomes of combining and applying the roughened airfoil data results to the DTU 10 MW RWT to investigate the impacts of roughness on turbine performance. The (clean) airfoil data for the RWT are

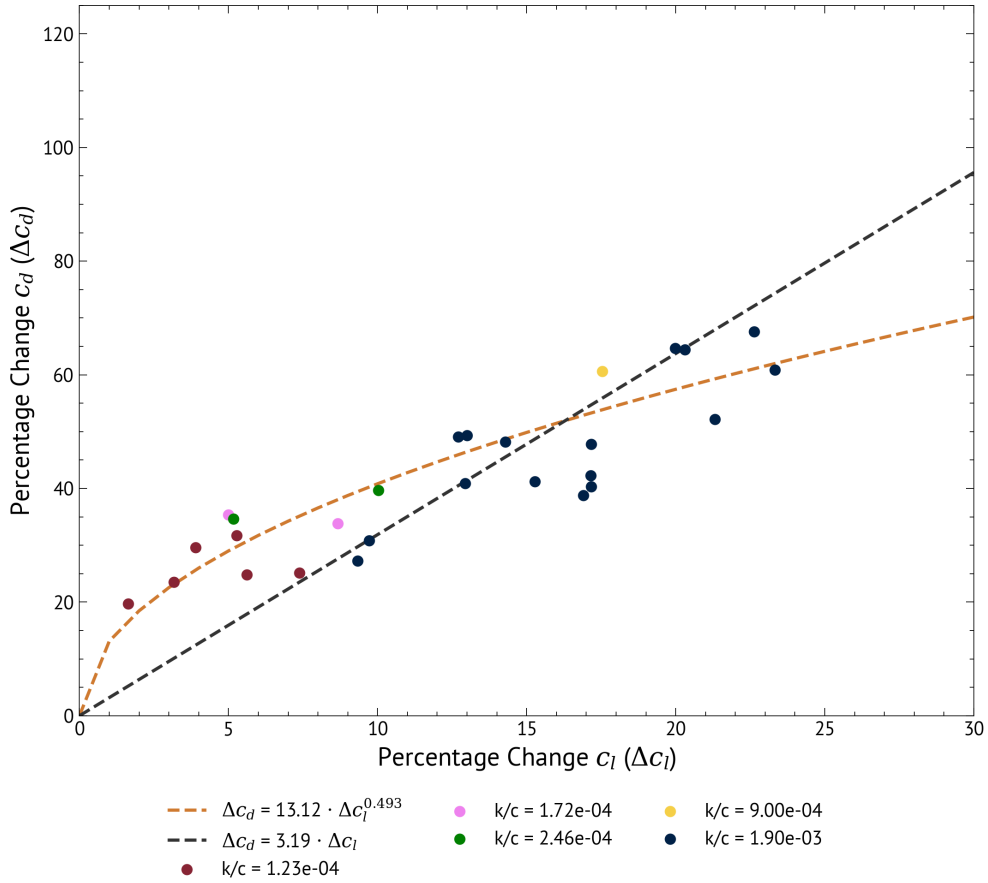


Figure 4.4: Percentage change in lift is plotted against percentage change in drag (at α^* for clean airfoils) for the OSU, Sandia and Jasinski et al. datasets with the linear (black) and power law (orange) best fit lines superimposed. The legend indicates the roughness height.

modified according to the model developed in Section 4.3 to create roughened airfoil data, which are then applied to the full rotor using BEM theory.

4.4.1 Turbine Model

The BEM method used in this chapter is detailed in Section 2.2.2 and has previously been validated with reference to BEM-based studies, see Chapter 2, Kelly et al. [77], Kelly et al. [76] and Vogel et al. [78] for further details.

The DTU 10 MW RWT was designed to provide a realistic, publicly available reference rotor for academic applications [97]. The DTU 10 MW RWT has a hub height of 119 m, rotor diameter of 178.3 m, a cut-in wind speed of $U_{ci} = 4$ m/s, a rated wind speed of 11.4 m/s, and a cut-out wind speed of $U_{co} = 25$ m/s. In

this chapter the DTU 10 MW RWT is modelled using the aforementioned in-house BEM method, with the publicly available airfoil and rotor data from DTU.

4.4.2 Performance Change Metrics

Performance loss for the DTU 10 MW RWT is calculated by comparing turbine performance (power coefficient C_P, thrust coefficient C_T and AEP) of clean and roughened airfoils. Note that the roughened airfoils are applied to the entire span of the rotor, implying a uniform distribution of roughness along the blade. While this is likely a simplified representation of blade roughness, it will indicate the general effects of roughness on turbine performance.

C_P was calculated by assessing turbine performance relative to the undisturbed kinetic energy flux, using Equation 4.6:

$$C_P = \frac{P}{\frac{1}{2}\rho U^3 \pi R^2} \tag{4.6}$$

where P is turbine power and R is blade radius. Similarly, C_T was calculated by comparing thrust against the undisturbed momentum flux, using Equation 4.7:

$$C_T = \frac{T}{\frac{1}{2}\rho U^2 \pi R^2} \tag{4.7}$$

where T is turbine thrust. AEP was assessed through numerical integration of Equation 4.8, where t is the time period (one year), $P(U)$ represents the power as a function of wind speed, and $f(U)$ represents the likelihood of a wind speed U occurring at a particular site, integrated across the range of operating wind speeds.

$$E = t \int_{U_{ci}}^{U_{co}} P(U) f(U) dU \tag{4.8}$$

The wind speed distribution $f(U)$ is evaluated using a two-parameter Weibull distribution [93]:

$$f(U; k_s, c_s) = \frac{k_s}{c_s} \left(\left(\frac{U}{c_s} \right)^{k_s - 1} \right) \exp\left(- \left(\frac{U}{c_s} \right)^{k_s} \right) \tag{4.9}$$

where k_s is the unitless shape parameter and c_s is the scale parameter in m/s.

Wind speed data from the Kentish Flats wind farm site in the UK were used in this chapter, and the data were retrieved from the Crown Estate's marine data exchange [94]. Kentish Flats is an offshore site and the wind data are from the 2002–2005 period, before the wind farm was built. The data were recorded at an altitude of 80m, in 10 minute average wind speeds, and have been extrapolated to hub height assuming a neutrally-stable log law profile. These data were best fit to a Weibull distribution with a shape parameter, $k_s = 2.83$, and a scale factor, $c_s = 10.52$. 68% of the distribution lies between $U_{ci} \leq U < U_{rated}$, 19% lies between $U_{rated} \leq U < U_{co}$, with the remaining 13% of the time lying between $U < U_{ci}$ or $U \geq U_{co}$ where the turbine is shut down. A different wind speed distribution, e.g. for a different location, will affect AEP by changing how often the turbine is in pre-rated, post-rated or non-operational configurations, however, the underlying aerodynamic mechanisms causing the performance drop will not be affected.

Effect of Roughness on DTU 10 MW RWT

Figure 4.5 displays the change in lift against change in drag values from all datasets as well as the (a) modelled C_P, (b) percentage change in AEP and (c) C_T values for the DTU rotor for the corresponding change in lift and drag. Also included are the lines of best fit from Figure 4.3, where moving towards the right along the curves indicate an increasing level of roughness.

Taking values from both ends of the experimental range, a plausible range of γ can be assumed to be $1 < \gamma < 25$ (i.e. a Δc_l of 1.00–25.00% and a Δc_d of 13.12–64.14% using the power law fit). Using this range, we can investigate different roughness cases for the DTU 10 MW RWT to evaluate the effect on performance. Two separate cases, $\gamma = 1.00$ and $\gamma = 25.00$, are applied to the clean airfoil lift and drag data in the range $1° \leq \alpha < \alpha_{sep}$, where α_{sep} is the angle where separation starts to occur in clean data. This results in a $C_P = 0.472$ ($\Delta C_P = -0.005$), $C_T = 0.799$ ($\Delta C_T = -0.005$), and an AEP reduction of 0.6% for the $\gamma = 1$ case, and a $C_P = 0.407$ ($\Delta C_P = -0.070$), $C_T = 0.664$ ($\Delta C_T = -0.140$), and a percentage AEP reduction of 9.6%, for the $\gamma = 25$ case respectively. This gives a plausible

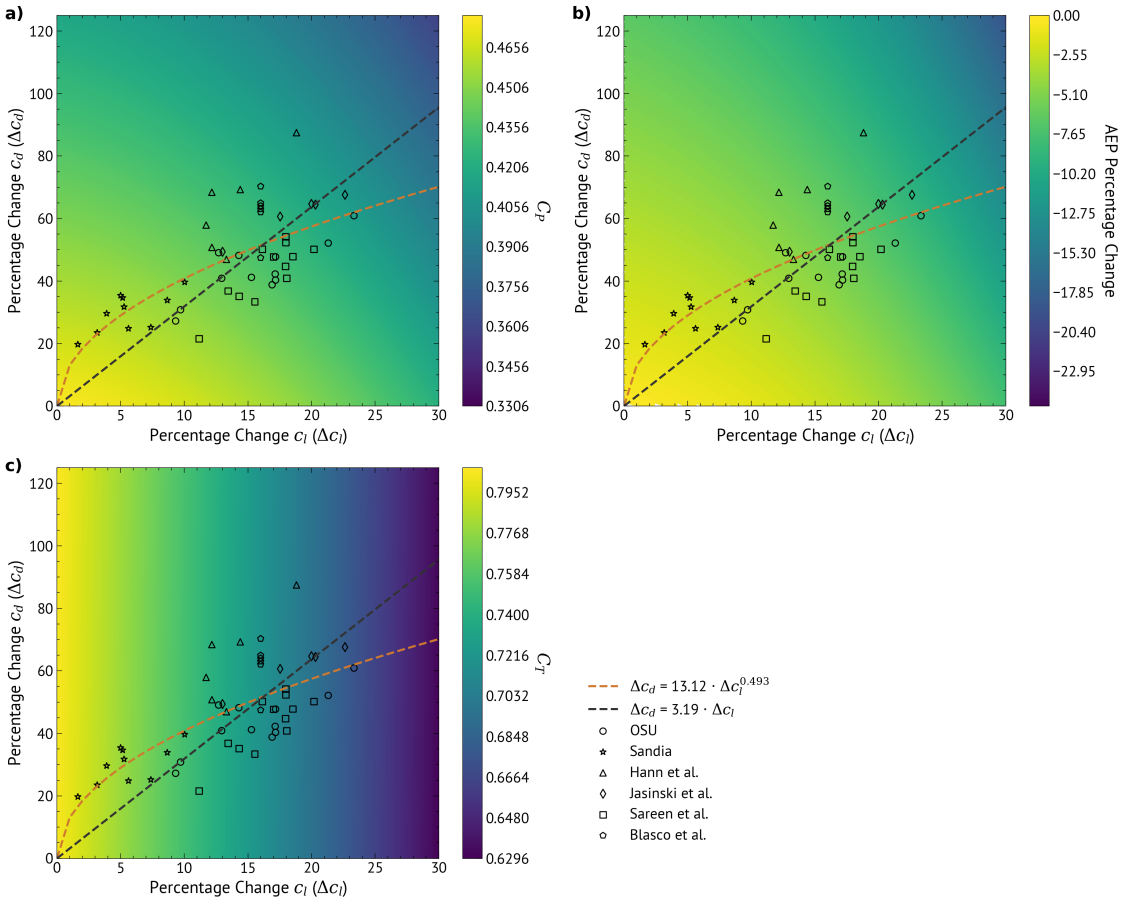


Figure 4.5: Percentage change in lift is plotted against percentage change in drag for all datasets with linear (black) and power law (orange) best fit lines are shown in all plots, and contour plots of the modelled C_P values (a), percentage change AEP (b) and C_T values (c) for the DTU 10 MW RWT.

range of performance decrease when the DTU 10 MW rotor blades are slightly ($\gamma = 1.00$) and severely ($\gamma = 25.00$) roughened.

The reduction in C_P as roughness increases means that the energy capture efficiency at lower wind speeds is reduced, and that the turbine reaches rated power at progressively higher, and less frequent, wind speeds. Consequently, the AEP worsens non-linearly as the level of roughness on the blades increases. The interaction between blade roughness and turbine control is discussed in more detail in Kelly et al. [76].

4.4.3 Impact of Roughness on Performance Over Time

The above analysis has yielded a roughness evolution parameter that can be applied to clean airfoil data to synthesize the effects of roughness where prior experimental data do not exist. To both contextualise the roughness evolution parameter values and relate them to energy yield, the AEP performance decline shown in Figure 4.5 (b) can be related to observed wind farm performance decline.

Multiple studies have explored how wind farm energy yields decline as they age. Staffell and Green [18] show that UK wind farms decline on average at 1.6% per year, Olauson et al. [19] show a 0.63% annual performance decline for Swedish wind farms, and Germer and Kleidon [20] show a 0.6% annual performance decline for German wind farms. These decline rates over time are shown in Figure 4.6 (a). These performance decline curves can be compared with the impact that roughness has on turbine AEP to estimate the AEP decline over time. It is assumed that the impact of wind farm wake effects does not change over time. Whilst wakes are not considered here the reduction of C_T for roughened airfoils would be expected to reduce the velocity deficit downwind of the turbine and hence reduce wake extent. This analysis then assumes for simplicity that the decline in performance is completely attributable to blade roughness, and that the rate at which outages (e.g. due to maintenance or curtailment) occur does not change over time. The implications of these assumptions are discussed further in Section 4.5.

The roughness evolution parameter γ can be plotted against AEP percentage change, using the intersection between the contours and best fit relationships for the DTU 10 MW RWT from Figure 4.5 (b) to produce Figure 4.6 (b). Under the assumption that changes in AEP over time are due to the development of roughness on turbine blades, a relationship between the age of a wind farm and γ is found using the common y-axis of Figure 4.6 (a) and (b). This demonstrates that the increase in the roughness evolution parameter gradually slows over time as the roughness becomes more developed on the blade. The rate of change is likely to be dependent on location, with offshore turbines more affected by the accretion of ice on the blade, whereas onshore turbines will tend to be more affected by insects and erosion

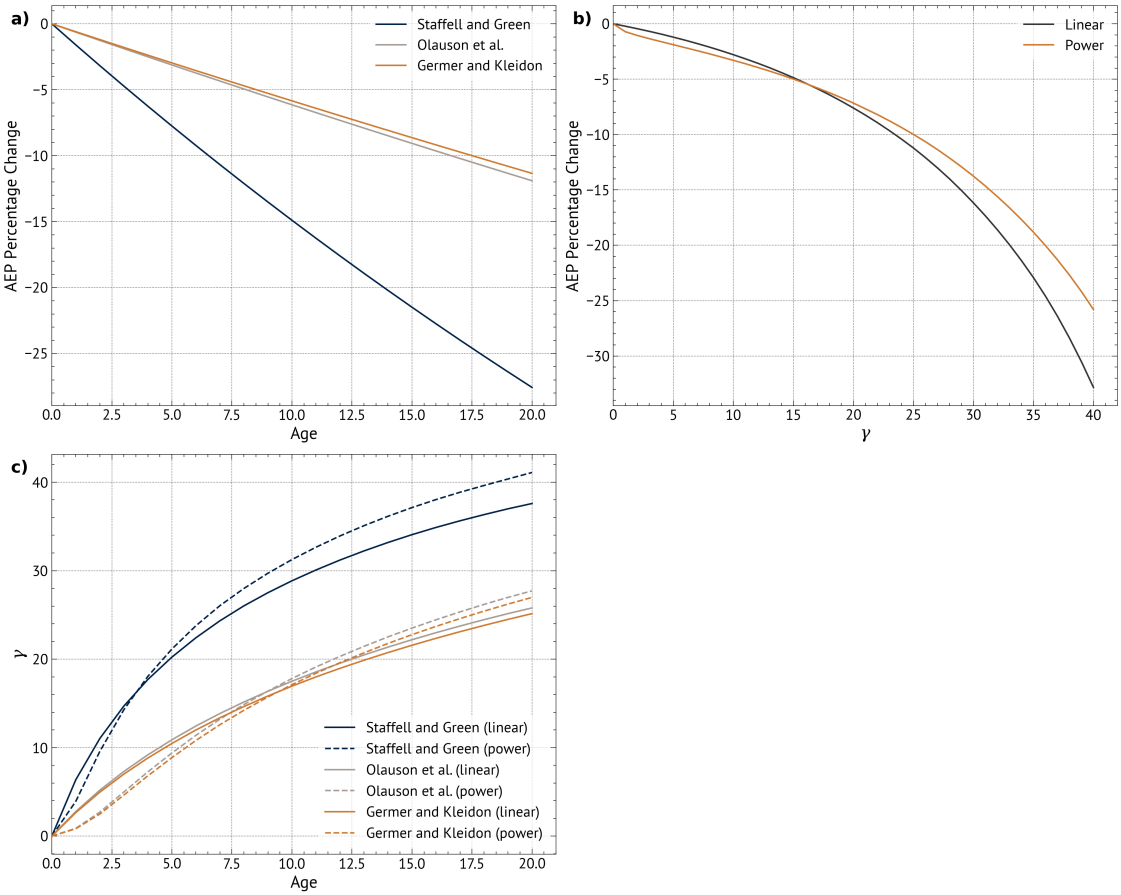


Figure 4.6: Wind farm age against AEP curves from Staffell and Green [18], Olauson et al. [19] and Germer and Kleidon [20] (a). γ against AEP as calculated using BEM analysis of the DTU 10 MW RWT (b). Wind turbine age against γ (c).

due to atmospheric dust. Combining 4.6 (b) and (c) can be used to relate wind turbine AEP performance loss due to roughness to turbine age. This relationship is shown in Figure 4.7, which bears a strong similarity to Figure 4.6 (a) - wind farm AEP variation over time - due to the assumption that the time variation in turbine performance loss is due entirely to airfoil roughness. This will vary with modified assumptions about how other factors affecting turbine performance that are beyond the scope of the present research, such as component failure, vary with age.

4.5 Discussion and Conclusions

A method has been created through the use of existing experimental data to create a 'roughness evolution parameter'. This parameter can be applied to clean airfoil

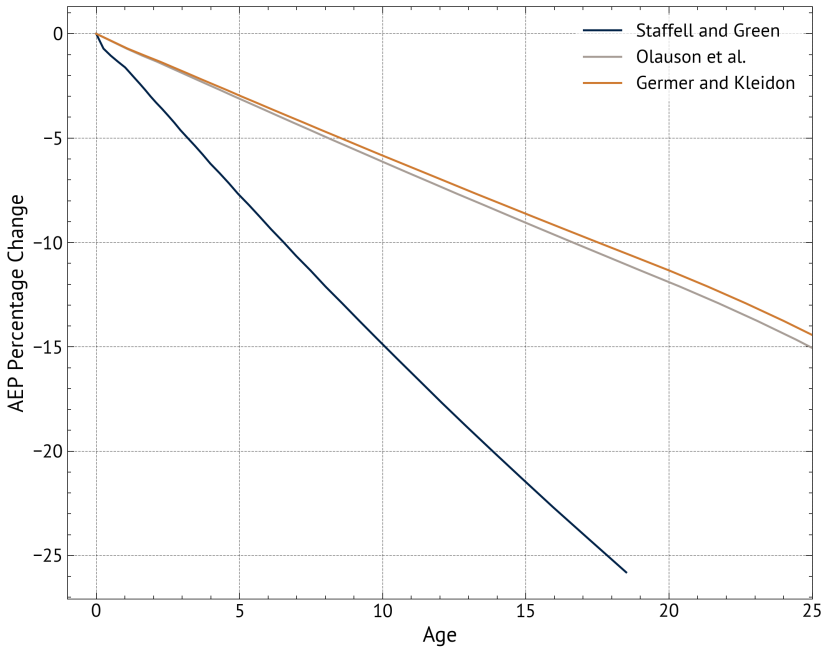


Figure 4.7: Wind turbine age against AEP performance change due to roughness based on wind farm performance decline rates from Staffell and Green [18], Olauson et al. [19] and Germer and Kleidon [20]

data to create the corresponding roughened data, which can then be applied to a rotor to examine the effect of blade roughness on turbine performance. This is useful because there is currently a lack of experimental data on roughened airfoils, which limits turbine studies in roughened conditions using new airfoil geometries. A demonstration case was investigated using the DTU 10 MW RWT that resulted in an AEP reduction range of 0.6–9.6%.

The creation of a roughness evolution parameter is a novel finding but would benefit from further data to verify the relationship. More experimental tests on airfoils with roughness applied should occur to evaluate the change in c_l and c_d. As this study has demonstrated, this should be completed for both a range of airfoil shapes and at different roughness heights. This is important to provide additional information on the relationship between change in lift and change in drag and allow further refinement of the roughness evolution parameter.

A framework for converting the roughness evolution parameter to AEP reduction over time has been proposed. The framework relates predictions of AEP (accounting for roughness) to observed changes in AEP from wind farms. Building on the wind

farm performance decline work of Staffell and Green [18], Olauson et al. [19] and Germer and Kleidon [20], the turbine AEP percentage change due to roughness was estimated to be in the range of 2.5–7.5% at age 5 and 8.5–22.5% at age 15 depending on which previous study is used.

This study has attributed the time-varying component of performance decline in wind farms to an increase in turbine blade roughness over time, which implies that the impact of factors such as wind farm wake interactions and component failure is constant in time. This simplifying assumption neglects the potential change in wake interactions as turbine performance changes and that probability of component failure is likely to increase non-linearly with turbine age.

Despite these simplifications, this study aims to provide a framework for a method used to relate roughness to performance decline. Alternative assumptions about how wind turbine performance varies with age can be incorporated through modifications to the curves presented in Figure 4.6.

The framework for converting the roughness evolution parameter to AEP reduction over time gives important context to the impact of roughness on performance. By providing a framework method, it allows future researchers to refine the assumptions about the impacts of roughness on AEP reduction as additional data become available.

The proper use of science is not to conquer nature but to live in it.

— Barry Commoner

5

Wind Speed Variability and Wind Farm Degradation

Contents

5.1 Wind speed variability 81
5.1.1 Comparing results to the Weibull distribution 87
5.2 Reanalysis Modelling 88
5.3 Performance Decline of Wind Farms 90
5.4 Data Sources 92
5.4.1 Ofgem 93
5.4.2 EIA 94
5.4.3 AEMO 95
5.5 Data processing 95
5.5.1 Raw data filtering 97
5.5.2 Wind Speed Data 98
5.5.3 Missing hub height data 98
5.5.4 Scaling wind speed to hub height 100
5.5.5 Converting wind speed to power 101
5.5.6 Curtailment 107
5.5.7 Removing outliers 109
5.6 Conclusion 110

This chapter represents a transition from examining performance decline of a single wind turbine as in Chapters 4 and 3, to studying wind farm level degradation. This requires investigating additional factors beyond wind turbine blade roughness, such as wake effects, curtailment and turbine downtime. To examine wind farm

performance, the power output of a specific farm must be examined relative to the wind speed in that location. The non-linear relationship between wind speed and power output, combined with the highly variable nature of the wind resource, makes this investigation difficult. To complete this analysis, an understanding of both the wind resource at each site and previous work in this area are crucial.

Additionally, this chapter describes the data sources used for the following two chapters of this thesis, and how they have been extracted, cleaned and modified. Two data types are extracted and used to quantify wind farm performance decline (Chapter 6) and to predict wind farm CF (Chapter 7); wind farm metadata and time series data. Wind farm metadata describe static features of the wind farm (e.g. number of wind turbines, size of wind turbines, year of commencement), and time series data detail historical information (e.g. age, CF and reanalysis wind speed). Data are retrieved from government sources and reanalysis models. All data used in this thesis are publicly available, and the important metadata is detailed in Appendix A.

Specifically, this chapter analyses wind speed variability and its importance in wind farm performance degradation, while also providing a review of previous wind farm performance decline studies. To understand wind speed variability, data from weather stations that are located in close proximity to two Australian wind farms are used. This provides a framework for the importance of understanding, and accounting for, the wind resource in wind farm models.

Because weather station data are not always available in proximity to wind farm sites, reanalysis models are introduced, which provide the wind speed data used in this thesis. Global reanalysis models are hindsight models, thus they provide data availability for all wind farms over recent decades, whereas weather station data are (at best) only available post installation. However, reanalysis models produce wind speed data that are averaged over a time period (usually the highest resolution available is 1 hour) and therefore exclude the wind speed variability in that period. This variability therefore needs to be approximated with a wind speed distribution; however, this distribution will change shape depending

on the averaging period and mean wind speed value. This chapter will use weather station data to evaluate the impact of averaging period and of mean wind speed value on wind speed distributions and to examine how to integrate the wind speed distribution insights into reanalysis modelling.

Understanding the variability of the wind resource is important when considering wind farm performance decline, because the aerodynamic performance of the wind turbine needs to be separated from the long-term wind resource variability in order to understand how performance changes over time. Previous studies that quantify wind farm performance decline are reviewed, showing how reanalysis models can be combined with wind farm output and statistical techniques to analyse wind farm performance[1] over time.

5.1 Wind speed variability

Wind speed variability makes understanding long-term wind farm performance difficult. If a wind farm encounters long periods of low or high wind speeds it can cloud the true performance of the wind farm, especially if it is simply being compared to previous power output data. For example, a highly degraded wind farm could have a high wind speed period (compared to the long-term average) and produce a relatively high average CF value over that period, however when wind speed is accounted for it will reveal the long-term corrected performance is actually much lower. Accounting for wind speed is complicated by the different wind speed characteristics at each site (average wind speed, shape of wind distribution, etc.), meaning broad generalisations are not possible. Furthermore, the wind speed distribution at a given site changes depending on the time period observed – e.g. the wind speed over a 1-minute period is unlikely to reflect an annual average.

To understand the variability of wind distributions, weather station wind speed data were used from two sites near wind farms in Australia (AUS). Data from the Edithburgh weather station (approximately 2.5 kms from the Wattle Point Wind Farm) and the Swan Island weather station (approximately 10 kms from the

[1] In Chapters 5 to 7, performance is measured in CF values.

Musselroe Wind Farm) were obtained from the Australian Bureau of Meteorology (BOM). These two stations were chosen because of their close proximity to a wind farm and because there exists only flat ground or ocean between the weather station and the wind turbines. Figure 5.1 shows the location of these weather stations as well as the closest wind farm to them. Data were obtained for the time period 2010 to 2020.

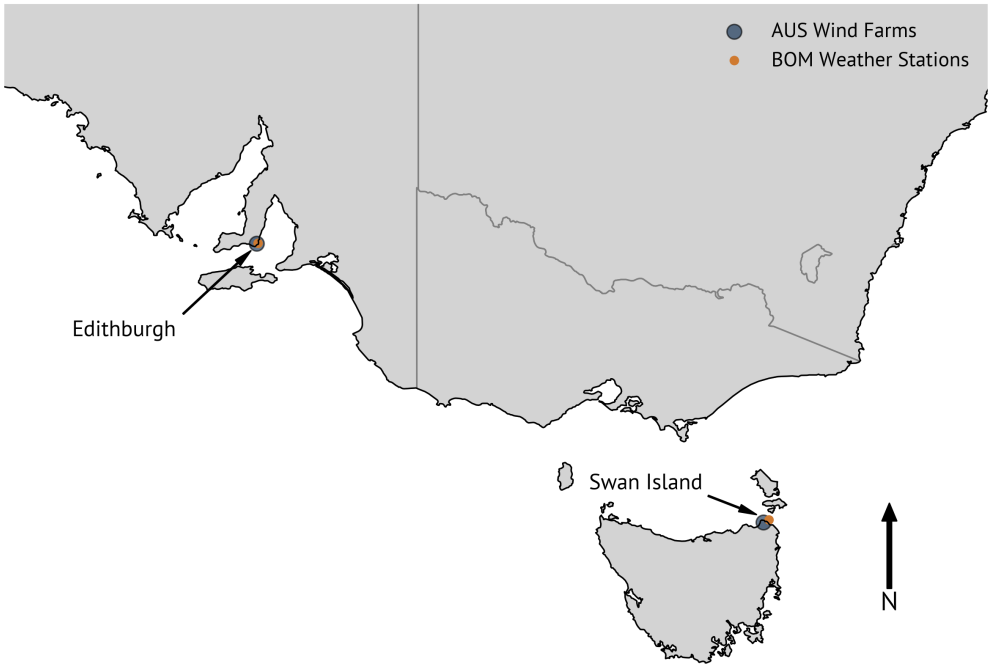


Figure 5.1: Location of the Edithburgh and Swan Island weather stations. The closest wind farm to these weather stations are also shown. This figure shows the south-east corner of Australia.

Data for the Australian wind farms were used because of the high time resolution data that are available. The BOM provides wind speed data, both magnitude and direction, at 1-minute resolution which can be averaged to 5-minute values to match the time resolution of the power output data. To ensure that the recorded weather station data were not being affected by the wake of the wind farm, only data where the wind was hitting the weather station before or at the same time as the wind farm were used.

To analyse the wind speed distributions from the Edithburgh and Swan Island weather stations, the time series wind speed data were processed using

the following steps:

1. Filter data for selected wind speed angles to ensure only wind data not impacted by the wind farm wake are used
2. Split time series wind speed data into selected time increment (1 hour or 1 month)
3. Bin each time segment into 1 m/s wind speed bins
4. Fit each wind speed bin with the most accurate distribution, as measured by the sum of squared errors and mean absolute error

The wind speed time series can then be split into bins based on different time averaging periods, as well as wind speed magnitude, to examine the shape of the wind speed distribution as these two parameters change. The two pertinent time increments for this study are 1 hour and 1 month, because 1 hour is the highest resolution of the reanalysis wind speed data used for AUS wind farms and 1 month is the time resolution of the United Kingdom (UK) and United States (US) wind farm power output data. Wind speeds were also separated by 1 m/s increments. To limit the number of bins, wind speeds less than 3 m/s were binned together, and wind speeds greater than 11 m/s were also binned together, because these represent the below cut-in wind speed and the approximate rated power wind speed for most wind turbines respectively.

Figure 5.2 shows the effect of different averaging periods on the wind speed distribution by plotting data with average wind speed between 7–8 m/s for the Swan Island weather station for the averaging periods of 1 hour, 1 day and 1 month. Each distribution has been fit by a best fit distribution, which will be explained below (in this section). This figure demonstrates the difference in shape of the wind speed distributions that occur for different averaging periods over the same average wind speed. It also shows how the spread of a wind distribution changes as the averaging period is varied. The wind speed distribution becomes more elongated as the averaging period increases and slightly positively skewed as the probability

of low wind speeds in the distribution increases. This is because a wider range of wind speed conditions will be recorded over a larger averaging period, and there cannot be negative wind speeds in the distribution so one tail of the distribution accumulates more values than the other.

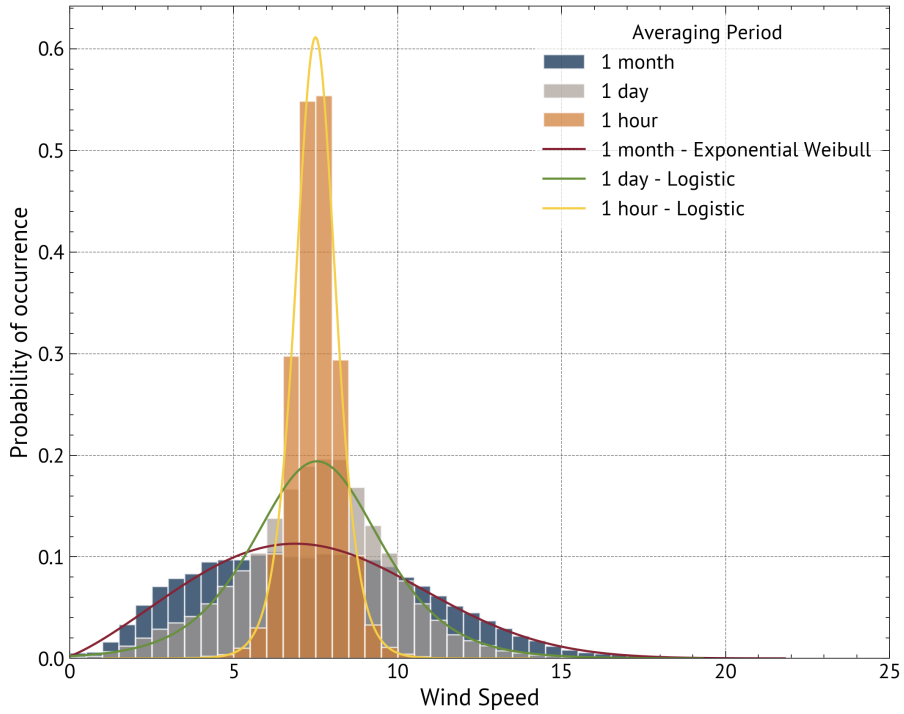


Figure 5.2: Wind distributions for average wind speed between 7–8 m/s over the averaging periods of 1 hour, 1 day and 1 month for the Swan Island weather station.

Figure 5.3 displays the impact of having different magnitude average wind speeds for the Swan Island weather station, with distributions for average wind speeds between 3–4 m/s, 6–7 m/s, and 9–10 m/s shown. The figure is also split into different averaging period subplots; 1 hour (a) and 1 day (b). The 1 month averaging period is not shown because there are not enough samples in all three wind speed magnitude bins to use. Again, the best fit distribution has been included in this figure and the method explaining this process is below (in this section). The shape of the distribution stays approximately constant when compared to changing the averaging period (Figure 5.2), but the location of the distribution changes significantly. The location of the plot changes due to the average wind speed recorded, e.g. if the average wind speed is 3 m/s, it is highly unlikely that there

will be a large proportion of values recorded above 10 m/s. At the 3 – 4 m/s wind speed for the 1 day averaging period the shape becomes slightly positively skewed due to an absence of negative wind speeds and therefore an increase in probability of values less than 1 m/s. For the larger time resolution averaging period, there is much greater overlap between distributions, which does not exist in the 1 hour case.

These two plots demonstrate how both the shape and location of distributions change significantly depending on the magnitude of the wind speed and the resolution of the averaging period.

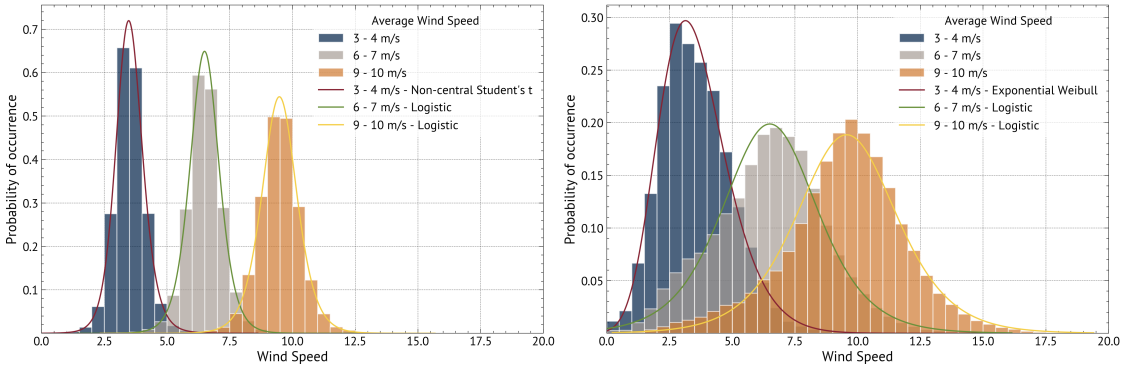


Figure 5.3: Wind distributions for varying average wind speeds (3–4 m/s, 6–7 m/s, and 9–10 m/s) over the averaging periods of 1 hour (left) and 1 day (right) for the Swan Island weather station.

13 continuous mathematical distributions were selected to fit each wind speed bin, each chosen to be a plausible wind speed distribution. A minimum sample size of 12 observations was used to ensure a representative distribution. A plausible wind speed distribution was considered to be one that is both continuous and passes through the origin. The distribution must pass through the origin because there cannot be a greater-than-zero probability of negative wind speeds (this would be recorded as a wind speed from the 180° opposite direction). Each distribution was fit to minimise both the sum of squared error and the mean absolute error. These metrics ensured the chosen curve fit was the most suitable in representing the underlying wind speed distribution. The chosen distributions were:

- Alpha distribution
- Chi distribution

- Exponential Weibull distribution
- Exponential power distribution
- Gamma distribution
- Inverted Weibull distribution
- Logistic distribution
- Non-central Student's t distribution
- Normal distribution
- Rayleigh distribution
- Student's t distribution:
- Weibull maximum distribution
- Weibull minimum distribution

Table 5.1 details the best fit curve (as measured by the sum of squared error) for each of the wind speed bins and both 1 hour and 1 month averaging periods for the Edithburgh and Swan Island weather stations. The 1 hour distribution best fit is dominated by logistic and Student's t distributions, whereas the 1 month distributions is predominately Weibull, exponential Weibull and inverse Weibull for the various wind speeds. This shows a clear difference between the averaging periods, given there is very little overlap between the distributions.

The same process was repeated over a larger range of averaging periods – 1 hour, 2 hours, 6 hours, 12 hours, 1 day, 2 days, 7 days, 14 days and 1 month. This was completed to analyse how the best fit distributions changed in between averaging periods ranging between 1 hour and 1 month. Across this many averaging periods, the most common distribution was the metric used instead of presenting each wind speed bin separately. Table 5.2 shows the most common distribution across the different averaging periods. The key similarities between Edithburgh and Swan Island are the change from Student's t and logistic best fit distributions for averaging periods of 2 days and less to Weibull variation best fits for averaging periods of 7 days and above. The distributions change because as the averaging period increases in length, the range of wind speeds broadens. This leads to a

Table 5.1: Best fit distributions for the Edithburgh and Swan Island weather station wind data for varying wind speeds and time increments. Where "N/A" exists, there were not enough distributions to fit (i.e. there were less than 12 months with average wind speed in this bin).

	Edithburgh weather station		**Swan Island weather station**	
Wind speed	**1 hour best fit**	**1 month best fit**	**1 hour best fit**	**1 month best fit**
< 3 m/s	Exponential Weibull	Non-central Student's t	Exponential Weibull	N/A
3 to 4 m/s	Non-central Student's t	Inverse Weibull	Non-central Student's t	N/A
4 to 5 m/s	Logistic	Inverse Weibull	Student's t	Exponential Weibull
5 to 6 m/s	Logistic	Weibull maximum	Logistic	Gamma
6 to 7 m/s	Logistic	Normal	Logistic	Chi
7 to 8 m/s	Logistic	N/A	Logistic	Exonential Weibull
8 to 9 m/s	Logistic	N/A	Logistic	Weibull minimum
9 to 10 m/s	Logistic	N/A	Logistic	Exonential Weibull
10 to 11 m/s	Logistic	N/A	Student's t	N/A
> 11 m/s	Exponential Weibull	N/A	Weibull maximum	N/A

higher probability of low wind speed values, which positively skew the data and cause the distribution to change shape.

As both the averaging period and the wind speed changes, different wind distributions are fit to the same dataset.

5.1.1 Comparing results to the Weibull distribution

Clearly, as the averaging period and mean wind speed change, the best fit distribution also varies. However, typically the Weibull distribution has been recognised as the most representative of wind speed distributions. This has largely held true in the

Table 5.2: Most common best fit distribution for the Edithburgh and Swan Island weather station wind data for varying averaging periods.

Averaging period	Edithburgh weather station most common fit	Swan Island weather station most common fit
1 hour	Logistic	Logistic
2 hours	Non-central Student's t	Student's t
6 hours	Student's t	Student's t
12 hours	Student's t	Student's t
1 day	Logistic	Non-central Student's t
2 days	Logistic	Logistic
7 days	Exponential Weibull	Non-central Student's t
14 days	Exponential Weibull	Weibull maximum
1 month	Exponential Weibull	Inverse Weibull

monthly averaging period analysis above; however for hourly data, distributions tend to be predominantly logistic or Student's t. By fitting the hourly wind speed distributions with both the best fit and a Weibull minimum distribution, a comparison can be drawn.

Figure 5.4 shows the comparison of the best fit (solid line) and Weibull minimum (dashed line) distributions across different magnitude average wind speed for the Swan Island and Edithburgh weather stations, with distributions for average wind speeds between 3–4 m/s, 6–7 m/s, and 9–10 m/s shown. While the Weibull minimum distribution is a less good fit (as measured by the sum of squared errors), it does roughly approximate the best fit distribution. The sum of squared errors for the different fits and locations are shown in Table 5.3. While the error is greater with the Weibull minimum distribution, it is valuable to approximate the distribution using this fit because not all wind farms will have access to weather station wind speed data and therefore longer time period distributions (Weibull) will need to be used.

5.2 Reanalysis Modelling

Measured wind speed data is not available for operational wind farm sites, and where it is available, it is at limited locations. Thus, wind farm performance decline

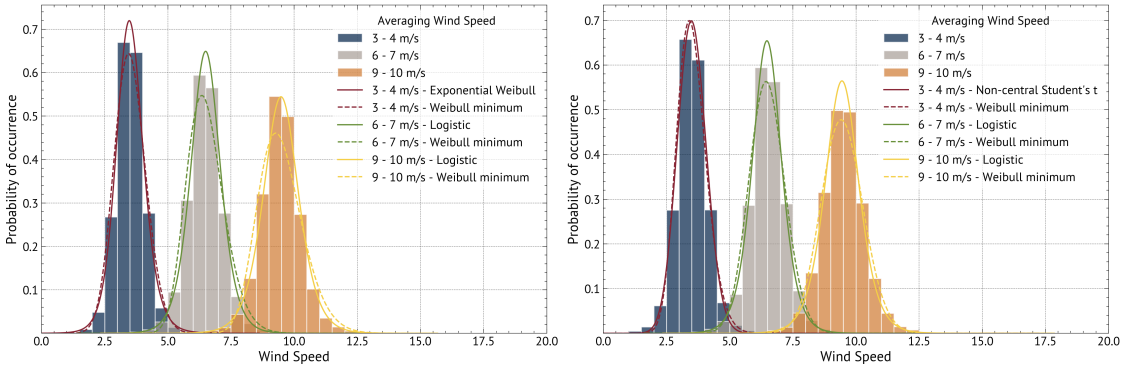


Figure 5.4: Wind distributions for varying average wind speeds (3–4 m/s, 6–7 m/s, and 9–10 m/s) over the averaging period of 1 hour fit with the best fit distribution indicated and a Weibull minimum distribution. The analysis is repeated for the Swan Island weather station (left) and Edithburgh weather station (right).

Table 5.3: Sum of squared errors (SSE) for the best fit distribution at varying wind speeds compared to a Weibull minimum fit for the Edithburgh and Swan Island weather station wind data.

	Edithburgh weather station		**Swan Island weather station**	
Wind speed	**Best fit SSE**	**Weibull minimum SSE**	**Best fit SSE**	**Weibull minimum SSE**
3 to 4 m/s	0.0025	0.0069	0.0004	0.0071
6 to 7 m/s	0.0003	0.0092	0.0006	0.0203
9 to 10 m/s	0.0002	0.0144	0.0003	0.0189

has been quantified using wind speeds from global atmospheric simulations, known as reanalysis models. Reanalysis models are constructed using global numerical weather prediction models that assimilate observations from a wide variety of sources including land surface stations, buoys, radiosonde balloons, aircraft and satellites [116, 117]. Due to their reliance on best-fitting to historical observation data, these are hindcast models and do not provide forecasts. Reanalysis data is, by construction, coarsely resolved and cannot represent small-scale wind fluctuations at a particular site [118]. Reanalysis models simulate many atmospheric and surface elements (cloud fraction, snow depth, soil temperature, etc.); however, only those relating to wind power production are considered here.

There are many reanalysis models which all perform slightly differently. In a recent review of reanalysis models' near-surface wind simulations, Ramon et al. [119]

investigated the relative performance of 5 state-of-the-art global reanalyses: the two European Centre for Medium-Range Weather Forecasts (ECMWF) reanalyses ERA5 [120] and ERA-Interim [121]; the Japanese 55-year Reanalysis (JRA55) [122]; NASA's Modern Era Analysis for Research and Applications-2 (MERRA-2) [117]; and National Centers for Environmental Prediction (NCEP)/National Center for Atmospheric Research (NCAR) Reanalysis 1 (R1) [123]. Table 5.4 lists an overview of the main characteristics of the reanalysis models concerning the coverage periods, data assimilation schemes, spatiotemporal resolution, and wind height levels. Ramon et al. compared the accuracy of reanalysis wind speeds to both surface and hub-height wind speed observations measured in 77 tall meteorological towers (also known as met masts) globally. The heights of the wind speed recordings at these met masts ranged between 18m and 122m, which aimed to cover a range of possible wind turbine hub heights. Ramon et al. concluded that comparison with in-situ observations show that the ERA5 surface winds offered the best agreement, correlating and reproducing the observed variability better than a multi-reanalysis mean in 35.1% of the tall tower sites on a daily time-scale. MERRA-2 was the next best performing reanalysis model and is the other predominant model used for wind speed estimates in academic studies.

To produce the global wind speed datasets, these models combine wind speed measurements with mesoscale models. For MERRA-2, wind observations with "fairly complete global coverage" are taken from weather stations, balloons, aircraft, ships, buoys and satellites, and processed by the model to give data with hourly resolution [117].

5.3 Performance Decline of Wind Farms

Various studies have used reanalysis wind speed data to predict actual turbine output and compare it with recorded values. Kiss et al. [124] appears to be first to compare nacelle measurements of wind speed and power generation from two

Table 5.4: Summary of the basic characteristics of 5 state-of-the-art reanalysis datasets.

Model Name	Institution	Period Coverage	Time Resolution	Assimilation Scheme	Horizontal Grid Spacing	Wind Heights
ERA-Interim	ECMWF	1979 to present	6 hour	4D-VAR	0.75° × 0.75°	Surface, 10m
ERA5	ECMWF	1980 to present	1 hour	4D-VAR	0.3° × 0.3°	Surface, 10m, 100m
JRA55	JMA	1978 to present	6 hour	4D-VAR	1.25° × 1.25°	Surface, 10m
MERRA-2	NASA GMAO	1980 to present	1 hour	3D-VAR-FGAT	0.5° latitude ×0.65° longitude	Surface, 10m, 50m
R1	NOAA/ NCEP and NCAR	1948 to present	6 hour	3D-VAR	1.875° latitude ×2° longitude	Surface, 10m

turbines in Hungary to the ERA-40 [125] reanalysis data set, finding "surprisingly good" agreement. Hawkins et al. [126] reproduced monthly load factors for the UK from a reanalysis data set calculated with the WRF (Weather Research and Forecast) model. Kubik et al.[127] matched the global NASA reanalysis to half-hourly power output from wind farms located in Northern Ireland, finding it to be more accurate than met mast data. Reanalysis models have since been used by multiple authors to produce wind resource data and model wind farm power production [128–130].

Staffell and Green [18] used reanalysis wind speed data to characterise performance decline of wind farms as they age. This study examined the rate of ageing of a national fleet of wind turbines using publicly available data for the actual and theoretical capacity/load factors from 282 UK wind farms. They used capacity factors, recorded monthly for the period of 2002–2012, and derived capacity factors based on MERRA-2 wind speeds at the location and hub height of each wind farm, accounting for the power curves of the specific models of turbine installed at each location/site. By accounting for individual site conditions, they concluded that average capacity factors across all 282 UK wind farms decline by a factor

of 1.6 ± 0.2% of their output per year, with average load factors declining from 28.5% when new to 21% at age 19 years.

The approach adopted in the Staffell and Green study was used by Olauson et al. [19] who completed a similar study, looking at the performance decline of wind turbines in Sweden. They found that Swedish wind turbines constructed before 2007 lose 0.15 capacity factor percentage points per year, corresponding to a lifetime energy loss of 6%, where the lifetime was defined as 20 years.

5.4 Data Sources

Using wind farm field data to understand and predict performance is commonplace in industry. Field data are the most important and relevant data for any wind farm and can be used to understand the impact of ageing on a given wind farm. However, most observed wind farm data from turbine supervisory control and data acquisition (SCADA) systems or meteorological masts, are not made publicly available [18]. However, wind farm electricity generation data are often made public through the governing electricity market body at a reduced resolution (an aggregated number for the entire wind farm, not individual turbines) and reduced data frequency than SCADA systems. Generation data are retrieved through government bodies; the Australian Energy Market Operator (AEMO) [131], the UK Office of Gas and Electricity Markets (Ofgem) [132] and the US Energy Information Administration (EIA) [133]. For each country, only wind farms with more than 5 years of data available were used in this study to ensure consistency with previous studies [18, 19] and avoid extracting false performance trends over a short period of time.

To extract insights into wind farm performance, the wind speed needs to be known for every data point. To do this, wind farm electricity data can be combined with wind speed data from weather models called reanalysis models. Reanalysis models are coarsely resolved weather models that are constructed using numerical weather prediction models that use observations from a wide variety of sources such as land weather stations and ocean buoys. For this study, two global reanalysis

datasets are used, MERRA-2 [117] and ERA5 [120], which ensure data availability and consistency for all wind farms.

A meta-dataset also needs to be produced for each wind farm, containing information such as wind farm rated power, size of individual turbines in a wind farm, number of turbines in a wind farm, location of a wind farm, data frequency, etc. This is important for many reasons, e.g., location is used to get the reanalysis wind speed time series for a given wind farm, time frequency is used to know how to extract and process the reanalysis model, etc. In general, these data are also provided by the government agencies; however, they are often self-reported and can become inaccurate if any details have changed over time, such as through expansion or re-powering of the wind farm. For this study, the datasets from government bodies were used as a starting point, and all entries were individually checked with reputable sources such as the wind farm operator website or turbine provider website.

5.4.1 Ofgem

All UK wind farms publish their monthly energy outputs in the Ofgem Renewables and CHP (combined heat and power) Register[2]. Data were retrieved from the beginning of the Ofgem records in 2003 to the start of 2020. This source of data has been used previously by Hughes [134], Staffell and Green [18] and Crabtree et al. [135] to examine wind farm power output over time.

The generation data have to be combined with metadata about each wind farm (size, commission date, etc), and capacity factors can then be calculated. This process was first completed by Hughes [134] and later improved by Staffell and Green [18], with their results publicly available (over a different time period compared to this thesis). However, for this study, metadata were compiled from scratch, by manually checking each wind farm's information from operator's websites or from central repositories, such as the website thewindpower.net. This was done to ensure the reliability of the data and to include the more recent wind farms.

[2]https://www.renewablesandchp.ofgem.gov.uk/

Figure 5.5 shows the location, size and start year of the 325 UK wind farms used in this study, of which 19 are offshore.

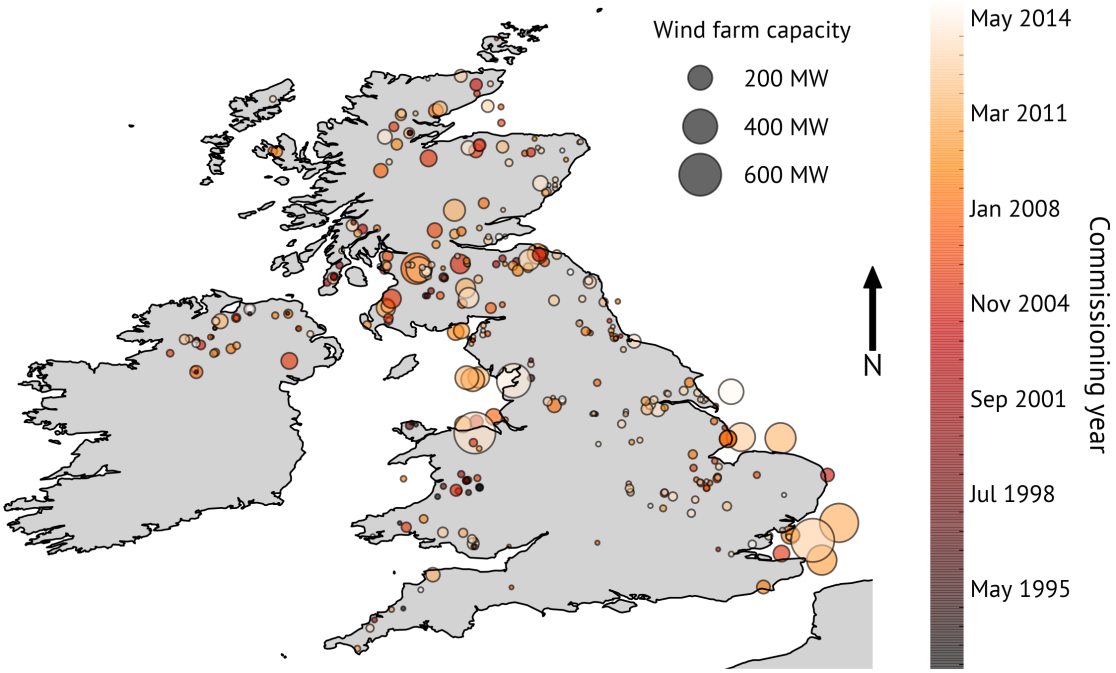


Figure 5.5: Location, size and commencement year information for UK wind farms that had at least 5 years of operation to 2020 used in this thesis.

5.4.2 EIA

Similarly, US wind farms publish their monthly output data using form EIA-923[3] (with previous data from EIA-906/920) and the wind farm metadata are available through form EIA-860[4]. These sources of data have been used previously by Hamilton et al. [21] to examine wind farm performance decline. The data for this study were taken for years 2001–2018, with 2018 being the most recent complete set of data for both forms at the time of writing. Figure 5.6 shows the location, size and start year of the 696 US wind farms (675 active and 21 retired) used in this study.

[3]https://www.eia.gov/electricity/data/eia923/
[4]https://www.eia.gov/electricity/data/eia860/

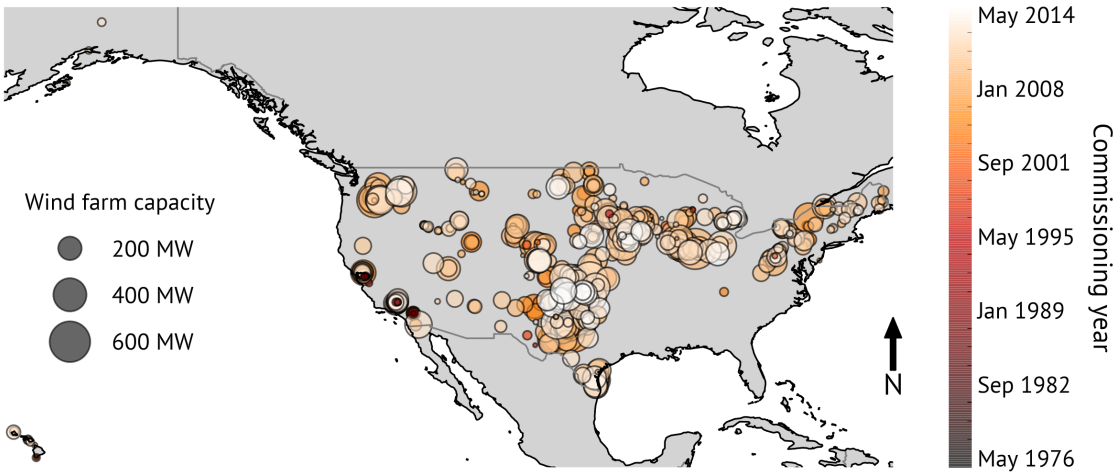


Figure 5.6: Location, size and commencement year information for US wind farms that had at least 5 years of operation by 2020 used in this thesis.

5.4.3 AEMO

Data about the Australian national electricity market (NEM) are made publicly available through the NEMWEB portal[5], from which the wind farm power data have been retrieved. This dataset contains five minute generation values starting from July 2009. For this study, data for 32 wind farms were used from 1 July 2009 to 1 Jan 2020. The metadata for the AUS wind farms are taken from AEMO's 'NEM Registration and Exemption List'[6] with each wind farm individually verified through its site page and additional information, such as location and turbine model, added. These sources of data have been used previously by Xenophon and Hill [136] to create a model of the Australian electricity market. Figure 5.7 shows the location, size and start year of the 32 Australian wind farms used in this study.

5.5 Data processing

This section details the data processing tasks that were undertaken for robust data usage and to convert raw data to long-term corrected data.

[5]https://aemo.com.au/en/energy-systems/electricity/national-electricity-market-nem/data-nem/market-data-nemweb

[6]https://aemo.com.au/en/energy-systems/electricity/national-electricity-market-nem/participate-in-the-market/registration



Figure 5.7: Location, size and commissioning year information for Australian wind farms that had at least 5 years of operation by 2020 used in this thesis. This figure shows the south-east corner of Australia.

Table 5.5 summarises the key statistics of each data set, and Figure 5.8 also provides a selection of summary statistics for the data and the mean capacity factor for each region. The mean capacity factor is calculated by averaging across wind farm mean capacity factors to avoid the data frequency affecting the calculation.

Table 5.5: Summary of the key statistics for the wind farms used in this study, separated by country – AUS, UK, US and ALL.

	AUS	**UK**	**US**	**ALL**
Temporal resolution	5 minute	1 month	1 month	N/A
Number of wind farms	32	325	696	1053
Total capacity (GW)	3.26	12.12	65.52	80.91
Average turbine size (MW)	2.15	2.04	1.71	1.82
Average number of turbines	46.31	16.77	62.15	47.66
Average farm size (MW)	101.87	37.30	94.14	76.84
Commission date period	2003–2014	1992–2014	1975–2015	1975–2015
Data recording period	2009–2020	2003–2020	2001–2019	2001–2020

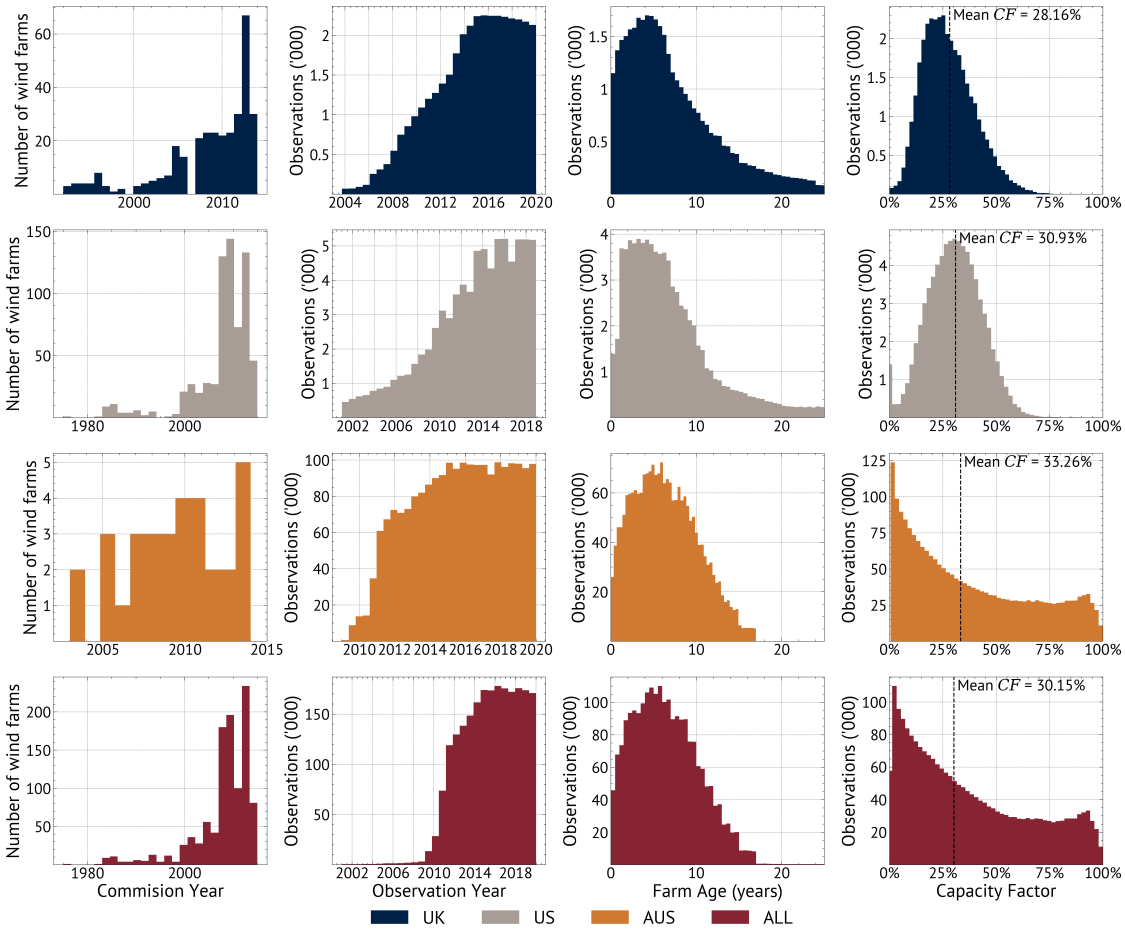


Figure 5.8: Summary data for the wind farms used in this study. Figures in each column represent the number of wind farms, number of observations by year, number of observations by farm age and distribution of capacity factors. Each row represents the different countries – UK, US, AUS, and ALL, from top to bottom. UK and US observations are monthly, AUS observations are hourly and ALL is a blend of monthly and hourly.

5.5.1 Raw data filtering

In the UK data downloaded from Ofgem (Section 5.4.1), errors were found in the generation data that were often due to the aggregation of multiple months, leaving a single month with a CF greater than 100%. These months were discarded from the data, as averaging the generation across these months would result in a flat output, which would not match the fluctuations in generation seen in reality. This inhibits the ability to assess how performance changes over these periods and would therefore not achieve the purpose of this study.

Similar to the UK wind farms, errors found in the US generation data retrieved

from EIA (Section 5.4.2) due to aggregation of multiple months were discarded. The metadata were manually checked and adjusted for reporting errors such as wind farms that had been repowered with more turbines at a later date.

5.5.2 Wind Speed Data

In this thesis, reanalysis models are used to produce wind speed data for each wind farm site. An overview of reanalysis models is provided in Section 5.2. As discussed in this chapter, ERA5 and MERRA-2 were the best performing models in previous literature and were chosen for this study. The data were extracted from their respective online portals.

Once reanalysis data were downloaded for each country over the required time horizon, the following steps were taken for each wind farm:

- For a given wind farm's location, a time series is created (for each hour between the wind farm's commission date and current date/closing date) by interpolating the two components of wind speed, north (v) and east (u), using the nearest four grid points at the highest two available heights (10m and 50m for MERRA-2 and 10m and 100m for ERA5).
- The wind speed magnitude (V) is calculated using $V = \sqrt{v^2 + u^2}$
- The wind speed is then calculated at the given wind farm's hub height using the log law (Equation 3.4) defined in Section 5.5.4.

5.5.3 Missing hub height data

It was not possible to find hub height data for all wind turbines in the dataset, which is required in order to scale the reanalysis wind speed correctly. A total of 138 out of the 1053 wind farms were missing hub height information for their turbines. Predicting the missing hub heights is a non-linear process because the hub height does not depend solely on the size of the turbine, although it is correlated. For example, turbines with the same rated power can be deployed at different hub heights in different locations. Other factors such as the number of turbines in a

wind farm or local environmental factors and regulations could also affect the hub height. Due to this, a gradient-boosting supervised machine learning algorithm [137] was used to predict the missing turbine hub heights[7].

This technique involves an ensemble of decision trees to predict a result, in this case turbine hub height. A decision tree is a way of breaking a problem into subsets depending on conditions. For a gradient-boosting algorithm, the complete dataset is divided into smaller subsets while at the same time an associated decision tree is incrementally developed. This allows non-linear interactions between variables to be more easily captured. A gradient-boosting algorithm has a number of changeable parameters: number of estimators, learning rate, sub sample size, max depth, and max features. A parameter space for these parameters was iterated over and the highest accuracy model was used from this set.

A gradient-boosting algorithm identifies patterns in the dataset based on real data (i.e. known inputs and outputs). The model is able to train on the known data to better capture trends and predict unknown data. For this to be effective, the training dataset must be reflective of the predicted data. To ensure the data are not over-fitted, a subset of the data is set aside that the model never 'sees' in its training process. This subset is then used to test the accuracy of the trained model. Once the best model, measured in terms of MAE, for the given dataset has been identified, the unknown values can then be predicted.

To complete this process, the input variables used were: wind turbine rated capacity, number of installed turbines in a particular wind farm, wind farm start year, wind farm country and wind farm mean capacity factor. 80% of the meta dataset with available hub height data (randomly sampled) was used for training the model and 20% for validation. The model was tested repeatedly to ensure the random sampling was different and did not affect the model accuracy. The model had a MAE of 4.4 m, which is 6% of the mean hub height of all wind farms. The errors were distributed randomly as demonstrated in Figure 5.9, which shows a scatter plot of observed vs predicted values using this gradient boosting model.

[7]For example implementation, see https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.GradientB

These errors are considered acceptable for the purpose of this thesis for two reasons: the average blade diameter in the predicted dataset is 79.24 m, meaning the MAE of 4.4 m is not material, and the wind speed interpolation difference between altitudes of approximately 4 m is negligible when at typical hub heights.

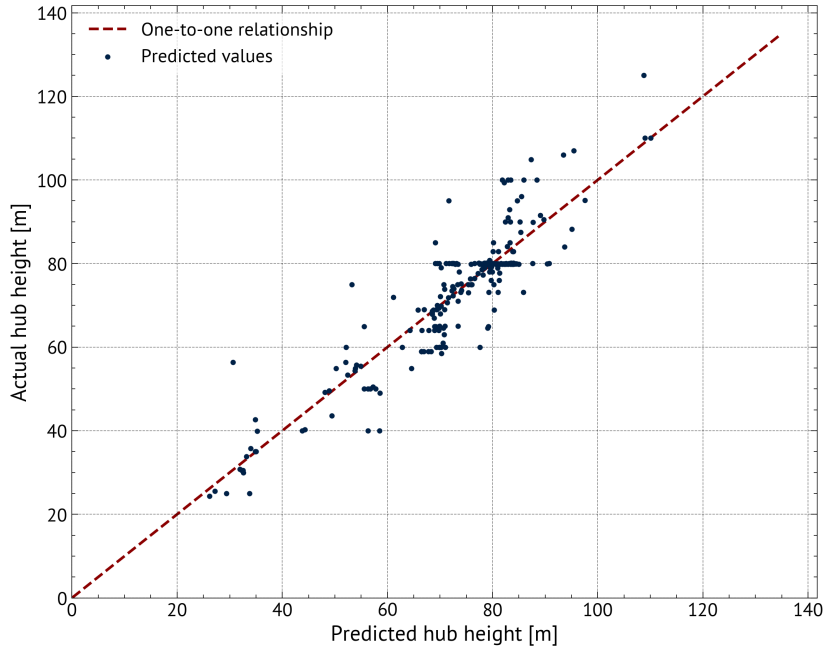


Figure 5.9: Scatter plot of predicted hub heights from the validation dataset of the gradient-boosting algorithm used to predict unknown hub heights.

5.5.4 Scaling wind speed to hub height

In general, wind speeds are of higher magnitude at higher altitude, making it important to accurately scale the reanalysis wind speeds to turbine hub-height to correctly predict power output. Both the ERA5 and MERRA-2 reanalysis wind speed data are available at two altitudes above the surface – 10m and 100m for ERA5, and 10m and 50m for MERRA-2. Typically, wind speed data can be interpolated using the empirically derived Hellman power law [138] (used by [19]) or the theoretically-derived log law profile [139] (used by [18]). If the hub height is higher than the higher reanalysis speed altitude, then the data are extrapolated to hub height[8].

[8]For the purpose of this work, the word 'interpolate' will be used when referring to scaling reanalysis wind speed data to hub height, however if the hub height is above the highest reanalysis

In this study, reanalysis wind speed is interpolated to a given wind farm's hub height using the log law under the assumption of a neutral atmospheric boundary layer. Using the wind speeds available at two altitudes from the reanalysis data (10 m and 100 m for ERA5, 10 m and 50 m for MERRA-2), the wind speed at hub height z can be calculated using the log law [32], which is defined in Equation 3.4.

Reanalysis data provided close to the surface (2 metres for MERRA-2, not provided for ERA5) were not used in this calculation due to the potential error that could occur given that local topology is not included in the model. The roughness height, z_0, is estimated for each time period and used to then predict the hub height wind speed for that time period. It is necessary to estimate z_0 for each time period to capture the effect of different wind directions.

This is completed for each hourly time period to capture the different wind shear profiles and to create a wind speed time series at hub height. For periods where the wind speed at the lower altitude was a greater magnitude than the higher altitude value, the neutrally-stable atmospheric boundary layer assumption is violated, and a linear interpolation was used (note this occurs for $< 0.9\%$ of time periods for MERRA-2 and $< 0.4\%$ of time periods for ERA5).

5.5.5 Converting wind speed to power

Power is calculated using a specific turbine power curve and wind distribution for each wind farm. Power curves are assigned to the turbines in each wind farm to facilitate the calculation of power given the wind speed. Publicly available power curves are taken from renewables.ninja (an extension of [140]), and matched with the wind farm metadata. Using this data, power curves were compiled for 97.7% of wind farms. Where the wind turbine power curve was missing, the closest matched power curve based on rated power and rotor diameter was chosen.

The power curve of a wind turbine will look different from the power curve of a wind farm due to wake interactions between turbines. This means that wind turbines that are not in the front row of a wind farm will experience a lower wind

wind speed, note that the data were technically extrapolated.

speed. Consequently, for a given wind speed (between cut-in wind speed and rated wind speed) a wind farm would have a slightly lower CF value than a single turbine. This means that in order to convert a wind turbine power curve to a wind farm power curve, modifications need to occur to reflect the lower CF at certain wind speeds. This is completed using a rolling average method called 'smoothing' which results in a more gradual transition between cut-in power and rated power. This also results in a slightly delayed cut-out wind speed because of the variation in observed wind speeds throughout the farm.

The power curves are smoothed to partially account for the effect of having multiple turbines within a wind farm. The importance of smoothing in the context of long-term correction (explained in Section 6.4) is to more accurately describe the region where power is less than rated (where $CF < 100\%$ in Figure 5.10). The power curve shape is a particularly important characteristic to capture for the LTC process.

To smooth the curve, it is assumed that within a farm, for a given reanalysis wind speed, there is a normally-distributed variation in wind speed similar to [18, 141]. The normal distribution accounts for fluctuations in wind speed due to turbulence and wind direction changes, but it assumes that wind speed recovers rather than continuously decreasing throughout the wind farm due to wake effects. However, in this work, it is assumed that the reanalysis wind speed is the highest wind speed (seen by the front row of turbines) and that the remainder of the normally-distributed wind speed is slower due to wake effects. This ensures that power is not produced before the cut-in wind speed and that at no point below rated power would the normalised wind farm power exceed that of an individual, unwaked wind turbine. In this thesis, the standard deviation used for each wind farm linearly varied depending on the number of turbines in that wind farm. An example of a normalised wind farm power curve (CF curve) compared to an individual turbine using an exemplar wind turbine power curve is shown in Figure 5.10.

When calculating energy production for a given time period, the power variation through the entire time period (one month for the UK and US, or one hour for

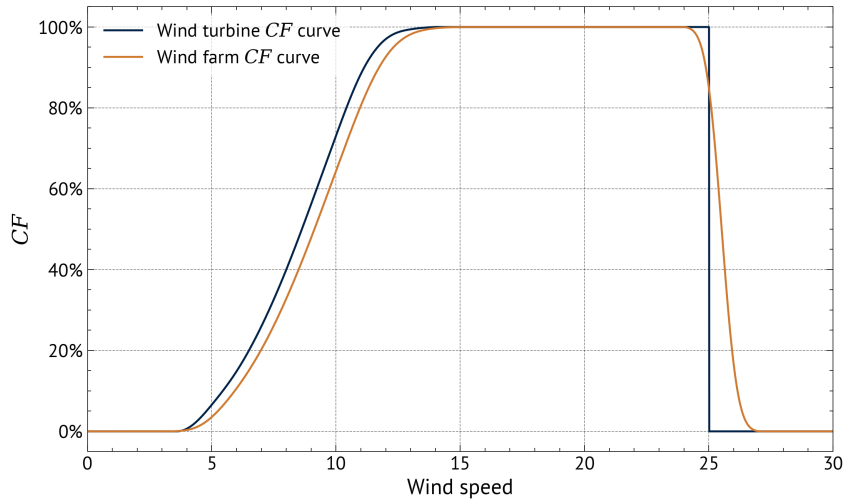


Figure 5.10: Capacity factor as a function of wind speed for an individual wind turbine (blue) and compared to a wind farm (orange) for an exemplar turbine CF curve.

AUS) must be represented. To calculate energy production for a given time period, wind speed is integrated over using Equation 5.1:

$$CF = \int_{U_{ci}}^{U_{co}} CF(U) f(U) dU \tag{5.1}$$

where $CF(U)$ represents the capacity factor as a function of wind speed (the CF curve) and $f(U)$ represents the likelihood of a wind speed U occurring at a particular site (the wind speed distribution).

As a result, to better estimate energy throughout the period, a wind speed distribution ($f(U)$) is used to capture variations in wind speed instead of the single, period-averaged wind speed value provided by the reanalysis model. The distribution used for each wind farm is based off the lifetime reanalysis wind speed distribution at that site and is modified depending on the magnitude of the reanalysis wind speed during each time period. Modifying the distribution according to the mean wind speed creates a more realistic wind speed distribution for each time period based on the available reanalysis data.

Section 5.1 showed that the best fit wind speed distribution varies based on the averaging period and mean wind speed. The wind speed distribution used to predict CF (in combination with the wind farm CF curve) for each time period is a Weibull distribution that is specific to each wind farm location and modified

according to the mean wind speed in each time period. This was chosen because it is a consistent method across all wind farms, it is representative of the long term distribution at that location, and even for the 1 hourly data, a modified Weibull approximates the best fit distribution (see Figure 5.4). This is an approximation for the hourly data; however, without access to weather station data at all sites, it was not possible to accurately complete a more thorough wind speed distribution analysis. Instead, the learnings from Section 5.1 were adapted to the reanalysis wind speed datasets used in this chapter.

To produce this Weibull distribution, first the hourly reanalysis wind speed data for the lifetime of the wind farm are fitted with a Weibull distribution [93]:

$$f(U; k_s, c_s, l_s) = \frac{k_s}{c_s}\left(\left(\frac{U - l_s}{c_s}\right)^{k_s - 1}\right) \exp\left(-\left(\frac{U - l_s}{c_s}\right)^{k_s}\right) \tag{5.2}$$

where l_s is the location parameter. Note that for all lifetime fitted distributions $l_s = 0$; however, this is a useful and potentially changeable parameter when calculating power for each time period.

Figure 5.11 (a) shows the fitted Weibull distribution for the Woolnorth wind farm in Australia. The shape and scale parameters for each wind farm are stored with the turbine metadata and are used when calculating wind farm power. Once the lifetime wind speed distribution has been determined, it is then modified for each time period depending on the mean reanalysis wind speed for that period. The wind distribution is adjusted according to the following rules:

- If the wind speed in time period t for a given wind farm w is greater than the average wind speed for the lifetime of that wind farm ($U_{w,t} > \bar{U}_w$) then the location parameter is increased by $l_{s,t} = U_{w,t} - \bar{U}_w$. An example of this is the dotted grey curve in Figure 5.11 (b).

- If $U_{w,t} < \bar{U}_w$, the scale parameter $c_{s,t}$ is recalculated based on the $U_{w,t}$ value using Equation 5.3 below. An example of this is the dashed blue curve in Figure 5.11 (b).

- In all cases the shape parameter stays constant as this characterises the specific wind farm site.

The scale factor for a given average wind speed can be calculated using the following equation from [142]:

$$c_{s,t} = \frac{\Gamma(1 + \frac{1}{k_s})}{U_{w,t}} \tag{5.3}$$

where Γ is the Gamma function [143].

When $U_{w,t} > \bar{U}_w$ the distribution is simply shifted to higher wind speeds by increasing the location parameter. This maintains the shape of the distribution and gives an approximation of the expected wind speed distribution during this time period. If the scale parameter is changed, it would 'stretch' the curve over a larger range of wind speeds and would therefore have a non-zero probability of almost every wind speed where the wind turbine is operational for that period. This is a less realistic approach as it limits the likelihood of the wind farm operating at its rated capacity at high wind speeds. This is supported by Figure 5.3 which shows that for higher average wind speeds, the location (mean) of wind speed distributions are largely shifted to higher values with minimal change to the shape of the distribution.

Adjusting the scale parameter when $U_{w,t} < \bar{U}_w$ corresponds to a 'tighter' spread of wind speeds at lower average speeds. The distribution cannot be shifted to lower speeds as in the case where $U_{w,t} > \bar{U}_w$, because this would create probabilities of negative wind speeds, which is by definition impossible. Consequently, only the scale parameter is changed in this scenario. This is supported by Figure 5.3 which shows that for lower average wind speeds, the wind speed distributions are narrower with higher peak probabilities and have minimal change to the shape of the distribution.

Once the wind distribution is defined for a given time period, CF can be calculated by numerical integration of Equation 5.1. The estimated CF output is then used in the LTC process described in Section 6.4. An overview of the CF production decision process is provided in Figure 5.12.

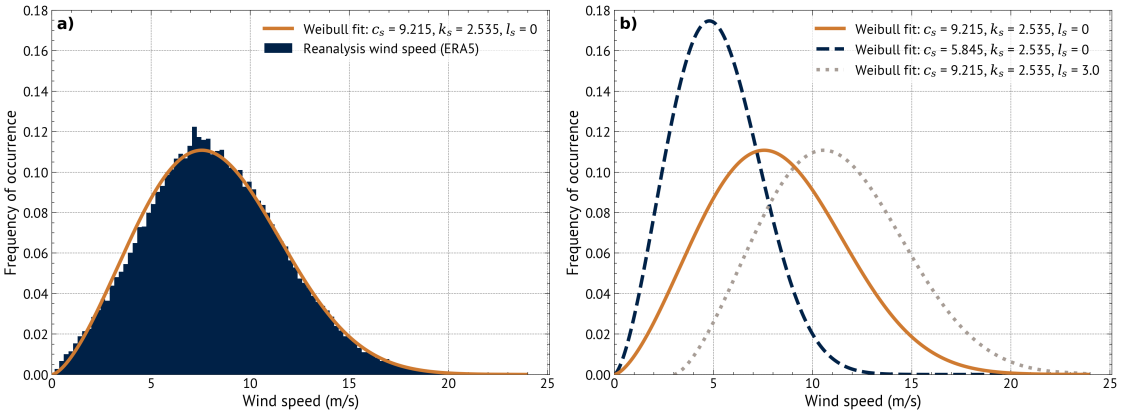


Figure 5.11: (a) Reanalysis hourly wind speeds for Woolnorth Wind Farm (from 2009 to 2020) in Australia, with Weibull distribution fit. (b) The resulting wind speed distributions used for estimating CF when the hourly wind speed $U = 8.19$m/s (orange curve), decreases by 3m/s (blue curve) and increases by 3m/s (grey curve).

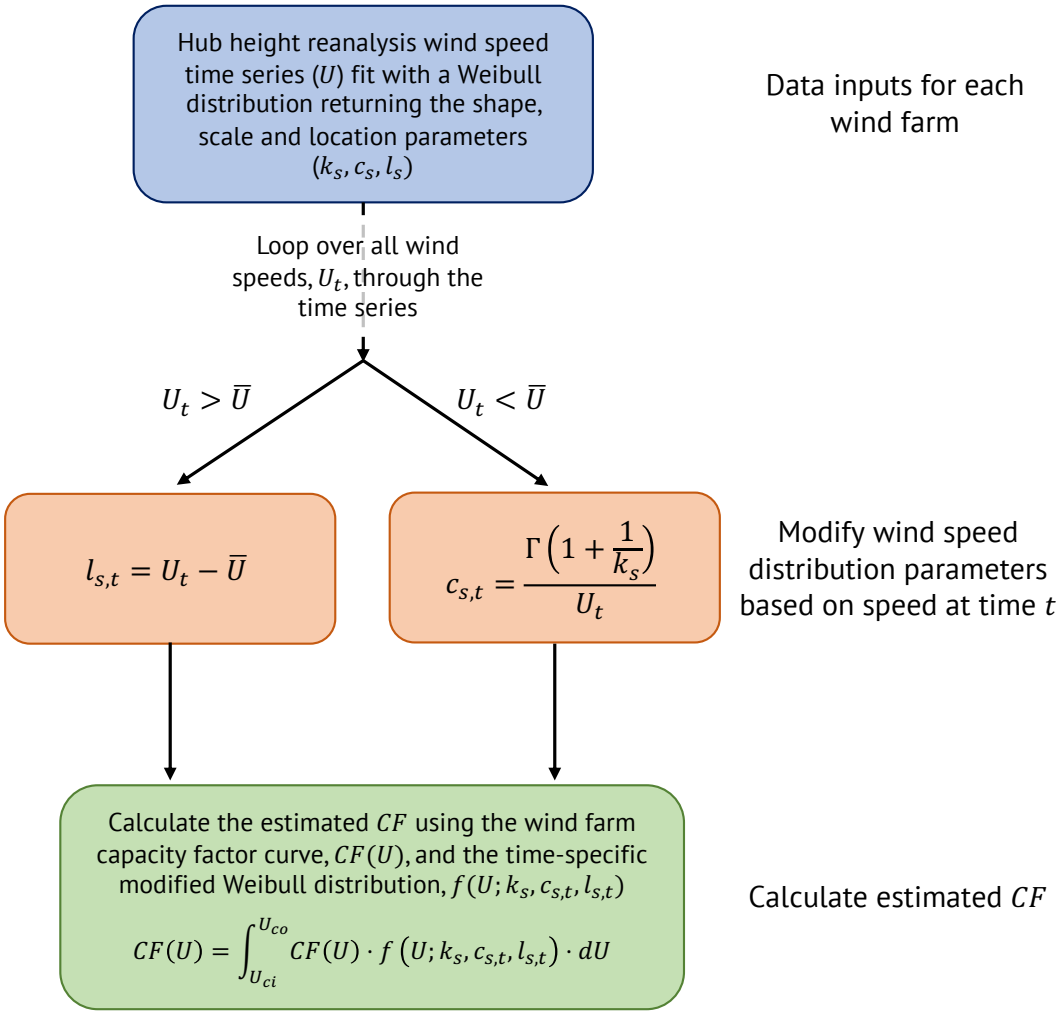


Figure 5.12: An overview of the CF production process.

5.5.6 Curtailment

The use of public electricity production datasets has the inbuilt risk of not always representing potential wind farm power output (i.e. electricity production does not equal potential wind farm power output due to wind speed). A particular wind farm may be subject to a network outage or limitations, or a low demand event may cause its production to be curtailed. Wind farms may even choose to curtail production if electricity prices are negative. Curtailment has become especially important in recent years because of the increasing penetration of variable renewable energy. Renewables cannot be scheduled in the same way that fossil fuel generators can, leading to periods of over supply and therefore the need for curtailment.

Curtailment data are available for semi-scheduled (a type of registration) generators in Australia (19 out of 33 wind farms in the dataset). Using this data, the increasing impact of curtailment is examined by looking at the percentage of wind farm rated capacity that is curtailed, on average, during a curtailment event. Figure 5.13 shows the increasing percentage of wind farm capacity that is curtailed annually. If not taken into account, curtailment will lead to an exaggeration of the performance decline rate of a given wind farm, because the reduced performance for a given wind speed stems not from the wind farm's lack of production power but from its deliberate reduction for non-aerodynamic reasons.

Hamilton et al. [21] considers curtailment when investigating performance decline of US wind farms, by adjusting power output nationally based on curtailment data in the Texas (Electric Reliability Council of Texas – ERCOT) grid. They assume that curtailment in Texas reflects the national fleet by extrapolating curtailment data from Texas to the national fleet. However, the opportunity to generalise the Texan data is limited because curtailment is often location-specific and depends on demand, generation and grid constraints. The ERCOT grid also has very minimal interconnection with surrounding states, meaning that network constraints in Texas may not reflect other regions of the US.

Curtailment data available for semi-scheduled generators in Australia have not been used to adjust data for several reasons. Firstly, the data do not cover all

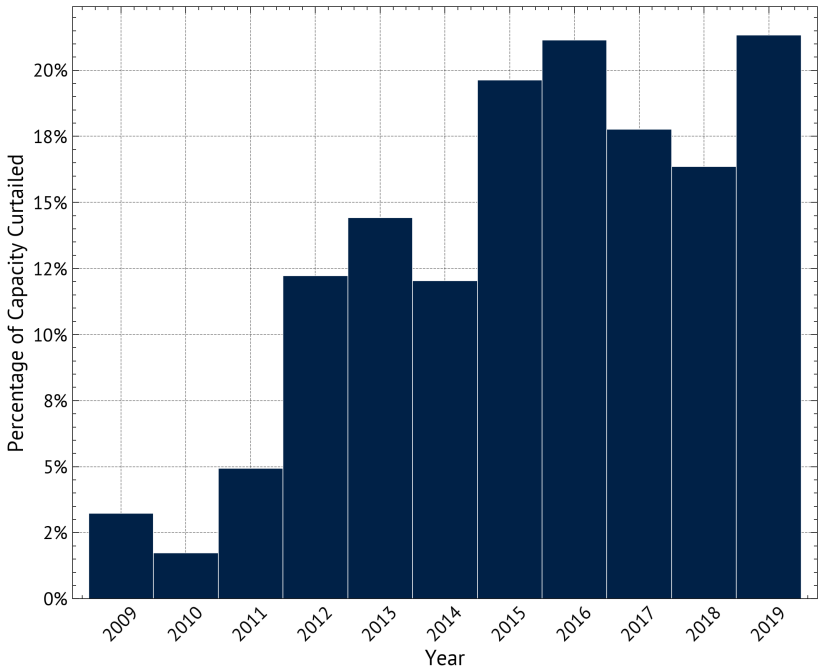


Figure 5.13: Percentage of wind farm rated capacity, on average, that is curtailed during recorded curtailment events annually for applicable Australian wind farms. Applicable Australian wind farms are those who are required to produce curtailment data (semi-scheduled generators, 19 out of 33 wind farms in the dataset).

Australian wind farms, and due to the highly localised nature of curtailment, it would not be appropriate to extrapolate the existing curtailment data to the farms where these data are not available. The Australian grid is very spatially diverse, and constraints can vary widely across wind farms at any given point in time, even if they are experiencing the same price signal. Further, because the output data from AEMO are electricity generation data, they only capture when the grid operator is curtailing the wind farm, not when the wind farm is self-curtailed (e.g. in times of negative prices). The curtailment data are produced by comparing the forecasted generation for the wind farm (what it is capable of producing given the forecast wind conditions) with the recorded generation in each time period. However, the forecasted generation data are estimated values and are therefore not 100% accurate, especially in large curtailment events where the wind farm has been curtailed for an extended period of time. Despite these points, the majority of curtailment events can be captured using these data, and it is a useful point to acknowledge, especially for future LTC studies where curtailment will continue to

increase. Figure 5.14 shows the curtailment events for the Musselroe Wind Farm which appear as bands in the capacity factor data.

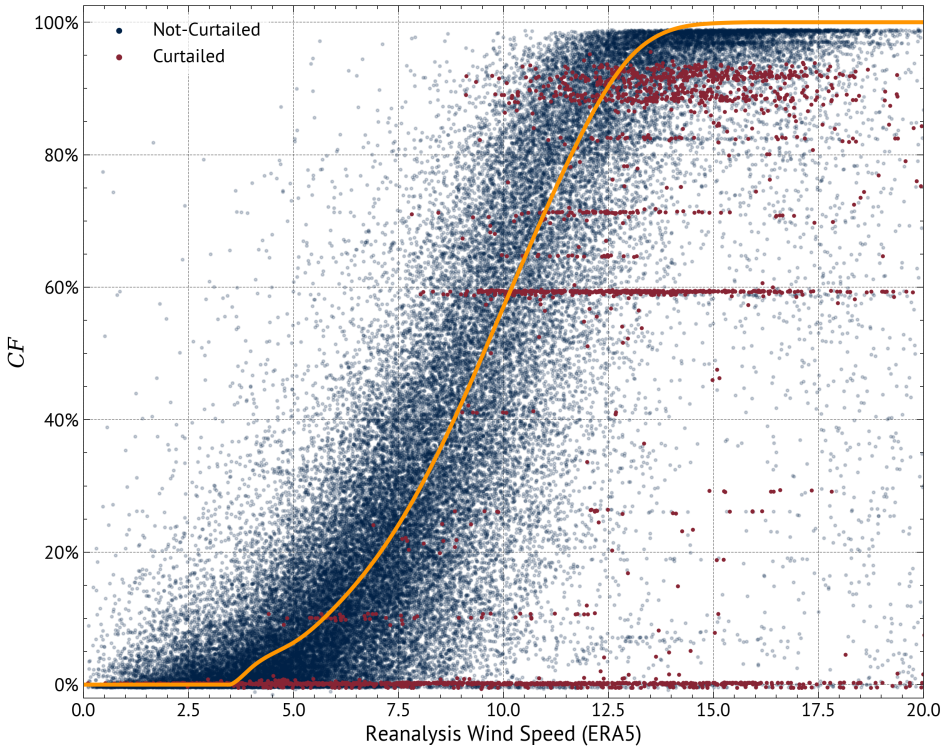


Figure 5.14: Wind speed against capacity factor with curtailment events shown in red, for the Musselroe Wind Farm. Note a CF of 100% is not achieved because the data represents exported power (to the electricity grid), not mechanical power.

5.5.7 Removing outliers

For each regression analysis completed in this thesis, outlier data points were removed. For example, if a regression of average CF over time was being completed for a cohort of wind farms and there was an instance of a wind farm with an especially large increase or decrease in performance, then that data entry was removed from the dataset. Such performance changes could result from factors such as maintenance or curtailment (or data quality issues) and are therefore not relevant to this study, which focuses on long term operational performance decline. It is also important that outlier results do not skew the population mean when averaging. A widely used method discussed in Leys et al. [144] is using the Median Absolute Deviation (MAD) to remove outliers. This method was implemented for

each instance of aggregation, averaging or clustering in this thesis. This method uses absolute deviation around the median to remove outliers.

The MAD for a sample set of n observations is defined as follows [145]:

$$MAD = b \cdot M_i(|x_i - M_j(x_j)|) \tag{5.4}$$

where the x_j is the n original observations and M_i is the median of the series. Usually, $b = 1.4826$, a constant linked to the assumption of normality of the data, disregarding the abnormality induced by outliers [146].

In every calculation used in this study, less than 1% of data was removed. Explicitly, in this study all observations that were outside the bounds of 3 MAD of the median were considered outliers and discarded.

5.6 Conclusion

This chapter has demonstrated through the use of recorded weather station wind speed data that the wind resource is highly variable and the best fit wind distribution at a given site will change with average wind speed and time averaging period. A comparison between best fit distributions and the Weibull distribution has also been drawn to show similarities across different average wind speeds. The lack of site-specific observed wind speed data, combined with the similarity between the best fit and Weibull minimum fit, has led to the conclusion that Weibull is sufficient in most cases. Reanalysis modelling has been introduced, and previous studies that have used reanalysis wind speed data to characterise wind farm performance decline have been reviewed.

We do not inherit the earth from our ancestors; we borrow it from our children.

— Oscar Wilde

6

Performance decline of wind farms in the UK, US and Australia

Contents

6.1 Introduction 112
6.2 Previous long term correction work 114
6.3 Raw Data Trends 118
6.4 Long-Term Correction 125
6.4.1 Existing LTC methods 127
6.4.2 Proposed LTC method 129
6.4.3 Testing the LTC assumptions 132
6.5 Importance of Wind Farm Size 137
6.6 Results 139
6.6.1 Reanalysis wind speed data 139
6.6.2 Wind speed trends 141
6.6.3 Comparing estimated and observed power output 144
6.6.4 Long-term correction and performance decline 147
6.6.5 Capacity-weighted results 160
6.7 Discussion 164
6.7.1 Application 167
6.8 Conclusions 168

This chapter quantifies the performance decline of wind farms in the UK, US and Australia. This is completed by first exploring the raw data trends before performing a long-term correction of the data to establish degradation rates whilst accounting for change of the wind resource over the same period.

6.1 Introduction

All mechanical systems degrade over time. General 'wear and tear' can result from environmental and operational damage, among other unavoidable factors. Understanding how each machine deteriorates and quantifying the rate at which this degradation occurs can help inform investment and operational decisions. Moreover, discovering why a system is degrading can lead to solutions to lessen the rate of decline. This concept applies to wind turbines and wind farms. Previous work has considered the reliability of individual components [147] and how aging impacts downtime [18, 19], however, power production decline rates over time have scarcely been quantified, partly because they are difficult to measure. Nevertheless, understanding performance decline and lifetime performance of a wind farm is crucial for estimating the LCOE, and improving lifetime performance of that asset. Analysing performance decline for all wind farms in a country or connected to the same electricity grid also allows for more accurate prediction of long term grid requirements.

Wind turbine performance declines over time because practically all mechanical systems are subject to deterioration as they age [98]. Machinery prognosis is the forecast of the remaining operational life, future condition, or probability of reliable operation of equipment based off acquired condition monitoring data [148], and it is commonly used for rotating machinery. However, while there have been some advances in the diagnostics of wind turbine faults, little research has been completed on the fault prognostics of wind turbines overall [99]. Only individual turbine components deemed critical, such as the gearbox, generator, main bearing, blades and tower have been thoroughly researched [100]. Data from European wind turbines has been analysed to understand a typical distribution of failure for an operational wind turbine and to characterise failure rate and downtime due to individual components. Figure 6.1 presents a compilation of

two databases[1] showing the distribution of failure rate and causes of wind turbine downtime, reproduced from [149].

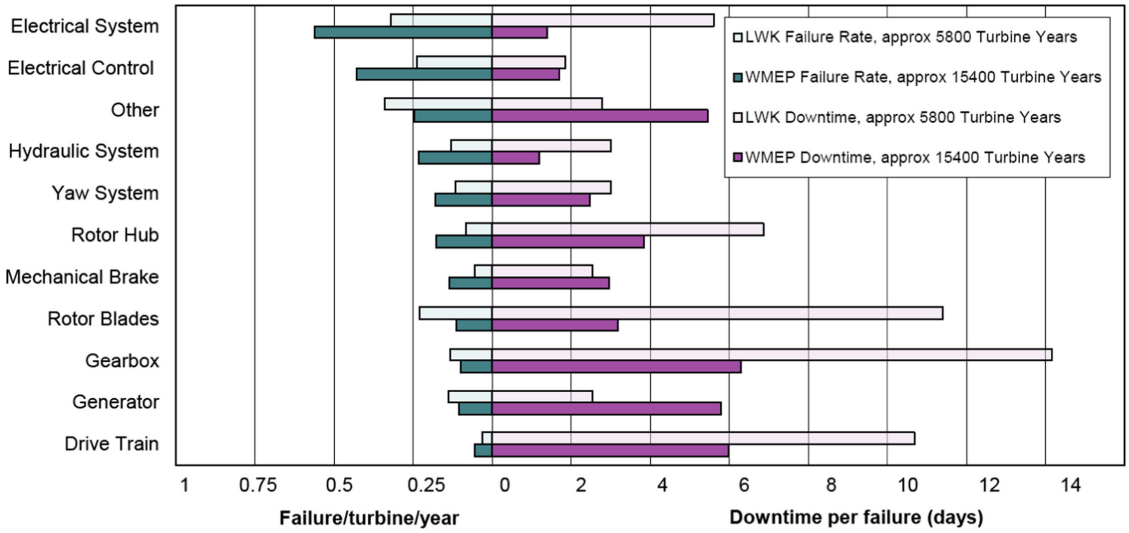


Figure 6.1: Failure rate and downtime of onshore European wind turbines [149].

While the failure rates of, and downtime due to, individual turbine components may be well researched, they are difficult to define for the entire wind turbine system, and therefore difficult to define for the entire wind farm as a whole. This is why some studies have focused on the aggregated degradation of the entire system to characterise a lifetime performance degradation rate. This is a distinctly different approach and presents its own challenges; wind speeds are constantly changing so a degraded turbine on a high average wind speed year might produce more power despite being less efficient, or vice versa. This is important because consideration of power without consideration of the power available from the resource masks the underlying degradation rate of the wind farm, which may result in incorrect analysis or assumptions about performance. This also complicates the analysis by making it difficult to know where the long-term average power output 'should' be if there is no wind speed variability.

[1]The *Wissenschaftliches Mess- und Evaluierungsprogramm* (WMEP) database contains data 1989 to 2006 and contains failure statistics from 1500 wind turbines, and failure statistics published by *Landwirtschaftskammer Schleswig-Holstein* (LWK) from 1993 to 2006 containing failure data from more than 650 wind turbines.

However, there has been a lack of work adopting this approach and focused on quantifying wind farm level performance decline, with a few notable exceptions. Staffell and Green [18] investigated performance decline across wind farms in the UK, Olauson et al. [19] looked at degradation of wind farms in Sweden, Hamilton et al. [21] quantified performance loss of wind farms in the US, and Germer and Kleidon [20] investigated German wind farms as they age. Quantifying performance loss can significantly impact business cases for wind farms, and it is crucial to develop a robust measurement framework that works for all available wind farm output data.

This chapter builds on previous methods used to define wind farm performance decline accounting for long term variation in wind speed and provides the first performance decline analysis across multiple countries – the UK, US and Australia. A framework is provided to allow this analysis to be replicated for different countries where wind farm power output data is available.

6.2 Previous long term correction work

All of the previous long term correction (LTC) studies highlighted the impact of variable wind speeds on wind farm performance. However, the majority of the previous studies used data that were averaged to a monthly level, and long-term correction predominantly focused on removing the effects of seasonal variations in wind speed. However, for higher time resolution data, there may exist higher frequency wind speed trends that also need to be considered. If this is the case, existing LTC methods will need to be adjusted. This chapter uses the analysis of wind speed distributions over different time periods shown in Chapter 5 to better understand the needs of a robust LTC analysis. Furthermore, not all wind farm power output data will have a time resolution of one month (e.g. Olauson et al. [19] use hourly data from Swedish turbines, and this study used 5-minute data from Australian turbines). Currently, there is no LTC method that robustly works across different time resolutions of wind farm power output data.

Hughes was the first to discuss performance decline using wind farm capacity factors in a paper for the Renewable Energy Foundation (REF) in 2012 [134].

Hughes analysed over 10 years of operating data (2002 to 2012) for UK and Danish wind farms and used an econometric analysis to calculate performance decline. The report suggested the normalised[2] capacity factors of UK wind farms have declined from a peak of 24% at age 1, to 15% at age 10, and 11% at age 15 (a rate of 0.93 CF percentage points per year), and equates to a 3.62% decline per year compared to the mean non-normalised CF for all UK wind farms across the 10 year period examined, 25.6%. Danish wind farms reduced from a peak of 22% at age 1 to 18% at age 15 (0.29 CF percentage points per year), a rate of 1.28% from the average non-normalised CF, 22.24%. The small set of Danish offshore wind turbines in the study reduced from a peak of 39% at age 0 to 15% at age 10 (2.4 CF percentage points per year), a rate of 6.54% per year compared to the non-normalised averaged CF of 36.7%. Hughes normalised age-performance curves by using a regression which corrected for month-on-month variations in national wind speed, each wind farm's location and age. Wind availability at each wind farm was treated as an unobserved factor whose contribution to monthly variations in wind output was estimated by a series of fixed period (month or year) effects in a statistical model of monthly output for all wind farms [134]. The lack of known local wind speeds is a crucial limitation of Hughes' work; however, Hughes shows that using fixed period effects determined by the regression is econometrically superior to using national wind speed averages because site-specific conditions vary compared to national averages. The Hughes study is significant because it was the first attempt to estimate the performance decline of wind farms on a national scale despite validated wind turbine field data being difficult to obtain [150].

Staffell and Green [18] extended Hughes' work by combining wind farm output data with reanalysis wind speed data and wind turbine power curve data. This allowed them to incorporate high temporal and spatial wind speed data into their analysis, which was a key limitation of Hughes' work. Their wind speed data was produced from the MERRA [151] dataset. By doing so, they were able to compare simulated wind farm power with observed wind farm output and use this

[2]Hughes defines normalised as the total output over a time period divided by the maximum potential output adjusted for wind availability in that same period

to perform a long-term correction of power data to produce a more robust estimate for performance decline. Simulated outputs were compared with observed monthly capacity factors at 282 UK wind farms, yielding the normalised performance for each wind farm. Their method accounted for local wind resource availability and produced weather-corrected capacity factors from which the effects of aging can be deduced. Their study was restricted to UK wind farms and concluded that wind farms lose $1.6 \pm 0.2\%$ of their output per year, with average capacity factors decreasing from 28.5% when new to 21% at age 19. Staffell and Green's examination of UK wind farm performance has since become the standard for subsequent studies of this type.

Olauson et al. [19] produced a performance decline analysis of Swedish wind farms, which was heavily based on Staffell and Green's paper. The primary objective of their paper was to present a similar study for Swedish conditions, albeit implementing slightly different methods. Olauson et al. included the impact of downtime as turbines age and factors explaining differences in trends into their regression. Including the downtime data was possible due to access to a unique database of hourly observations for almost all Swedish wind turbines. In total, their study comprised 1,100 monthly and 1,300 hourly time series datasets that partially overlapped. Multiple reanalysis models were also used (MERRA [151], ERA-Interim [116] and EMD ConWx [152]) to ensure the results were robust across models. An updated long-term correction method – where wind speed was first linearly scaled to the long term average before predicting "fictive" turbine power – was used as well as a seasonal regression to better capture the long term trends; however, generic (unspecified manufacturer) wind farm power curves were used which limits the accuracy at each wind farm location. Olauson et al. corrected data based on the average fictive data of a particular month divided by the average of the entire period, and thus did not incorporate hourly trends into their analysis. Their analysis showed that Swedish wind turbines constructed before 2007 lose 0.15 capacity factor percentage points per year, corresponding to a lifetime energy loss of 6% [19].

Germer and Kleidon [20] used a linear mixed-effects model to evaluate different mechanisms that lower turbine yields in observed wind turbine generation in

Germany between 2000–2014. They used monthly energy yield data and compared them to predicted energy yields calculated using the COSMO-REA6 [153] reanalysis dataset. The wind turbine database includes the location of German onshore wind turbines since 1988 and the monthly energy yields for a subset of individual wind turbines. However, the database only provides locations in terms of postal code and does not report the exact positions of wind turbines or wind farms, which provides uncertainty around the accuracy of the reanalysis wind speed and other metrics, limiting this analysis. The average size of a German post code is not provided, however Germer and Kleidon state that "postal code areas differ considerably in size" [20]. If a wind turbine is located in the middle of the farm it will be heavily influenced by inter-turbine effects such as wake effects, and local wind speeds can vary within a given postal code area. Germer and Kleidon also used installed capacity per postal code area as a variable, but they noted that this might not be the best proxy for regional installed capacity as the postal code areas differ considerably in size. The final combined dataset included 531 single turbines and 921 wind farms. Due to data limitations, the 'ideal turbine yield' was calculated using a generic power curve instead of a site-specific one. This results in uncertainty around the 'loss' values found in this analysis, as the CF-wind speed curves can vary significantly between different wind turbines [154]. Germer and Kleidon concluded that turbine performance decreased by $0.63 \pm 0.01\%$ per year. A key difference in this study is that generation data were not long-term corrected to produce a performance trend and were instead used to compare actual generation to 'ideal' generation (power generation produced using reanalysis wind speed data and a power curve), which resulted in an average of 26.3% overall losses.

Hamilton et al. [21] used similar methods to both Staffell and Green [18] and Olauson et al. [19] to investigate the performance trends at US wind farms. They explored the influence of plant characteristics on the rate of performance decline through a multivariate analysis. However, they also included the effect of policy mechanisms and curtailment into their analysis. They found a significant drop in performance of 3.6% at the 10 year mark, as plants lose eligibility for the production

tax credit and operational strategies shift, which is unique to the US. Overall, they found a degradation rate of 0.53%/year for older vintage plants (>10 years old) and 0.17%/year for newer plants (<10 years old). They were also able to remove the effects of market-based curtailment (before completing the LTC analysis) of wind farms based on data recorded in Texas, by producing new estimates informed by market price signals. Additionally, they used the same method as Staffell and Green [18] to long-term correct the generation data, where the wind speed data was sourced from the ERA5 reanalysis model.

A summary of the key elements of the previous research is in Table 6.1.

Table 6.1: Summary of the key statistics from the relevant previous studies on country-wide wind farm performance decline. Note the Olauson et al. performance decline rate is based off the assumption that the average CF is equal to 30%.

Paper	**Country**	**Reanalysis model(s)**	**Power curves**	**Degradation rate (%/year)**
Hughes (2012) [134]	UK	–	–	3.63
	Denmark	–	–	1.28 (onshore), 6.54 (offshore)
Staffell and Green (2014) [18]	UK	MERRA [151]	Site-specific	1.6 ± 0.2
Olauson et al. (2017) [19]	Sweden	ERA-Interim [116], EMD ConWx [152]	Generic	0.5
Germer and Kleidon (2019) [20]	Germany	COSMO-REA6 [153]	Generic	0.63 ± 0.01
Hamilton et al. (2020) [21]	US	ERA5 [120]	Site-specific	0.53 (age > 10), 0.17 (age < 10)

6.3 Raw Data Trends

This section presents trends from the raw wind farm power output data, with the purpose of understanding wind farm electricity generation trends. These trends are different from the individual wind farm performance trends because the data has

not been normalised relative to the wind speeds occuring (or hindcast) at the wind farm site. However, these are the trends that have been observed by grid operators and therefore still have a tangible meaning. The results from this section will be compared to the long-term corrected results presented in Section 6.6.

Figure 6.2 shows the distribution of all observed capacity factors against age, separated by country, and fit with a linear regression. The yearly means are produced by aggregating the observed CF (CF_{obs}) from each wind farm to annual values (based on age, not date), and then calculating the mean wind farm CF_{obs} (not weighted by installed capacity). The wind farm capacity factor degradation rate was calculated using Equation 6.1:

$$CF = \beta_0 + \beta_1 \cdot age + \epsilon \tag{6.1}$$

The regression was fit for ages > 1 because the first year of operation often included periods where wind farms are operating at less than available capacity whilst still getting grid-connected, which would skew the results. This regression results in a decline of 0.23, 0.64 and 0.36 absolute percentage points per year for the UK, US and AUS respectively. Combining all three countries results in a performance decline of 0.49 CF percentage points per year. Higher order polynomial regression fits were considered, especially in the case of the US data; however, the shape of the resulting distributions across countries, combined with the comparison to previous results, resulted in a preference for linear regression. The linear regression reasonably describes the degradation rates seen here, but also implies at a certain age CF will reach 0 and be negative thereafter (albeit at an age where the wind farms would already be decommissioned), which is unrealistic. It may also be expected that a wind farm will decline at a fixed percentage rate per year (as presented in previous studies [18–20]), which would therefore mean the trend is exponential, however this is not shown in the raw data. The raw CF data indicates the trend is approximately linear.

There exists a clear negative trend for each country, showing that power production from wind farms in these countries has decreased as wind farms age.

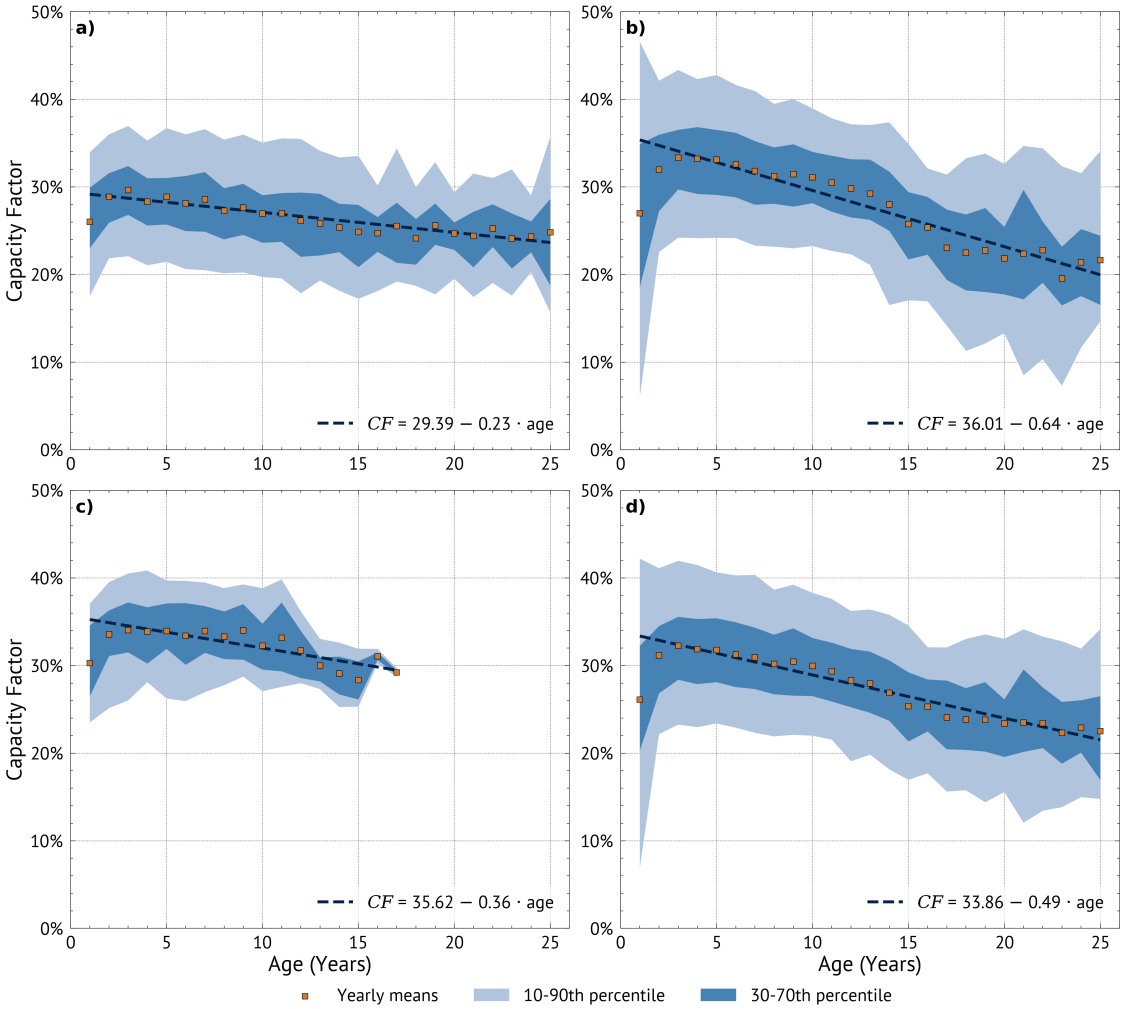


Figure 6.2: Yearly averaged observed capacity factors against age for wind farms in (a) UK, (b) US, (c) AUS and (d) all three countries combined. The shaded regions indicate yearly percentile ranges of CF_{obs} and the yearly means are shown with squares. Each country has a linear regression line computed from age 1 onwards.

Given that these plots are against age and not time, it is worth noting that any trends in wind speed over each wind farm's lifetime are included in these trends. That is, for two adjacent farms that have commenced operation in different years, the observed wind speeds in a specific year will appear at different ages in Figure 6.2. Nevertheless, in this case, each wind farm could still be materially biased by region trends at different ages. The US has approximately double the performance decline rate as the other two countries and has an increased decline rate between ages 14 and 17, although it is not sustained. Weather conditions may exacerbate these trends; hence the need for verification with a long-term weather correction model.

Grouping all turbines together can also produce an effect that impacts these results. In the dataset used in this study, there are inevitably a larger number of data points for the earlier ages because the number of wind farms being commissioned continues to increase and the oldest generation of commercial wind farms are only now starting to be retired. Furthermore, the newest turbines in this dataset that have been operating for 5–10 years will typically be more efficient than their older counterparts at extracting energy from the wind due to their larger size, and technological, operational and design improvements, which would push their CF higher on average. Figure 6.3 shows the average CF for the UK, US, AUS and ALL wind farms grouped by their commission year, as well as indicating the overall capacity installed that year and the average length of the time series data in years. This indicates that CF has been steadily rising over time, and while this may not affect the overall performance decline rate of newer wind farms, it will skew the CFs of farms at earlier ages in Figure 6.2 higher. It is also known from Figure 5.8 that the majority of observed data are aggregated in the ages between 0–10, and the majority of wind farms were commissioned since 2010 in this dataset. Therefore, it is hypothesised that Figure 6.2 will be slightly skewed as the newer wind farms with higher capacity factors will increase the average CF of younger wind farms and exaggerate the regression line.

To break down some of the trends occurring at wind farms in different age ranges, the data can be grouped by their final recorded age and individual wind farm degradation rates can be quantified. Data can also be grouped by average CF, size of wind farm, number of turbines and whether the wind farm is onshore or offshore. Individual wind farm degradation rates are calculated using the observed capacity factor for each wind farm and Equation 6.2:

$$CF = \beta_0 + \beta_1 \cdot age + \beta_2 \cdot \cos(2\pi \cdot age) + \beta_3 \cdot \sin(2\pi \cdot age) + \epsilon \tag{6.2}$$

where coefficients β_2 and β_3 are included to account for residual seasonal patterns in CF and produce a more reliable linear slope. In this thesis, β_2 and β_3 is measured for the grouping of turbines in focus (e.g., single turbine or cohort/grouping of

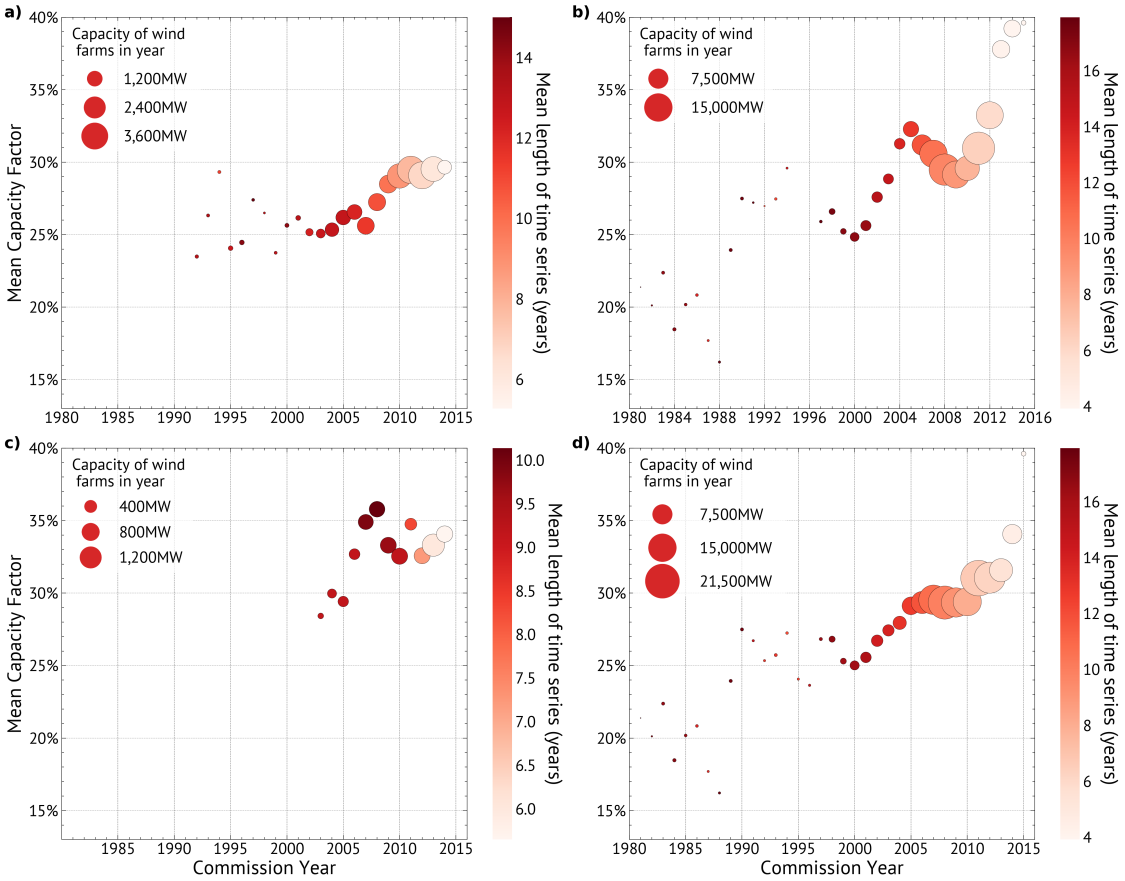


Figure 6.3: Mean lifetime CF for wind farms in the (a) UK, (b) US, (c) AUS and (d) ALL grouped by their commission year. The size of the markers represents the total capacity commissioned in the year, and the colour represents the average length of the data series.

turbines). Age is used instead of time for these regressions, such that regression results for all wind farms can be aggregated. This equation is used because the full datasets (monthly and hourly data) exhibit seasonal behaviour. Figure 6.4 illustrates the importance of including the seasonal components in the regression to accurately capture the linear trend.

Figure 6.5 summarises the annual degradation rates for individual wind farms for each country, grouped by age groups. The degradation rates vary for each country, but are generally bunched just below 0%. It is expected that the degradation rates will become increasingly negative for the older wind farms (darker shading), however this effect is hard to visualise due to the larger spread and larger sample sizes for the younger wind farms. Some wind farms show a positive degradation

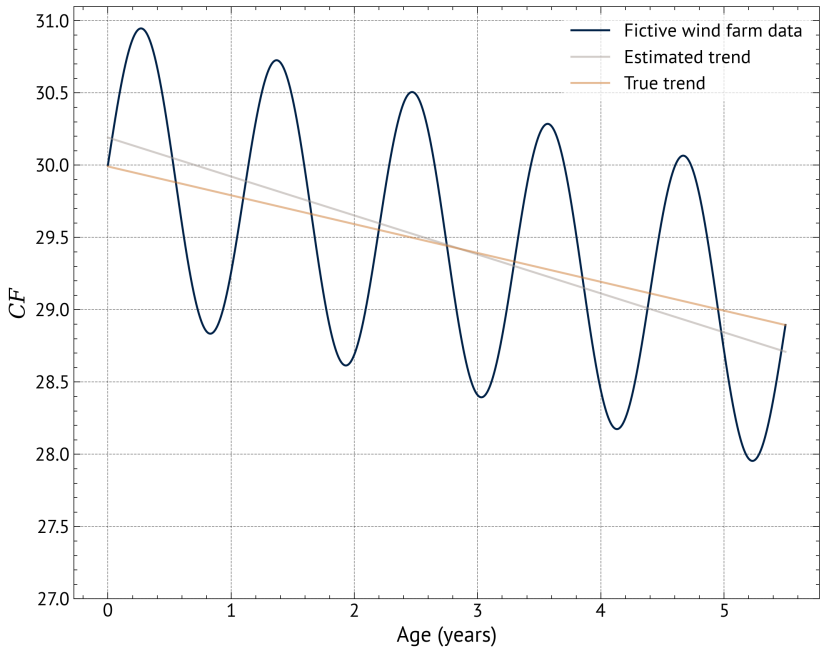


Figure 6.4: A fictive example of CF against age, demonstrating the importance of including seasonal components to the regression equation to accurately capture the linear trend.

rate (performance improvement) which can occur for younger wind farms which are yet to degrade or for reasons such as poor early age performance followed by improved operation. Several outliers exist in both directions which stretch the distribution out. A Student's t-distribution is therefore used to fit the entire distribution, centred on -0.207, -0.158, 0.104 and -0.168 for UK, US, AUS and ALL respectively. The standard deviations for these distributions are 0.728, 1.365, 0.515 and 1.094 respectively. The fit for AUS is largely irrelevant because of the small sample size, especially with the lack of wind farms older than 15 years. The mean of these distributions show that on average (using observed CF data), wind farms are degrading over time; however, the standard deviation of these distributions show there are a large range of performance trends across the entire dataset.

The effect of improved technology, operation and design can be further tested by looking at the average degradation rates (calculated using Equation 6.2) for wind farms grouped by commission year. Figure 6.6 shows this and also highlights the capacity of wind farms installed per year (size of markers) and the average length of the time series (colour of markers). More recently commissioned wind farms have

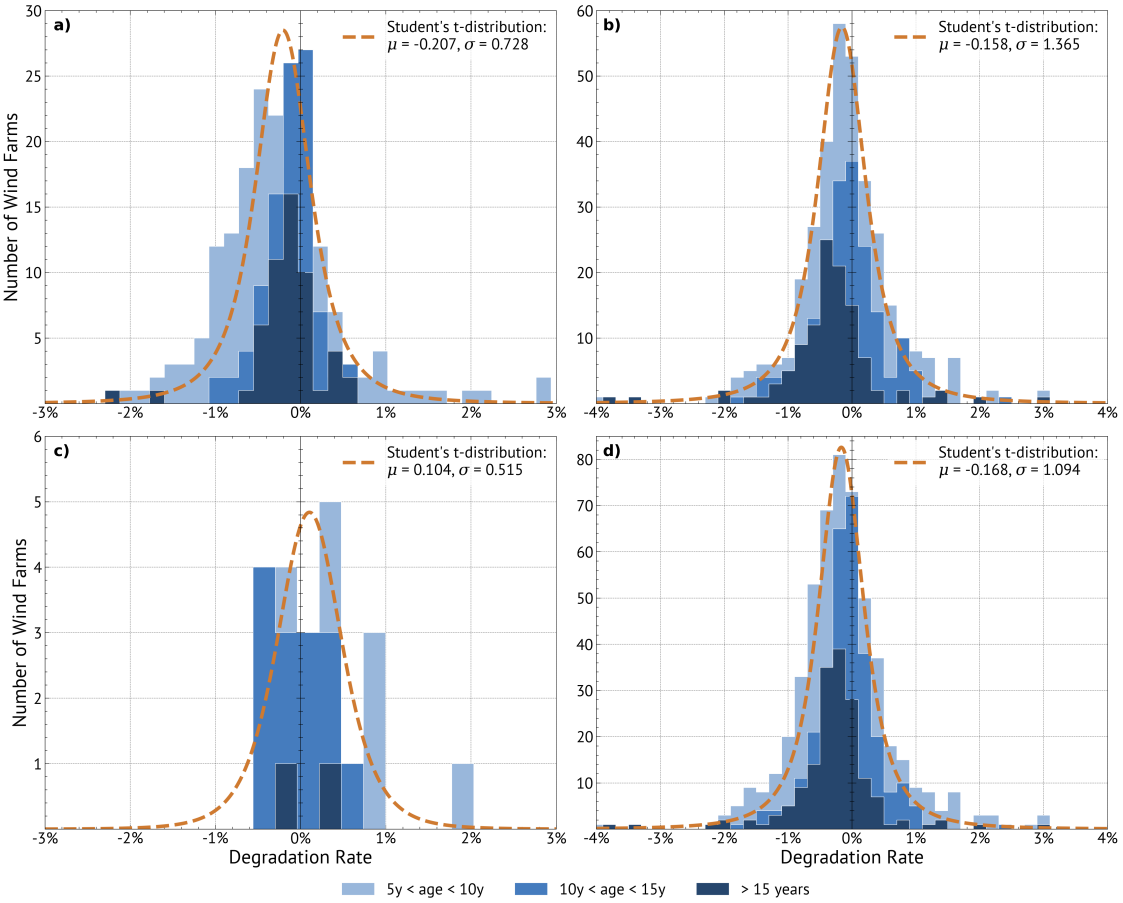


Figure 6.5: Distribution of lifetime degradation rates for individual wind farms for (a) UK, (b) US, (c) AUS and (d) ALL.

slightly less degradation or even an increase in performance over time (degradation rate > 0). The increase in mean degradation rate over time supports the hypothesis that the raw data degradation rates may be inflated due to the superior performance of the more recently commissioned wind farms. The trends in Figures 6.3 and 6.6 support the need to correct the data in order to test the results from Figure 6.2.

The trends presented here use observed wind farm power data, but they have not been corrected. To produce robust results, the data must be corrected to remove wind speed variability. This is abundantly clear when using the same method as Figure 6.2 but with monthly averages instead of yearly averages in Figure 6.7. Even when aggregating over all wind farms, seasonality trends persist, revealing the impact of varying weather conditions. Seasonality and other potential sources of variability create doubts about the underlying trend and must therefore

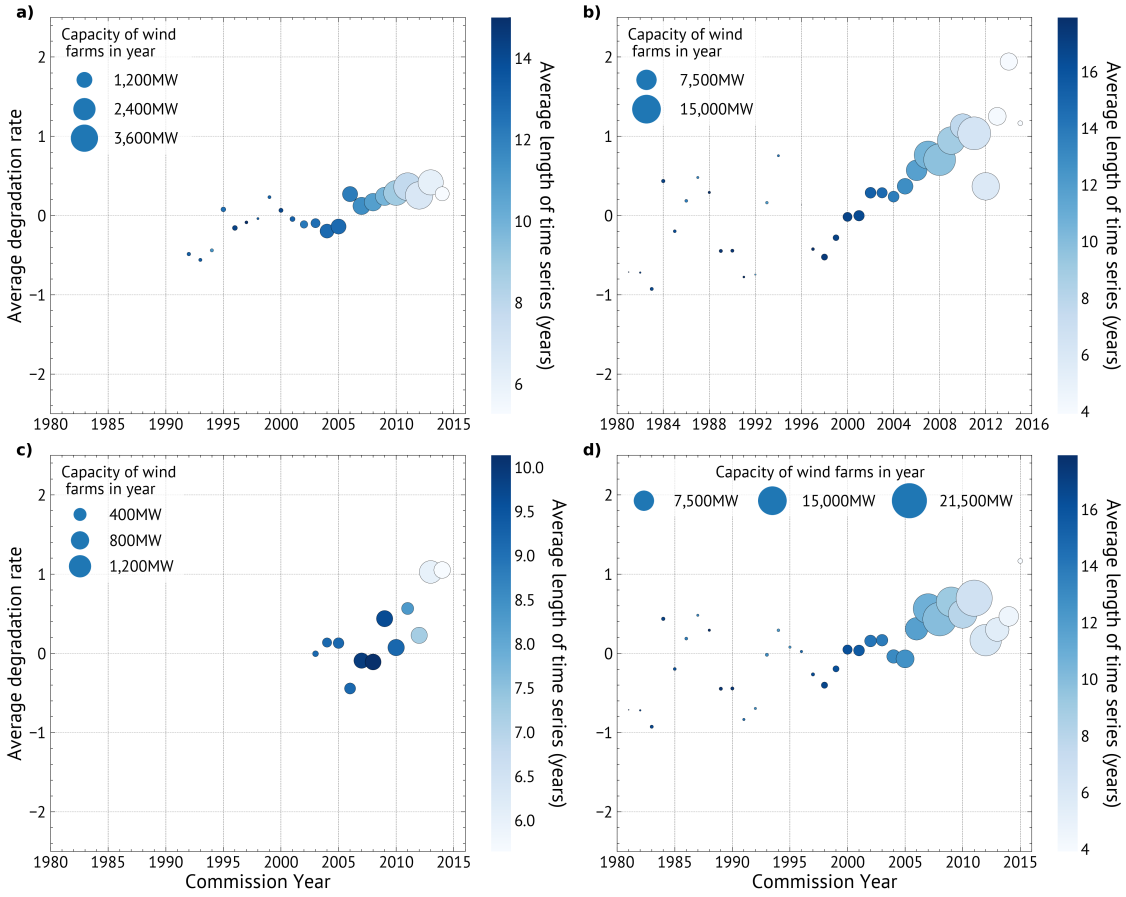


Figure 6.6: Mean degradation rate (regression slope from Equation 6.2) for (a) UK, (b) US, (c) AUS and (a) ALL wind farms grouped by their commission year. The size of the markers represents the total capacity commissioned in that year, and the colour represents the average length of the data series.

be methodically accounted for. The long-term correction method for this study is described in Section 6.4.

6.4 Long-Term Correction

To robustly analyse the change of performance across UK, US and AUS wind farms, considerations must be made for potential sources of variability in the data. These sources of variability include: trends in the reanalysis wind speed data, inaccuracies in reanalysis wind speed data, turbine downtime, curtailment and electrical losses. The wind speed variability and inaccuracies in the reanalysis wind speed data have been addressed using a novel LTC method.

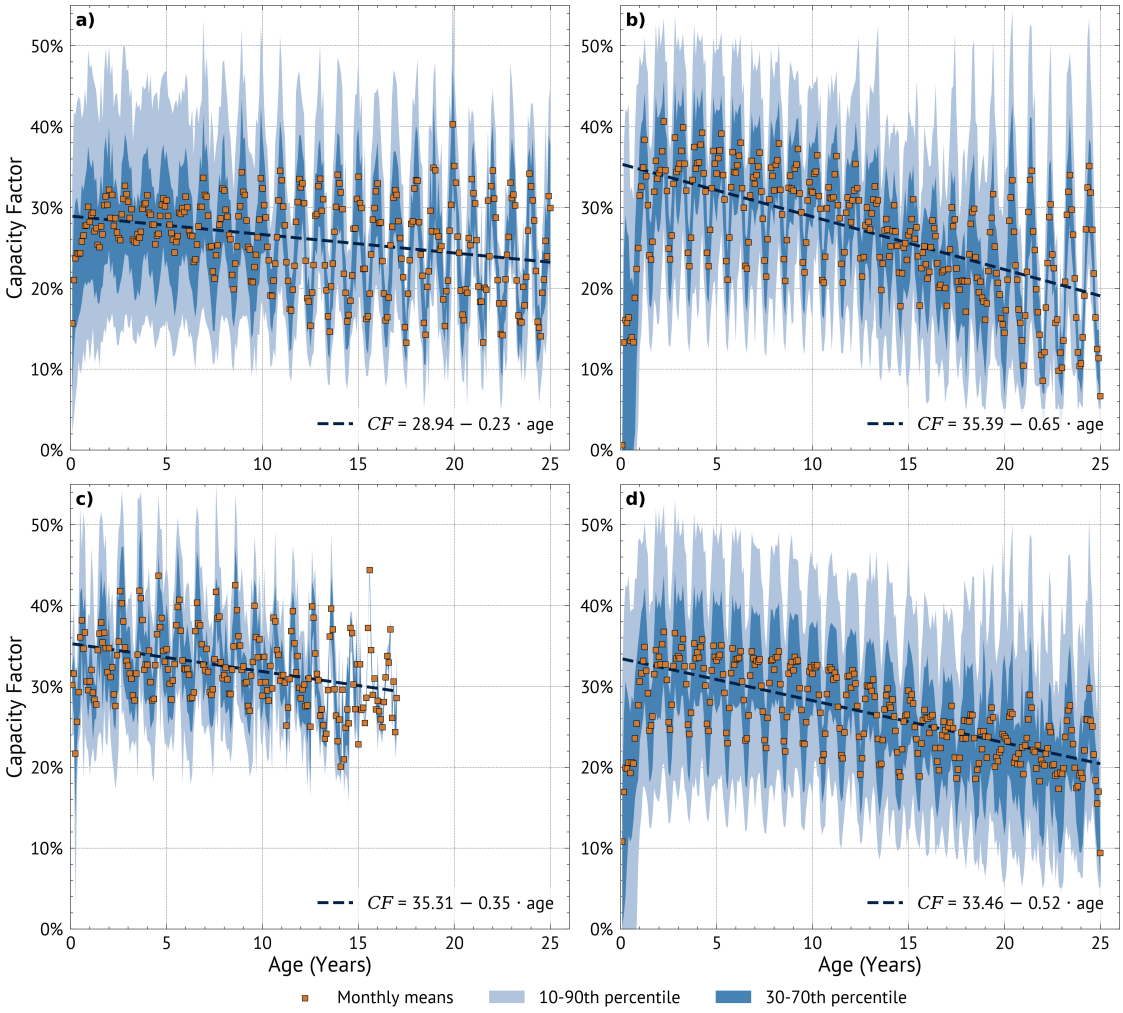


Figure 6.7: Monthly averaged observed capacity factors against age for (a) UK, (b) US, (c) AUS and (d) all three countries combined. The shaded areas indicate percentile ranges and the monthly means are showed. Each country has a regression line computed from age 1 onwards.

The purpose of LTC of wind power time series data is to separate out the trends in power production from underlying causes of variability. It is important to recognise that this LTC is trying to remove wind speed trends from the data, such as seasonality, rather than predict the power output of a wind farm at a given time. This ensures that performance decline findings are not influenced by any underlying changes of the wind speed data.

6.4.1 Existing LTC methods

Long term correction methods compare observed wind farm performance (CF_{obs}), which is affected by variations in wind speed and performance degradation, with estimated wind farm performance (CF_{est}) based on wind speed data and (idealised) wind farm vs wind speed performance curves. A typical LTC process consists of the following steps:

1. Calculate a long-term average benchmark metric (using either reanalysis wind speed, or estimated CF using reanalysis wind speed)

2. For every data point, calculate the CF_{est} for the wind farm based on reanalysis wind speed and using the wind farm power curve

3. If the CF_{est} is higher than the benchmark, adjust the CF_{obs} down by the relative amount higher than the benchmark. If the CF_{est} is lower than the benchmark, adjust the CF_{obs} higher.

The process is simple; however, the difficulties lie in understanding whether the reanalysis data are accurately representing the real wind speeds in a particular time period, calculating the CF_{est} accurately, and adjusting the CF_{obs} by the correct magnitude. The purpose is to estimate a long-term corrected CF (CF_{LTC}) time series that sits close to the underlying performance trend of the wind farm, without knowing what the trend line is.

Typical LTC methods involve the use of a "wind index" (WI) to adjust recorded data (step 3 above). A wind index provides an indication of the mean wind speed during a specific historical period, relative to that of a designated long-term period [155]. A WI can be computed from different sources, such as surface wind speeds [156], wind farm generation data [157] or reanalysis wind speeds [158, 159]. The wind index can be calculated using the following equation [160]:

$$WI = \frac{X_t}{\sum_{t=1}^{N} X_t \Big/ N} \tag{6.3}$$

where X is the wind speed data (i.e. recorded measurements, wind farm power output, etc.), t is a time series point, and N is the total number of recordings.

Once calculated, the wind index can be used to 'correct' recorded data by shifting a recorded power output towards the long-term mean. This is completed by dividing the recorded power output by the wind index value for that time period. To correct power output data, an "estimated" power needs to be computed using wind speed data and an appropriate power curve because power output is not linear with wind speed (meaning that wind speed cannot be used directly to correct data). In effect, $X = CF_{est}$ in Equation 6.3 where CF_{est} is the estimated capacity factor.

To calculate CF_{est}, reanalysis wind speed data, scaled to the wind farm's hub height (see Section 5.5.4) are used along with the power curve for that farm. In previous studies, the CF has been estimated using the reanalysis data as a 'point' wind speed. However, a point wind speed does not represent the variability of wind speed in the time period of the data (either one hour or one month in this study), and is therefore neglecting the wind speed variability trends shown in Section 5.1.

Once CF_{est} is calculated, the LTC calculation is performed using the following calculation:

$$CF_{LTC,t} = \frac{CF_{obs,t}}{CF_{est,t} \Big/ \left(\sum_{t=1}^{N} CF_{est,t} \Big/ N \right)} = \frac{CF_{obs,t}}{WI_t} \tag{6.4}$$

Previous methods either fail to implement or neglect to incorporate several aspects that may be expected to affect the calculated trends: previous methods only show plausible results using monthly data, which therefore neglects areas with more granular datasets; power production is typically completed using a 'point' wind speed, which does not reflect the time period it represents; and inaccuracies in reanalysis models are not factored into the analysis. The importance of these issues, and potential solutions, are outlined below.

6.4.2 Proposed LTC method

In this work, a novel framework is proposed for long-term correction of time series power production data. An overview of the 'Kelly method' LTC process is shown in Figure 6.8.

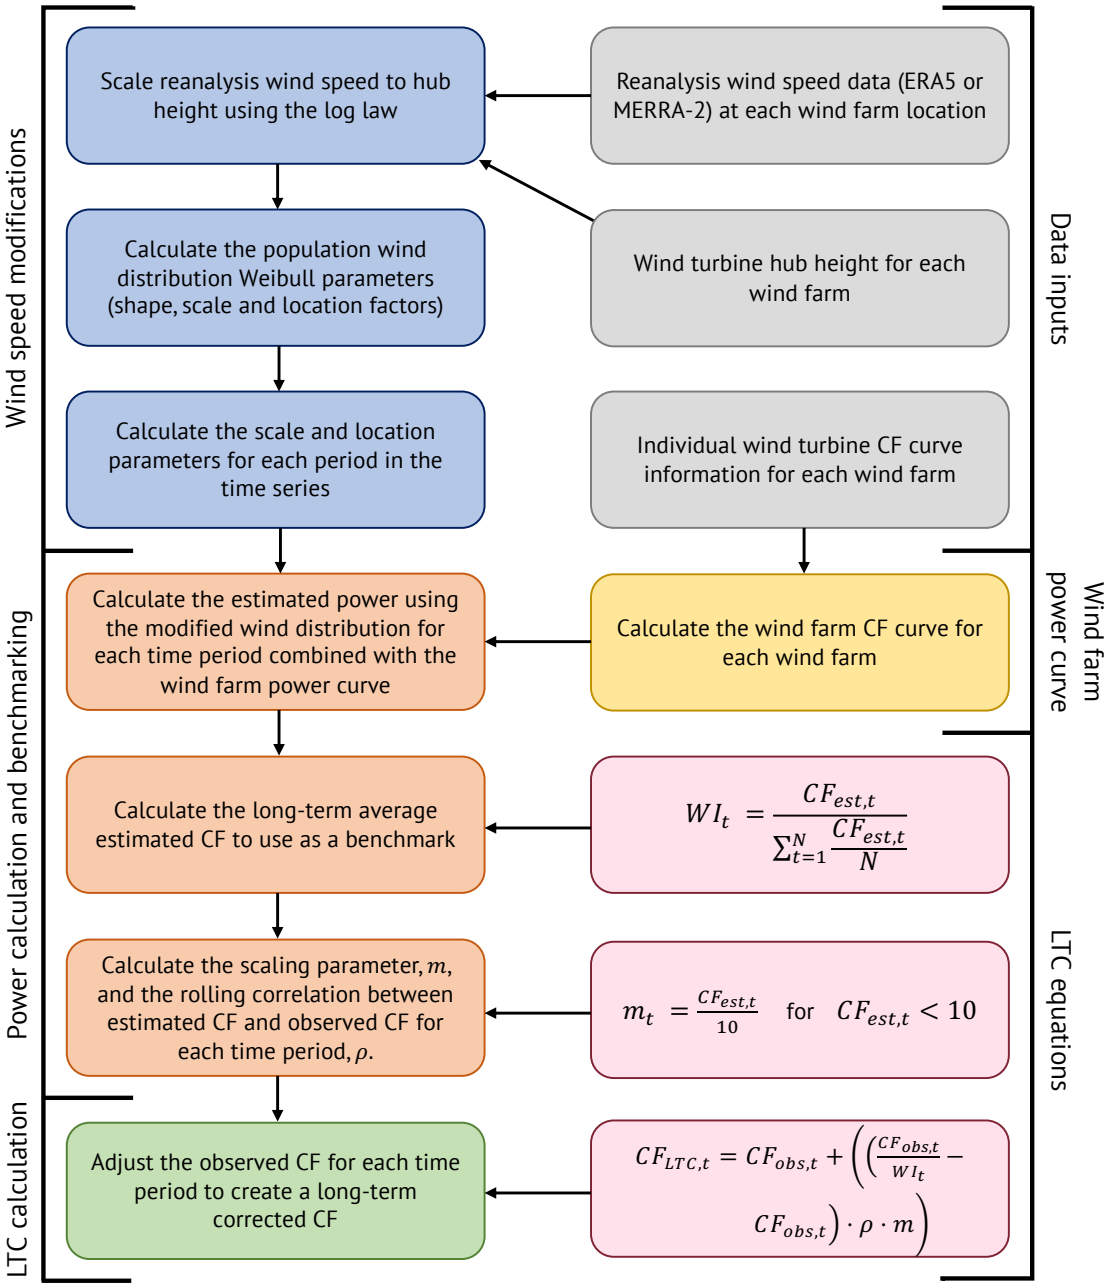


Figure 6.8: LTC process flow chart.

This novel framework builds on existing methods [18, 19], with the following

areas targeted for improvement:

- Time resolution of data: creating a method applicable to all resolutions of time series data (i.e. hourly, daily, monthly)
- Realistic power production: accounting for local wind distribution characteristics
- Reanalysis errors: accounting for periods where reanalysis and actual data (generation data) are diverging in their trends

Thus, a robust approach has been developed that can be used to compare performance decline across different countries and different datasets.

The proposed method recognises that the reanalysis wind speed represents a period of time (either an hour or a month for this study), and so a wind speed distribution around the reanalysis wind speed value is combined with the power curve to calculate CF_{est}. This is especially important in regions near the cut-in or cut-out speeds, where a point distribution is an 'all or nothing' approach to whether the wind farm generated power in that time period. The wind speed distribution for each time period depends on both the reanalysis wind speed for that time period and the specific wind farm as detailed in Section 5.5.5.

If Equation 6.4 were to be used to calculate CF_{LTC}, it would be a similar method to previous studies, albeit with a novel approach to evaluating power production for each time period. However, there are two fundamental flaws in this model: low values of $CF_{est,t}$ resulting in low WI_t values, and the potential inaccuracies in the reanalysis wind speed estimate.

The problem with a low $CF_{est,t}$ value is that it results in $WI_t \approx 0$, which causes CF_{LTC} to equal an unrealistically high value (often with $CF_{LTC} > 100$). This is especially important when using hourly data, as the range of reanalysis wind speeds, and therefore CF_{est}, is much larger than when using monthly-averaged data. To limit the impact of low WI_t values, a scaling variable, m, is used for $CF_{est,t} < 10$, where m is defined as:

$$m_t = \frac{CF_{est,t}}{10} \quad \text{for} \quad CF_{est,t} < 10 \tag{6.5}$$

and $m = 1$ for all other instances. m has the effect of mitigating the scaling effect of the wind index for low $CF_{est,t}$ values. As $WI_t \to 0$, then $m \to 0$ and the original $CF_{obs,t}$ remains unchanged. m has a very small impact on the LTC calculation for $CF > 5\%$ and no impact on the LTC calculation for $CF > 10\%$. Any CF_{LTC} values post-scaling that were still greater than 100 were changed to $CF_{LTC} = 100$.

Given that reanalysis data are estimates of historical wind speeds, using $CF_{est,t}$ to scale power is not always accurate as the reanalysis and historical wind speeds sometimes diverge. Therefore, while the purpose of LTC methods is to correct for weather trends, the only confirmed data in the LTC analysis are the observed wind farm generation data. This is important because all other calculated data are estimates based on many assumptions and should be recognised as such. Consequently, an algorithmic approach is adopted to account for anomalous results. In cases that were more uncertain (i.e. CF_{est} and CF_{obs} diverging significantly), corrections were made less severe or removed, thereby defaulting to the underlying generation data.

A rolling Pearson correlation [161] (ρ_t) between $CF_{obs,t}$ and $CF_{est,t}$ is calculated over 3 months for monthly data (UK, US) and 7 days for hourly data (AUS) and used as a measure of reliability in the reanalysis model. A 3 month time period is chosen because monthly-averaged data are less volatile than hourly data, and thus fewer data points are needed. For hourly data, a week was chosen because it corresponds to a typical weather forecasting period and provides a large data sample to compare. The rolling correlation was then used to weight the amount of correction that was applied to CF_{obs} (e.g. if $\rho = 0.8$ and the correction from Equation 6.4 was determined to be $+0.5$ CF percentage points, then the final correction would be $+0.5 \cdot 0.8 = +0.4$ CF percentage points).

A final representation of the LTC method is represented in Equation 6.6, and a full description of the LTC algorithm can be found in Algorithm 1.

$$CF_{LTC,t} = CF_{obs,t} + \left(\left(\frac{CF_{obs,t}}{WI_t} - CF_{obs,t} \right) \cdot \rho_t \cdot m_t \right) \tag{6.6}$$

This method has been designed to remove errors from typical WI/LTC methods and work for multiple data granularity levels. A scaling method for reducing exaggerated scaled CF_{LTC} values was used. Furthermore, the focus has been shifted from relying on the reanalysis data, to acknowledging that the CF_{obs} data provides the only certainty, so that when CF_{est} is poorly correlated with CF_{obs}, then CF_{obs} should be more heavily weighted. By correcting wind farm output data using this novel method, a robust performance decline analysis can be undertaken.

6.4.3 Testing the LTC assumptions

There are two elements of the proposed method that need to be tested to ensure that the method is statistically sound: the data modification (LTC method), and the regression method. The LTC method must not alter the data in a way that will skew the time trend coefficient (i.e. the LTC calculation must have a statistically insignificant correlation with age). Additionally, the regression method (linear) must satisfy the common Gauss-Markov assumptions [162], which indicate the appropriateness of the linear model.

Data modification

To test whether the LTC method alters the underlying trend, Equation 6.6 can be equated to Equation 6.1:

$$CF_{LTC,t} = CF_{obs,t} + \left(\left(\frac{CF_{obs,t}}{WI_t} - CF_{obs,t} \right) \cdot \rho_t \cdot m_t \right) = \beta_0 + \beta_1 \cdot age_t + \epsilon \tag{6.7}$$

Simplifying this results in:

$$CF_{obs,t} + \Big((\Delta CF) \cdot D_t \Big) = \beta_0 + \beta_1 \cdot age_t + \epsilon \tag{6.8}$$

where D_t is the 'dampening correction' $(\rho_t \cdot m_t)$ and ΔCF is the difference between the observed value and previous LTC methods. If the correlation between the

Algorithm 1 Kelly LTC algorithm

function SCALEWINDSPEED($U_{1,w},\ U_{2,w},\ z_1,\ z_2,\ h_w$)
  for $U_{1,w,t}, U_{2,w,t}$ in $U_{1,w}, U_{2,w}$ **do**
    if $U_{2,w,t} \leq U_{1,w,t}$ **then**
      $U_{s,w,t} = m \cdot h_w + c$ ▷ Linear scaling for inversed conditions
    else
      $z_{0,w,t} = \exp((U_{2,w,t} \cdot \log(z_1) - U_{1,w,t} \cdot \log(z_2))/(U_{2,w,t} - U_{1,w,t}))$
      $U_{f,w,t} = (U_{1,w,t} \cdot \kappa)/(\log(z_1/z_{0,w,t}))$
      $U_{s,w,t} = (U_{f,w,t}/\kappa) \cdot \log(h_w/z_{0,w,t})$ ▷ The log law
    array($U_{s,w}$) $\leftarrow U_{s,w,t}$
  return $U_{s,w}$ ▷ Hub-height scaled wind speed

function WEIBULL($U,\ c,\ k$) ▷ Weibull distribution
  return $(k/c) \cdot (U/c)^{k-1} \cdot \exp(-(U/c)^k)$

function ESTIMATEPOWER(U_s, POWERCURVE, c_w, k_w)
  for $U_{s,t}$ in U_s **do** ▷ Loop over all U_s values
    $c_w = U_{s,t}/\Gamma(1 + (1/k_w))$ ▷ where Γ is the gamma function
    for U in range$(0, 30)$ **do** ▷ Loop over all operational wind speeds
      array($CF_{est,w,t,U}$) $\leftarrow$ WEIBULL(U, c_w, k_w) $\cdot$ POWERCURVE$(w, U) \cdot dU$
    $CF_{est,w,t} = \text{sum}(CF_{est,w,t,U})$
    array($CF_{est,w}$) $\leftarrow CF_{est,w,t}$
  return $CF_{est,w}$, mean($CF_{est,w}$) ▷ Store the average CF_{est} value

function WINDINDEX($CF_{est,w}$, mean($CF_{est,w}$))
  for $CF_{est,w,t}$ in $CF_{est,w}$ **do** ▷ Loop over all $CF_{est,w,t}$ values
    if $CF_{est,w,t} < 0.01$ **then**
      $WI_{w,t} = 1$
    else
      $WI_{w,t} = \text{mean}(CF_{est,w})\ /\ CF_{est,w,t}$
    array(WI_w) $\leftarrow WI_{w,t}$
  return WI_w

function CORRELATION(r, $CF_w,\ CF_{est,w}$)
  for $CF_{w,t},\ CF_{est,w,t}$ in $CF_w,\ CF_{est,w}$ **do**
    $\rho_{w,t}$ = PEARSONCORR($CF_{w,(t-r):t}$, $CF_{est,w,(t-r):t}$) ▷ Rolling correlation calculation
    array(ρ_w) $\leftarrow \rho_{w,t}$
  return ρ_w

Algorithm 1 Kelly LTC algorithm (continued)

function LTC

for w in WFs **do** ▷ Loop over all wind farms

$U_{s,w}$ = SCALEWINDSPEED($U_{1,w}$, $U_{2,w}$, z_1, z_2, h_w)
$CF_{est,w}$, mean($CF_{est,w}$) = ESTIMATEPOWER(U_s, POWERCURVE, c_w, k_w)
WI_w = WINDINDEX($CF_{est,w}$, mean($CF_{est,w}$))
ρ_w = CORRELATION(CF_w, $CF_{est,w}$)

for $CF_{w,t}$, $WI_{w,t}$, $\rho_{w,t}$ in CF_w, WI_w, ρ_w **do**

$\Delta CF = CF_{w,t} \cdot WI_{w,t} - CF_{w,t}$

if $CF_{w,t} \leq 10$ **then**
$m = CF_{w,t}/10$
else
$m = 1$

if $\rho_{w,t} \leq 0$ **then**
$\rho_{w,t} = 0$ ▷ CF cannot be negative

$\Delta CF = \Delta CF \cdot m \cdot \rho_{w,t}$

$LTC_{w,t} = CF_{w,t} + \Delta CF$

if $LTC_{w,t} > 100$ **then**
$LTC_{w,t} = 100$

array(LTC_w) $\leftarrow LTC_{w,t}$

return LTC_w

($\Delta CF \cdot D_t$) and age_t is approximately equal to zero then the data correction is not affecting the estimate of the aging trend, except by correcting for errors.

Figure 6.9 shows the correlation between ($\Delta CF \cdot D_t$) and age_t for each wind farm in the dataset, with the average being -0.036 and -0.031 for the ERA5 and MERRA-2 LTC calculation respectively. The average is approximately equal to 0, however there are some outliers. If all outliers that are ≥ 0.2 and ≤ -0.2 are dropped from the LTC analysis (2.09% and 1.61% of all wind farms for the

ERA5 and MERRA-2 calculations respectively), the LTC results do not materially change due to the average correlation ≈ 0. This indicates a statistically sound data modification process. To understand this further, $(\Delta CF \cdot D_t)$ can be regressed against age_t which shows that there is no clear trend between the data correction and age. Figure 6.10 shows that there is no trend between these elements and that there are very little fluctuations between wind farms, indicating that the average of the correlations between $(\Delta CF \cdot D_t)$ and age_t of -0.036 and -0.031 for the ERA5 and MERRA-2 LTC calculation respectively, are statistically insignificant.

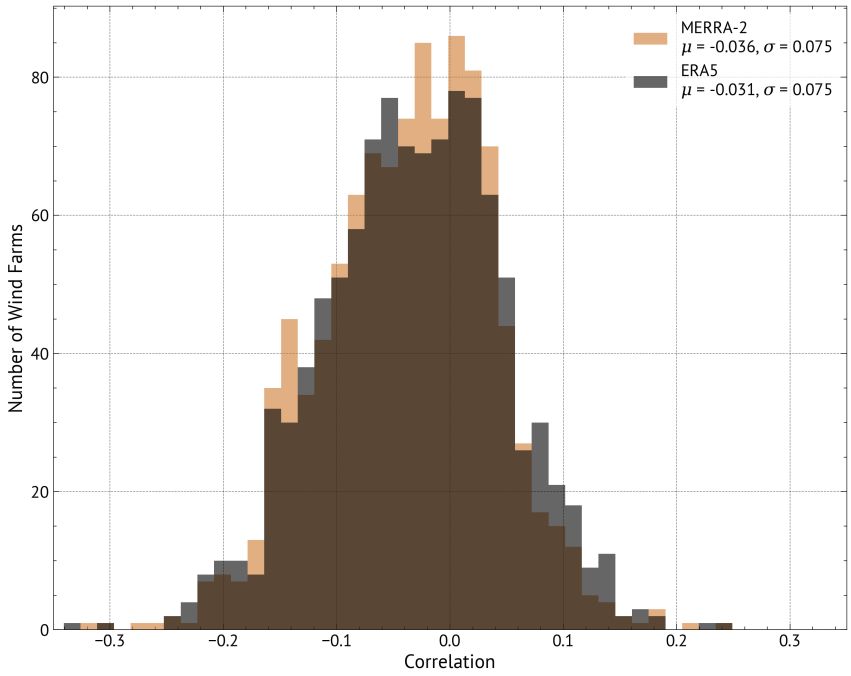


Figure 6.9: Correlation between the error correction in the LTC process and age for each wind farm in the dataset.

Linear regression

In order to draw conclusions from linear regressions, several assumptions need to be met [162]:

1. *Linearity*: The relationship between X and the mean of Y is linear.

2. *Existence*: For any value of X, Y is a random variable with finite mean and variance.

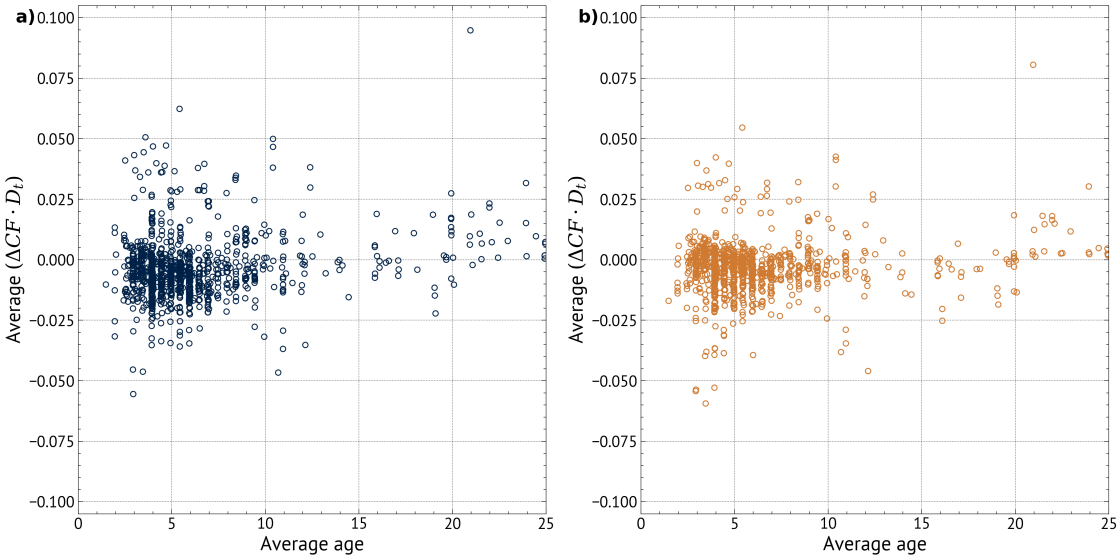


Figure 6.10: Average age of each wind farm plotted against the LTC correction process ($\Delta CF \cdot D$) described in Equation 6.8 calculated using reanalysis wind speed data from (a) ERA5 and (b) MERRA-2.

3. *Normality*: For any fixed value of X, Y is normally distributed.

4. *Homoscedasticity*: The variance of the residual is the same for any value of X.

5. *Independence*: The Y-values are statistically independent.

The residuals were analysed for each regression in order to see if the data satisfied these requirements. Residuals are calculated as the difference between the predicted values (linear regression) and observed values. Both normality and homoscedasticity can be observed from the residuals, which should be normally distributed and have constant variance. For all data in this thesis, existence is satisfied. Independence of age and important regression variables are also satisfied in all regression examples, which was shown in Figure 6.10. Figure 6.11 shows fictitious examples of regressions that fail linearity and homoscedasticity, as well as a data compliant example using a linear regression.

The normality and homoscedasticity conditions were occasionally violated for some wind farms, in most cases due to extended farm downtime. These cases were uncommon and do not effect the final results. Furthermore, downtime is an important component of performance degradation, so these data were not removed.

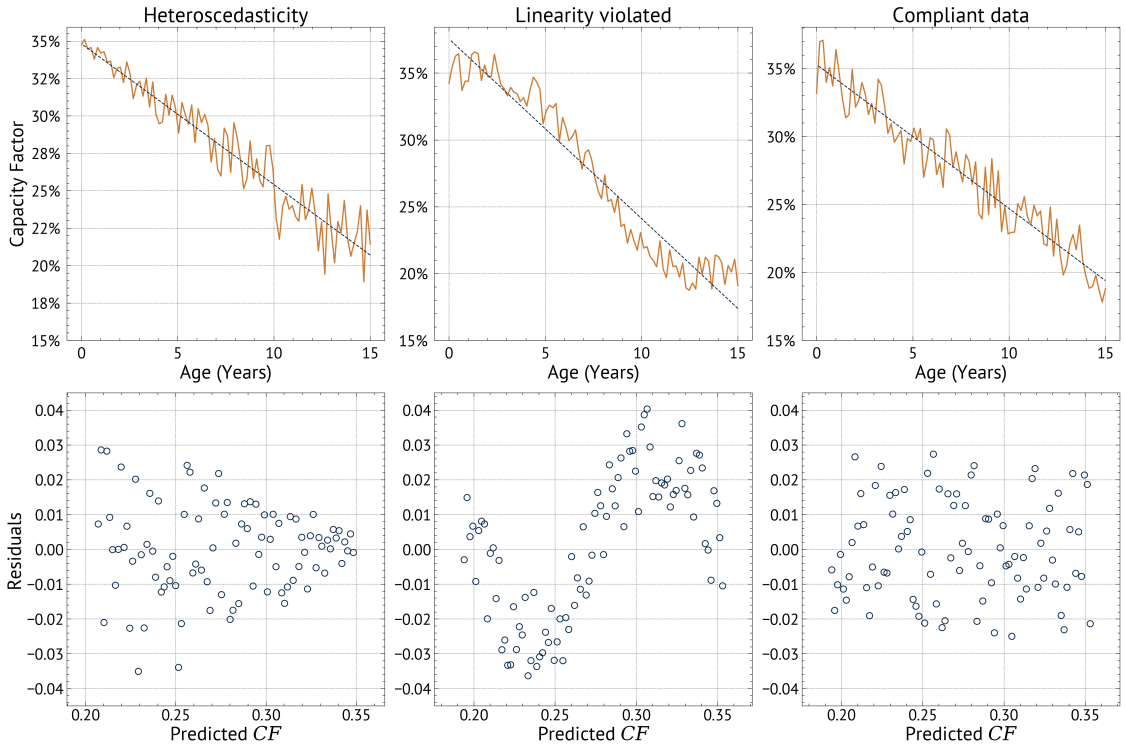


Figure 6.11: Fictitious data showing examples of heteroscedasticity (left), linearity violated (middle), and compliant data (right). The upper row is plotting age against CF and the bottom row shows predicted CF (from the linear regression) against residuals.

Heteroscedasticity was observed in the LTC analyses where older wind farms (age > 15), built using older wind turbine designs (e.g. before the year 2000) had large variability in CF_{obs} and CF_{LTC} values. These results challenge the linear regression assumption; however, when using newer technologies, heteroscedasticity does not exist, indicating that it is not a persistent trend. Violating linearity is more serious because it indicates that the regression model may not be representative of the true trend; however, this very rarely occurred due to the implementation of Equation 6.2, which accounts for seasonality terms. Importantly, when wind farm data are aggregated by country (the key findings of this chapter), the linear regression rules are satisfied.

6.5 Importance of Wind Farm Size

In many ways, the largest wind farms are the most important – they supply a larger proportion of demand and they are the most capital intensive. Thus, their

degradation rates are more important from both grid management and investment perspectives. It is a matter of scale: if a 200 MW wind farm degrades at a higher rate than a 20 MW wind farm, this needs to be clearly understood and factored into future-looking analyses. Considering the largest wind farms also gives greater importance to more recent wind farms, as the average size of a wind farm has increased over time (Figure 6.12).

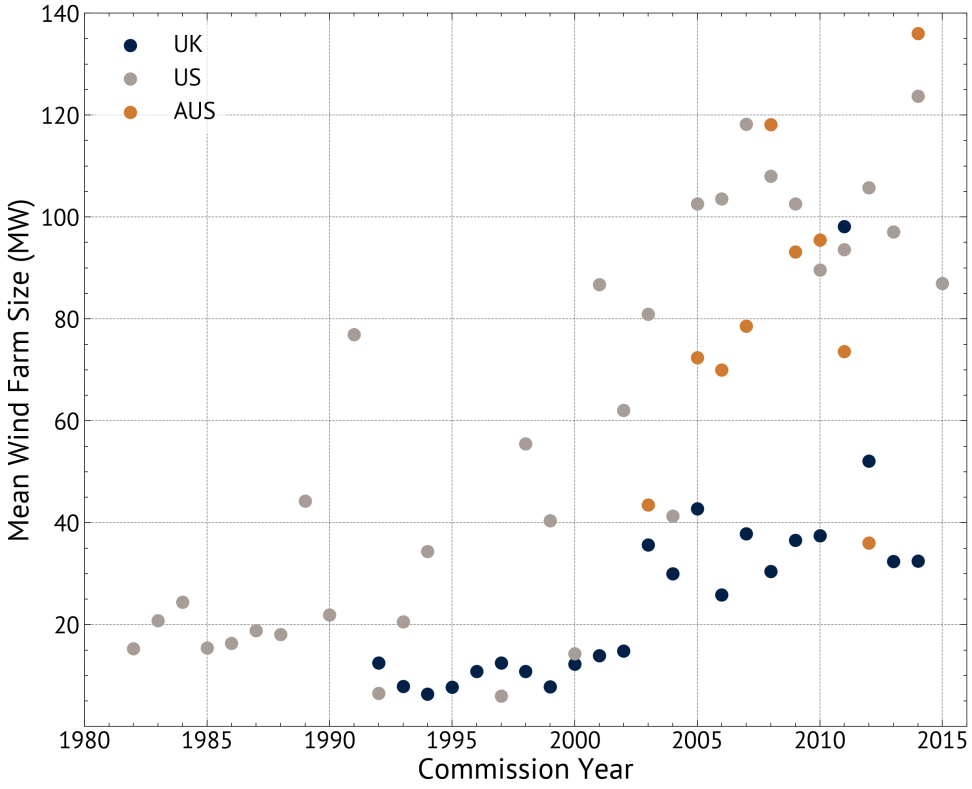


Figure 6.12: The mean wind farm capacity by commission year for UK, US, and AUS.

Therefore, weighted averages and weighted regressions have been completed to highlight performance decline in the wind farms that most significantly contribute to electricity demand. Weighted averages and regressions were completed with the total installed capacity of the wind farm being the weighting factor. Weighted averages were taken when aggregating CF data across wind farms for a time period (monthly or yearly) using Equation 6.9:

$$\overline{CF} = \frac{\sum_{i=1}^{n} w_i \cdot CF_i}{\sum_{i=1}^{n} w_i} \tag{6.9}$$

where $\overline{CF}$ is the weighted average capacity factor, n is the number of wind farms being averaged over, and w is the weight (the rated capacity of the wind farm). In cases where percentiles were used, these were also calculated using weighted-average ranges.

For this study, two sets of results will be presented distinctly, a non-weighted and weighted set.

6.6 Results

6.6.1 Reanalysis wind speed data

To make this study more robust to the input data source, the two most predominant reanalysis wind speed models were compared. A byproduct of having two sets of wind speed data is that a comparison between their outputs can be completed to capture how they are correlated and the difference in magnitude between their values. This is important because the actual wind speed for each time period at each wind farm is unknown and the reanalysis models are only an approximation.

If the reanalysis models are highly correlated (as calculated using a Pearson correlation [161]), it gives more confidence that the reanalysis wind speed trends represent actual local conditions; however, they could also have the same underlying modelling errors. Figure 6.13 displays the correlation between ERA5 and MERRA-2 wind speeds for all wind farms using the same time resolution data as the available power data (monthly for the UK and US, hourly for AUS), separated by country. It can be seen that at most wind farms, the reanalysis models are highly correlated (> 0.9); however, there are cases, mainly in the US, where the reanalysis models are not highly correlated. This is especially important because it would be expected that the US data would be more highly correlated than the AUS data because they are calculated using monthly data, meaning the reanalysis models diverge significantly. It is also clear that the reanalysis models are very highly correlated in the UK.

By quantifying the differences between the ERA5 and MERRA-2 data sets, potential biases can be identified. The data could be well correlated but at different wind speed magnitudes, which has implications for wind farm power production.

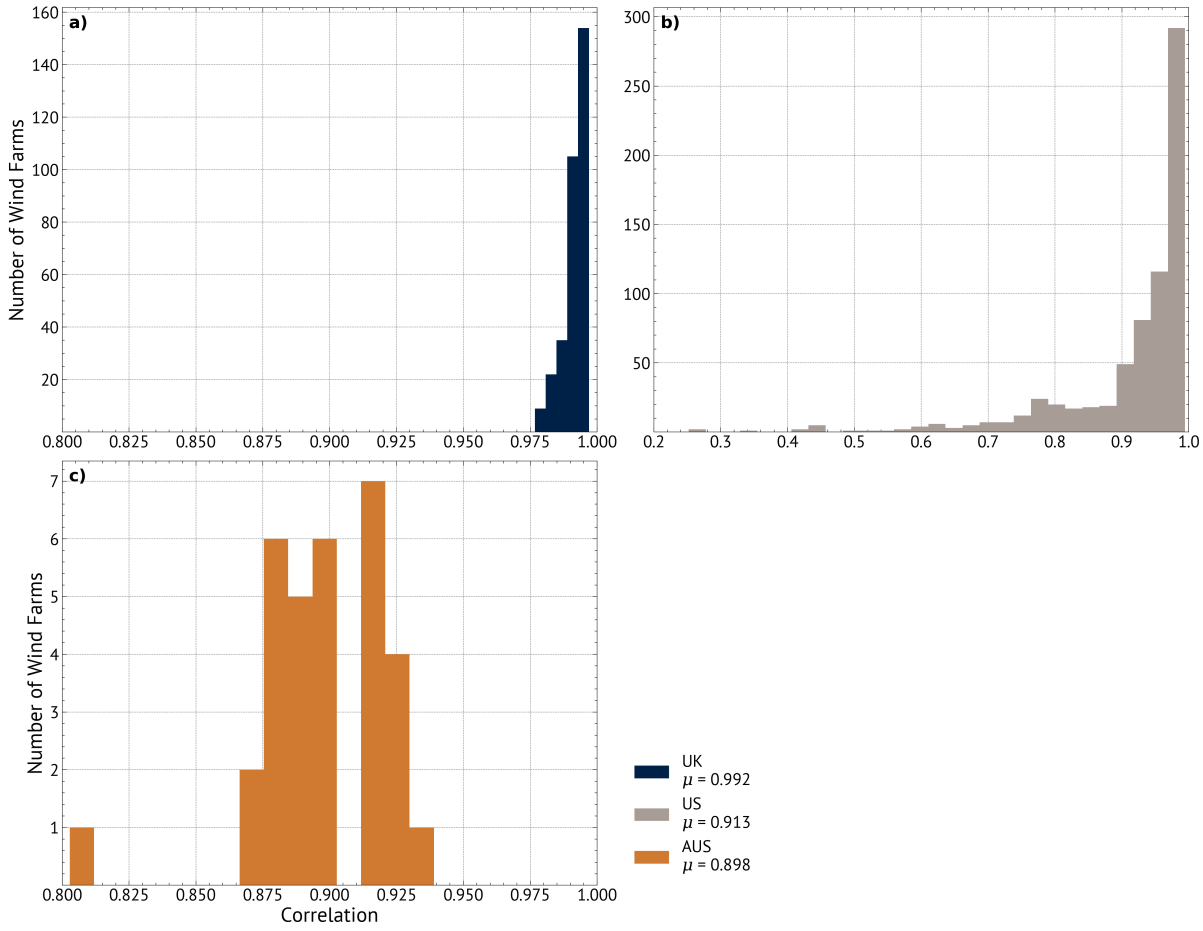


Figure 6.13: The Pearson correlation between ERA5 and MERRA-2 wind speed data at each wind farm for (a) UK, (b) US, and (c) AUS.

Figure 6.14 shows the mean difference between the wind speed time series for the two reanalysis datasets for each wind farm location. The difference in wind speed is calculated by averaging the difference between the wind speed values for each time period, shown in Equation 6.10:

$$\Delta U = \frac{\sum_{i=1}^{N} U_{ERA5,i} - U_{MERRA-2,i}}{N} \tag{6.10}$$

where i represents each time period and N is the length of the time series. This shows that for the majority of wind farm sites, MERRA-2 has an approximately 1 m/s higher wind speed on average. This is reflected in the mean wind speed across all wind farm time series data, by country, which is shown in Table 6.2.

Table 6.2: Average wind speed across the wind farm sites, separated by country – AUS, UK, US and all three countries combined (ALL). All values are displayed in units of m/s.

Country	ERA5	MERRA-2
UK	6.843	7.968
US	5.827	6.693
AUS	6.400	6.949
ALL	6.158	7.094

This means at least one of the reanalysis datasets is biased compared to observed wind speeds and also serves to emphasise the fact that reanalysis wind speeds are not always accurate. For LTC methods, biases are less important than the correlation of predicted CF with observed CF, because corrections are typically made solely from the CF_{est} data series, with no knowledge of the CF_{obs} data series. That is, the CF_{est} data series is used to correct the CF_{obs} based on the long term average of CF_{est}; as long as the two CF series are highly correlated, the correction will occur in the right direction even if the magnitudes are different (see Section 6.4). This therefore makes the correlation between CF_{est} and CF_{obs} the crucial element.

6.6.2 Wind speed trends

One reason raw wind farm power outputs need to be long-term corrected is to remove any underlying wind speed trends. This is important because if the average wind speed at a site increases or decreases over time, then so does the wind farm performance trend. Emphasising this, Sinden [163] calculated that variability in energy delivered by UK wind farms per year was 7.4%. The impact of varying wind speed is accounted for in the LTC algorithm; however, some of the initial trends shown in Figures 6.2 and 6.7 can be compared to raw wind speed trends. This is done to see if there are direct correlations between the two that might provide a simple explanation of the observed trends.

The basic assumption is that the long-term average wind speed at a given location is unchanging over time. To test this, the wind speed time series at all wind farm sites is analysed. A linear regression was performed on the wind speed at every wind farm site and the slopes were aggregated by country to form Figure

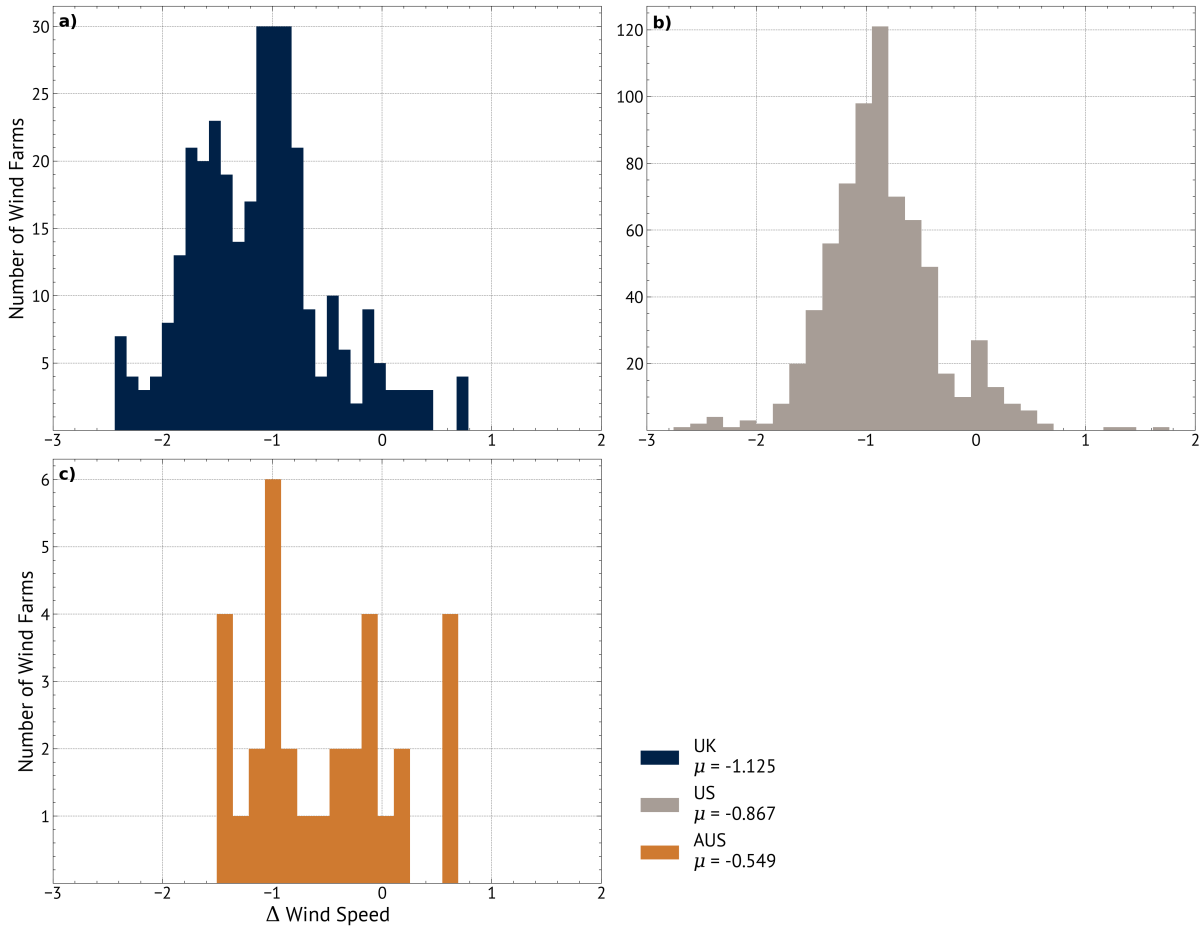


Figure 6.14: The mean difference in m/s between ERA5 and MERRA-2 wind speed data for each wind farm for (a) UK, (b) US, and (c) AUS.

6.15. A linear regression was used to extract the long-term trend at each site, but it does not indicate that wind speed is best represented by a linear regression. The average slope has also been indicated in the legend for each subplot. This figure shows a congregation of trends around zero but also shows that, on average, the mean wind speed has decreased for UK and US wind farms and has increased for AUS wind farms. This is significant because it indicates that for the UK and US, the decrease in performance shown in Figures 6.2 and 6.7 may be exaggerated by a long-term decrease in wind speed. It also indicates that the degradation seen in AUS wind farms may be more severe than the CF observed raw data suggests. The trends seen in Figure 6.15 occur for individual wind farms across different locations and time periods depending on the specific wind farm. Therefore, these do not necessarily indicate wind speed trends over time, but rather show

wind speed trends as wind farms age.

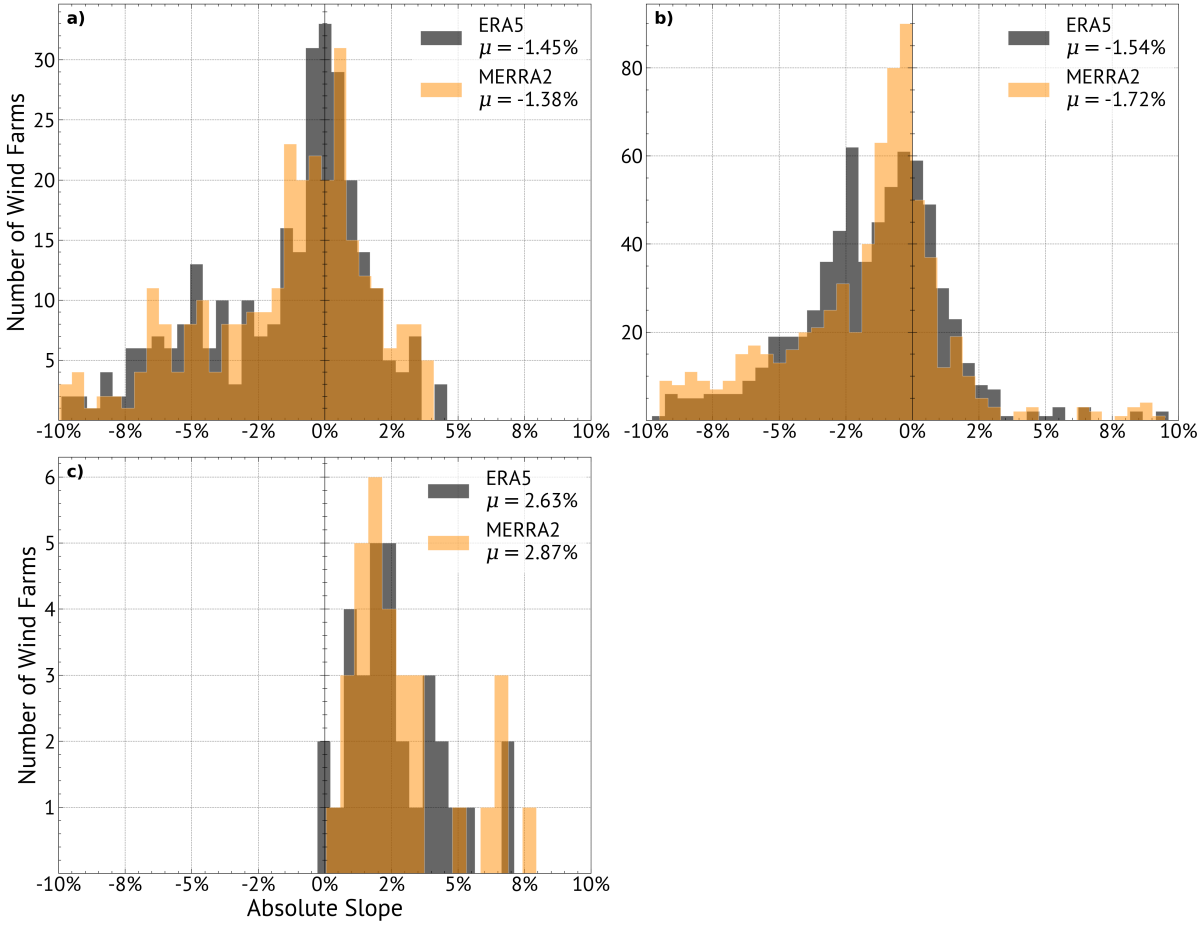


Figure 6.15: The absolute wind speed time series regression slopes at individual wind farm locations using both ERA5 and MERRA-2 data for (a) UK, (b) US, and (c) AUS. On average, at individual wind farm locations wind speed is shown to decrease for the UK and US, and increase for AUS.

To obtain trends over time, wind speeds are averaged by time (monthly increments) and country to produce Figures 6.16 and 6.17 for ERA5 and MERRA-2 data respectively. These figures show monthly averaged data, and a linear regression was fitted to the data to understand the trend of the data (i.e. was the wind speed in each country generally increasing or decreasing over the time period used in this study). This does not indicate that a linear regression most appropriately fit the seasonality of the data, it was simply used to get an understanding of a long term trend. The trend has been calculated in units of m/s per year. It is clear from Figures 6.16 and 6.17 that wind speeds across US wind farms showed a higher variance at any point

in time than the UK or AUS wind farms. This is because the US wind farms in this dataset exist over a much larger geographical area than the UK or Australian wind farms, and they therefore experience a larger variation in wind speeds at a given time. Note, the Australian wind farms used in this study exist only in the south-east corner of the country. This analysis was also completed on annual data, and the trends for both monthly and annual averages are displayed in Table 6.3 and were calculated over the time periods 2003–2020, 2001–2019 and 2009–2020 for the UK, US and AUS respectively. Trends are calculated using a linear regression for both monthly and annual averages, and the absolute trend values relative to the mean wind speed (Table 6.2) for that country are also displayed in brackets in Table 6.3.

Table 6.3: Summary of wind speed trends using ERA5 and MERRA-2 data at wind farm sites, separated by country – UK, US and AUS. The trends were calculated over the time periods 2003–2020, 2001–2019 and 2009–2020 for the UK, US and AUS respectively. Trends are calculated using a linear regression for both monthly and annual averages, but are both displayed in m/s per year units with the absolute values relative to the average wind speed shown in brackets.

Country	ERA5 monthly	ERA5 annual	MERRA-2 monthly	MERRA-2 annual
UK	0.024 (0.35%)	0.005 (0.07%)	0.034 (0.43%)	0.017 (0.21%)
US	0.099 (1.70%)	0.086 (1.48%)	0.089 (1.33%)	0.077 (1.15%)
AUS	0.007 (0.11%)	-0.018 (-0.28%)	-0.015 (-0.22%)	-0.034 (-0.48%)

6.6.3 Comparing estimated and observed power output

After estimating the power for each time period for each wind farm using the method described in Section 5.5.5, the results can be compared with observed power output. Ideally, these values would be very similar, especially in the early years of operation (before degradation); however, reanalysis wind speeds are not completely accurate, and given the differences in reanalysis wind speed from the ERA5 and MERRA-2 models (Section 6.6.2), there are discrepancies. These

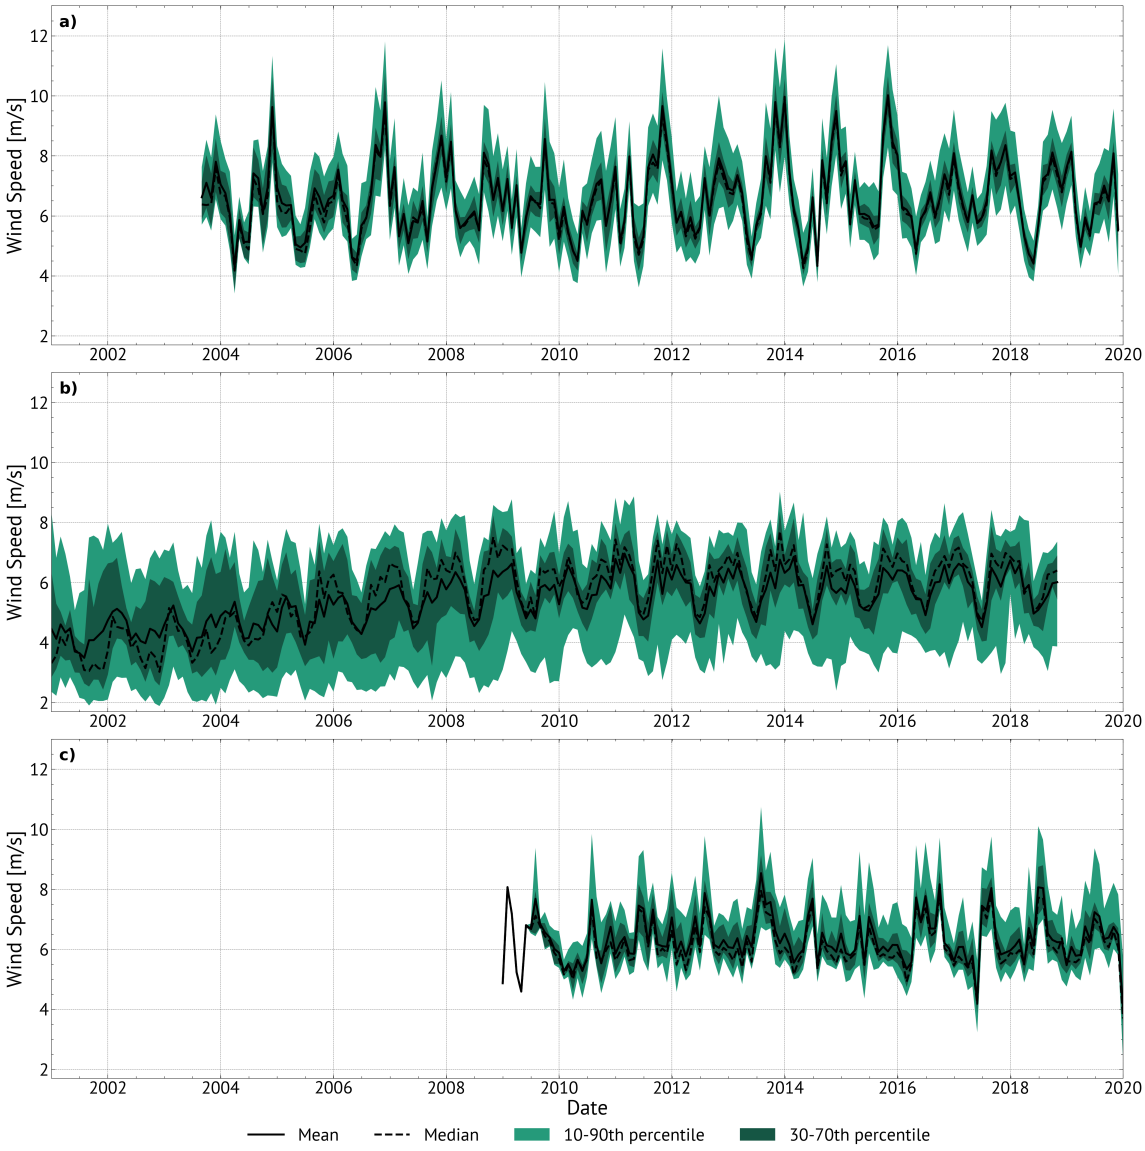


Figure 6.16: Monthly averaged ERA5 wind speed trends for wind farm locations in (a) UK, (b) US, and (c) AUS.

differences, where present, can be magnified by the non-linear relationship between power production and wind speed.

Figure 6.18 shows the average predicted capacity factor ($\overline{CF_{est}}$) vs the average actual capacity factor for each wind farm. There is clearly a spread of results around the one-to-one relationship line, indicating errors in predicting the power. Some level of error is unavoidable due to inaccuracies in the reanalysis wind speed data and the unknown operating conditions at the wind farm (e.g. degradation, curtailment, etc.). These is no clear bias in the US or AUS results; however, the

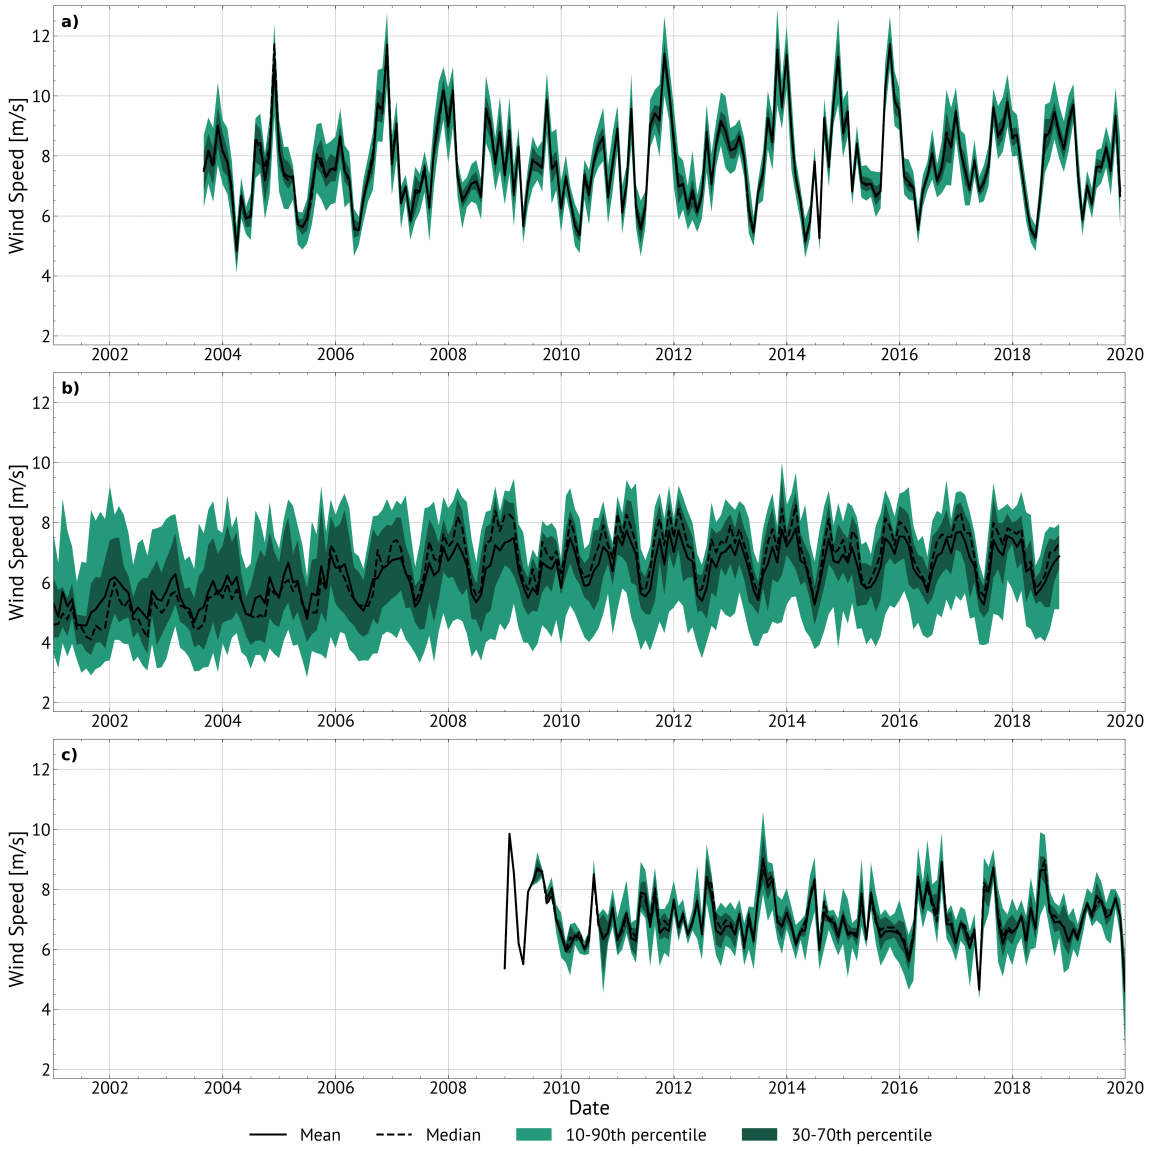


Figure 6.17: Monthly averaged MERRA-2 wind speed trends for wind farm locations in (a) UK, (b) US, and (c) AUS.

estimated power for the UK is systematically over-predicted, indicating a consistent bias in reanalysis wind speeds.

Despite differences in averaged outputs, the estimated CF and observed CF are strongly correlated (mean wind farm has a Pearson correlation between time series of greater than 0.7) for all countries. Figure 6.19 shows rolling correlation for the UK, US and AUS. The UK and US have a 3-month rolling correlation and AUS has a 1-week rolling correlation (the same as used in the LTC method). The AUS histogram (c) has a broader peak due to more granular data leading to a

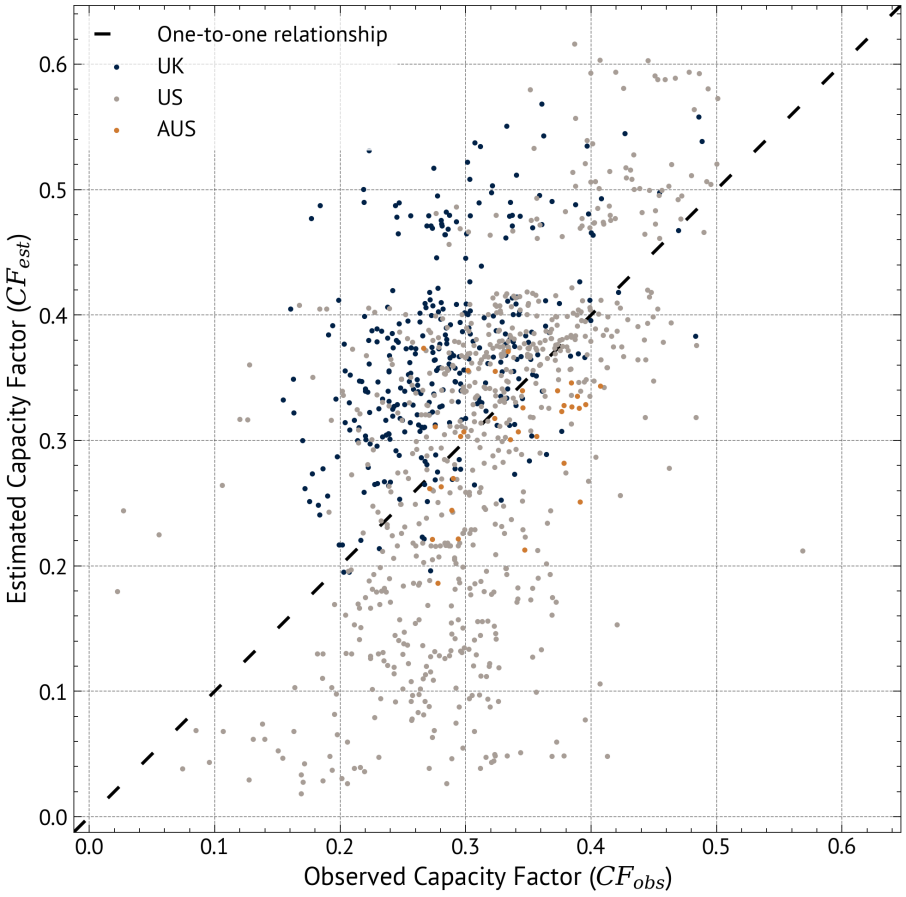


Figure 6.18: Scatter plot comparing average estimated capacity factors calculated using ERA5 data against average actual capacity factors for each wind farm in the UK (blue), US (grey) and AUS (orange).

wider range of correlation values. The mean correlation for the UK, US and AUS is 0.843, 0.706 and 0.786 respectively. This is a high level of correlation, although it is still appropriate to use the LTC method described in Section 6.4 to mitigate the effects of the non-correlated data points.

6.6.4 Long-term correction and performance decline

The purpose of using a LTC method is to remove the effects of varying wind speed to see the underlying performance trend – effectively separating the noise from the signal. Working from this premise, the novel LTC method presented in this thesis was compared to previously-cited methods using variation away from a linear trend line. The linear regression line was fit for each wind farm because the underlying performance trend for each wind farm is assumed to be approximately linear over

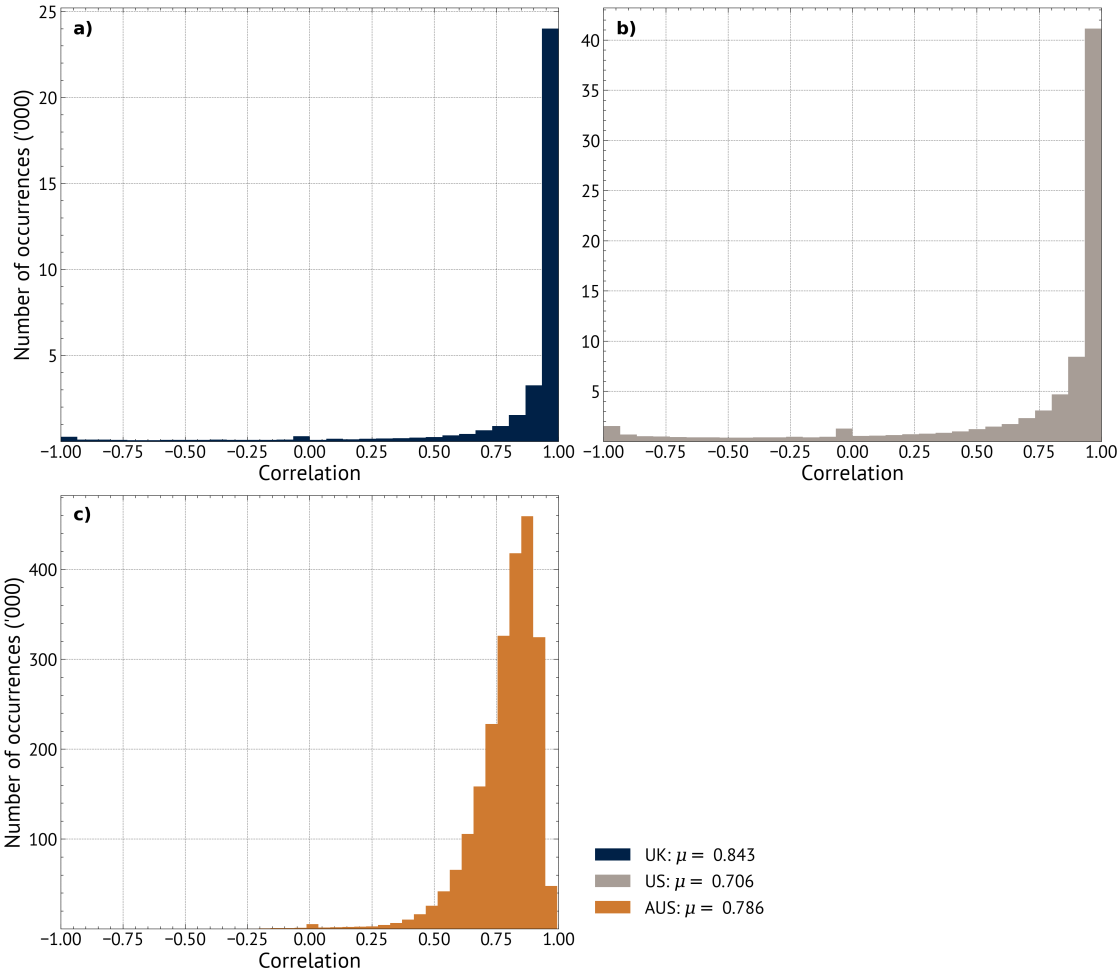


Figure 6.19: Rolling correlation between CF_{est} and CF for (a) the UK, (b) US, and (c) AUS. The correlation period for the UK and US is three months and is one week for AUS.

the lifetime of a wind farm, which reflects the raw data and results from previous studies. MAE was used as the metric for quantifying the variation away from the regression line. MAE is calculated using Equation 6.11 where $CF_{reg,i}$ represents the equivalent point on the regression line to $CF_{LTC,i}$. MAE provides an average measure of error between two series and thus can be used as a mean distance from the regression line, or a measure of average variability.

$$MAE = \frac{\sum_{i=1}^{N} |CF_{reg,i} - CF_{LTC,i}|}{N} \tag{6.11}$$

Table 6.4 quantifies the mean MAE across all wind farms for this novel method (labelled 'Kelly method') against unmodified data and previous methods. Data are presented using both the MERRA-2 and ERA5 reanalysis wind speeds as inputs.

The Olauson et al. method provides an improvement from the Staffell and Green LTC method; however, there remains a lot of variability compared to the novel Kelly method. The MAE increases for previous methods compared to raw data in AUS and ALL because the presence of higher frequency data leads to over correction in the LTC process for small CF_{obs} values.

Table 6.4: Summary of MAE values resulting after the implementation of LTC methods and fitting a linear regression line. The reanalysis wind speed datasets used for the LTC calculation are given in the parentheses.

Method	**ALL**	**UK**	**US**	**AUS**
Raw data (CF)	0.0855	0.2547	0.0858	0.0776
Staffell and Green method (MERRA-2)	0.1087	0.2299	0.0881	0.1127
Staffell and Green method (ERA5)	0.0932	0.2291	0.0701	0.0978
Olauson et al. method (MERRA-2)	0.0823	0.1895	0.0736	0.0814
Olauson et al. method (ERA5)	0.0834	0.2070	0.0744	0.0819
Kelly method (MERRA-2)	0.0517	0.1847	0.0366	0.0527
Kelly method (ERA5)	0.0491	0.1725	0.0368	0.0492

Figure 6.20 is an example of the LTC method for the Ardrossan wind farm in Scotland between age 10–15 years (corresponding to years 2013–2018), across three LTC calculation stages: reanalysis wind speed, power prediction, and LTC values. This figure gives a 5-year example breakdown of the Kelly LTC process and how it differs from previous methods. There is a notable difference between the variability of the Kelly LTC data compared to the raw data and previous LTC methods. Previous methods tend to over/under predict power output for high/low wind speeds because they use a point wind speed to calculate power; this is unrealistic across the time period it represents because it does not capture the wind distribution experienced by the wind farm over that period. This will be exaggerated for wind speeds near the cut in speed especially because the CF_{est} can suddenly fall to zero, which is unrealistic especially over the time period of a month. This will cause the LTC calculation to over-correct the CF_{obs} value, often creating an inverse CF chart, but not actually correcting the value towards the

underlying trend. The Kelly method avoids over-correction by correcting the power values towards a more central value regardless of the season.

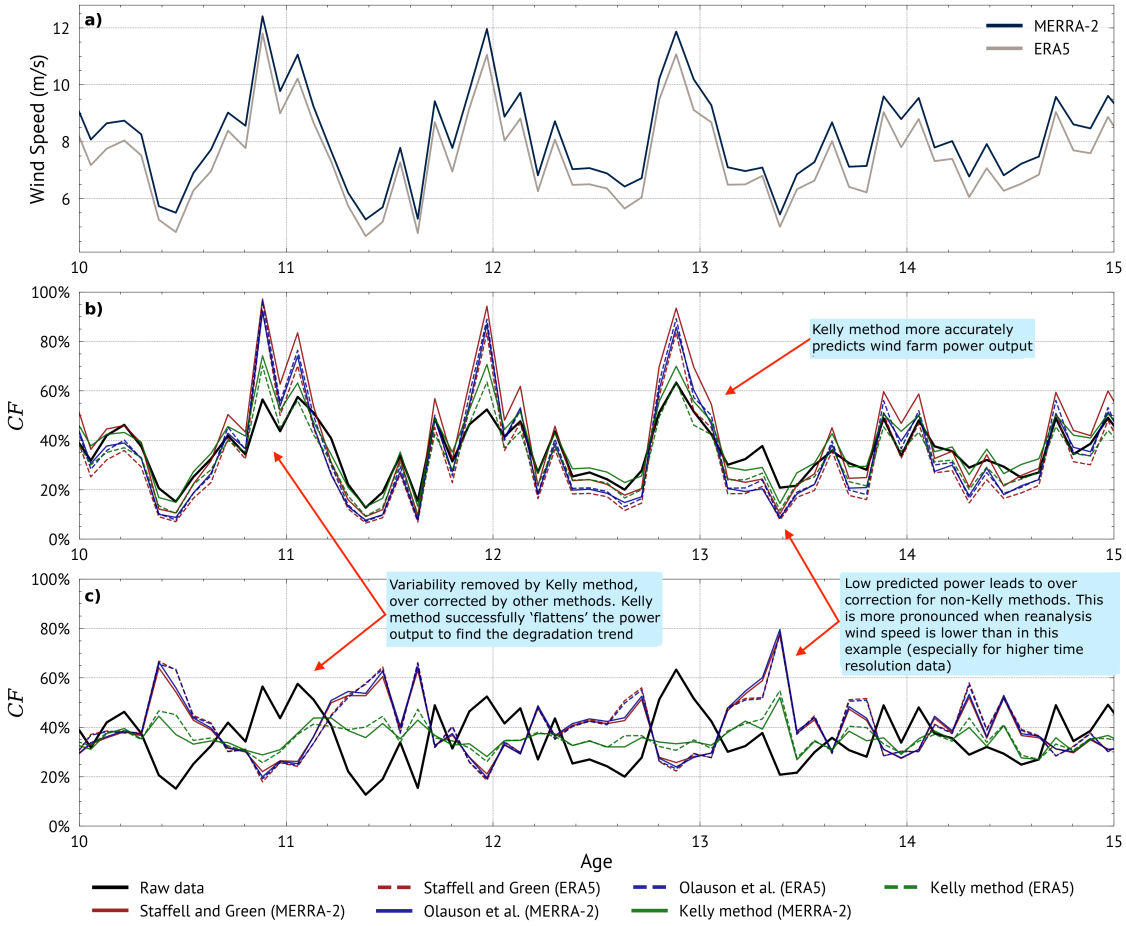


Figure 6.20: A comparison of LTC methods in across a 5 year period for the Ardrossan wind farm in Scotland. Shown is (a) the reanalysis wind speed, (b) predicted CF and (c) CF_{LTC}.

Figure 6.21 recreates Figure 6.20 for a 1 day period for the Capital wind farm in Australia at age 7 (corresponding to the 24 hours following 10am on the first of January 2017). This shows some of the variability that can occur with a LTC process when the data have a higher time resolution. A low wind speed event is shown where the wind farm fluctuates between a reanalysis wind speed above and one below the cut-in wind speed (4m/s for this wind farm which has 67 Suzlon S88 2.1 MW wind turbines; power curve shown in Figure 5.10). Thus, the methods which use a 'point' wind speed power estimation have large fluctuations between very low power, which over-corrects the CF_{obs} data due to a low WI value, and

zero power, which results in no correction. This is clearly a problem due to the difficulty in uncovering a long-term trend from data fluctuating between $CF_{LTC} \approx 0$ and $CF_{LTC} \approx 100$. The Kelly LTC method is able to handle this situation through two mechanisms: implementing a scaling factor, m (Equation 6.6), which regulates how much the data are corrected for low CF_{est} values; and by using a wind speed distribution to calculate power, meaning that power will not go to zero when the reanalysis wind speed is close to the cut-in wind speed. This figure is an example of how the Kelly LTC method can apply to data of varying time resolutions.

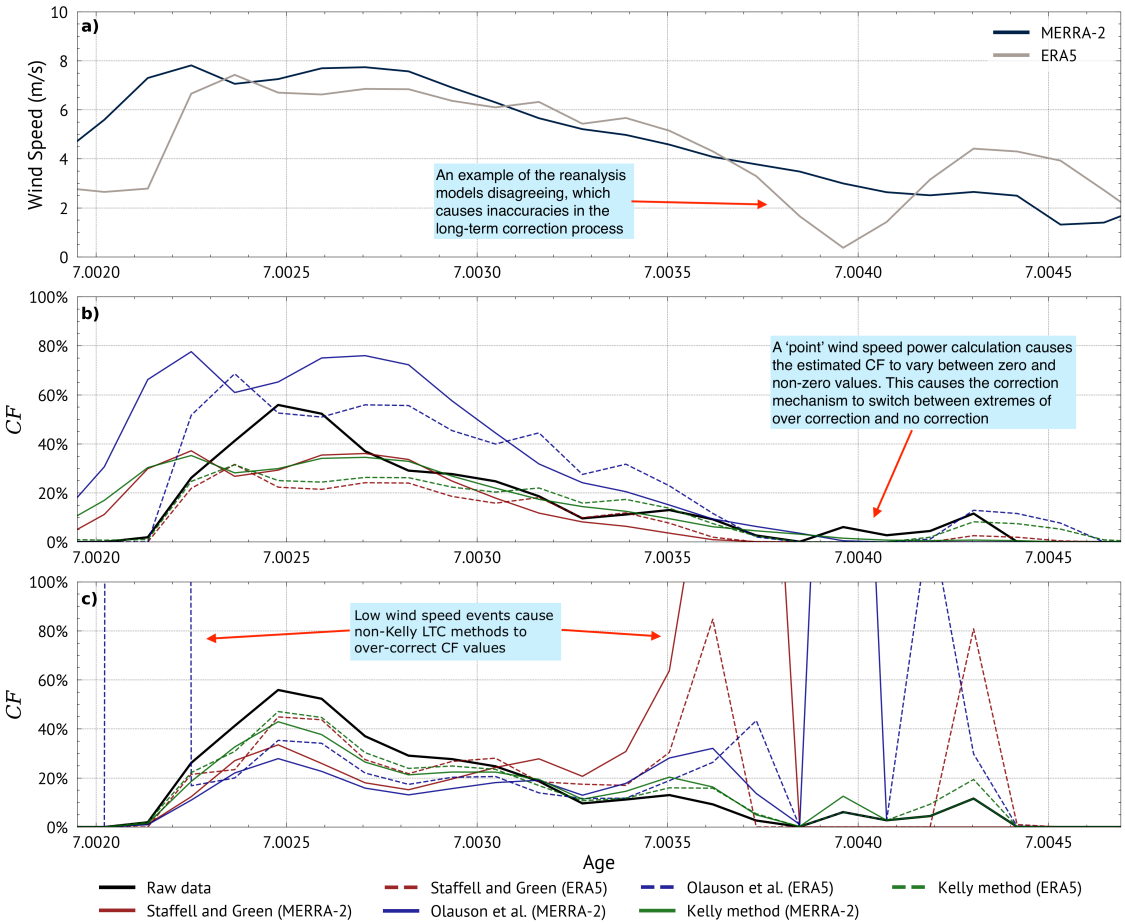


Figure 6.21: A comparison of LTC methods across a 1 day period for the Capital wind farm in Australia at age 7 (i.e., the time period between age 7 years and 0 days to age 7 years and 1 day). Shown is (a) the reanalysis wind speed, (b) predicted CF and (c) CF_{LTC}.

The result of implementing the Kelly LTC method using ERA5 reanalysis wind speeds to correct national wind power data is seen in Figure 6.22. The corrected

data show a similar trend to the raw CF data, but with an increased degradation rate in general. For this UK and US this is not the expected result as wind speed, on average, were reducing over time. However, this could be due to the change of proportion of time operating in region of peak C_P for which there exists a more noticeable reduction in power for degraded systems than above rated flow speed. Table 6.5 details the changes in degradation slopes between the raw and corrected data. The yearly means show a close agreement with the linear regression line; however, the US shows a sharp dip between ages 13 and 15. Hamilton et al. [21] saw a similar trend in US wind farms; however, the 'jump' down in performance at age 10 was attributed to the loss of the production tax credit that occurred at this age. The results shown here show that the performance jump occurs at a later age, suggesting that it may due to a technological, operational or design improvement of younger turbines or to a poor performing cohort of turbines.

Figure 6.23 splits the US wind farms into two sample sets: those with a final age (in the dataset) of greater than 13, and those with a final age less than 13. This split has 142 wind farms greater than age 13 with a median commencement year of 2001, and 554 wind farms less than age 13 with a median commencement year of 2010. This figure shows that it is the older cohort of wind farms that causes the exaggerated drop in performance at age 13. With the younger wind farms removed, the US trend is more linear (although there is still a relative resistance to degradation until age $\approx 12 - 13$), and performance is declining at a lesser rate (0.62 CF percentage points per year compared to 0.72) without the higher average CF_{LTC} values in the dataset skewing the early averages. The wind farms with age less than 13 show a lower performance decline with a rate of 0.04 CF percentage points per year. Effectively, after a certain wind farm commencement date the technology is more resistant to ageing effects, though it is unknown how the younger cohort of wind farms will degrade as they continue to age.

US wind farms degrade at the highest rate (0.72 CF percentage points); however, this is slightly exaggerated by the difference in average CF_{LTC} for the younger cohort of turbines. Generally, US wind farms (Figure 6.22 (b)) have a relatively

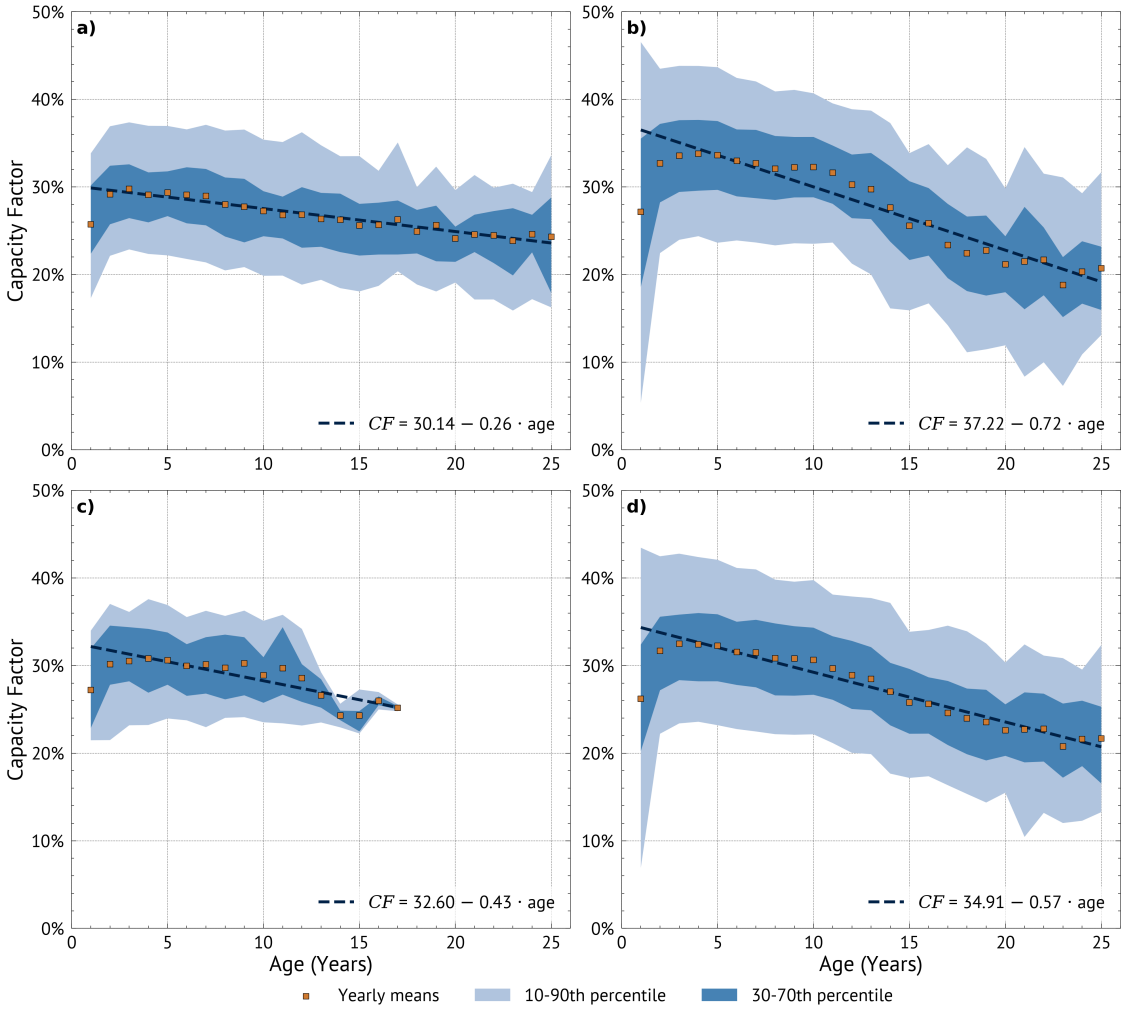


Figure 6.22: Wind farm performance plotted against age for the long term corrected capacity factors (using the Kelly method and ERA5 reanalysis wind speeds) for wind farms in (a) the UK, (b) US, (c) AUS and (d) ALL (as per Figure 7.2 for raw data only). A linear regression line was fitted to the yearly averaged CF values across all wind farms.

high early-age average CF with similar degradation to the UK and AUS until age 13 (0.30 CF percentage points per year from age 1 to 13). After the drop between ages 13 and 15, the degradation rate is 0.57 CF percentage points per year (for ages 15 to 25). Therefore, the degradation rate is exacerbated by the drop in performance between ages 13 and 15.

Figure 6.24 shows the average degradation rate across all farms separated by commission year and country – UK (a), US (b), AUS (c) and ALL (d). As

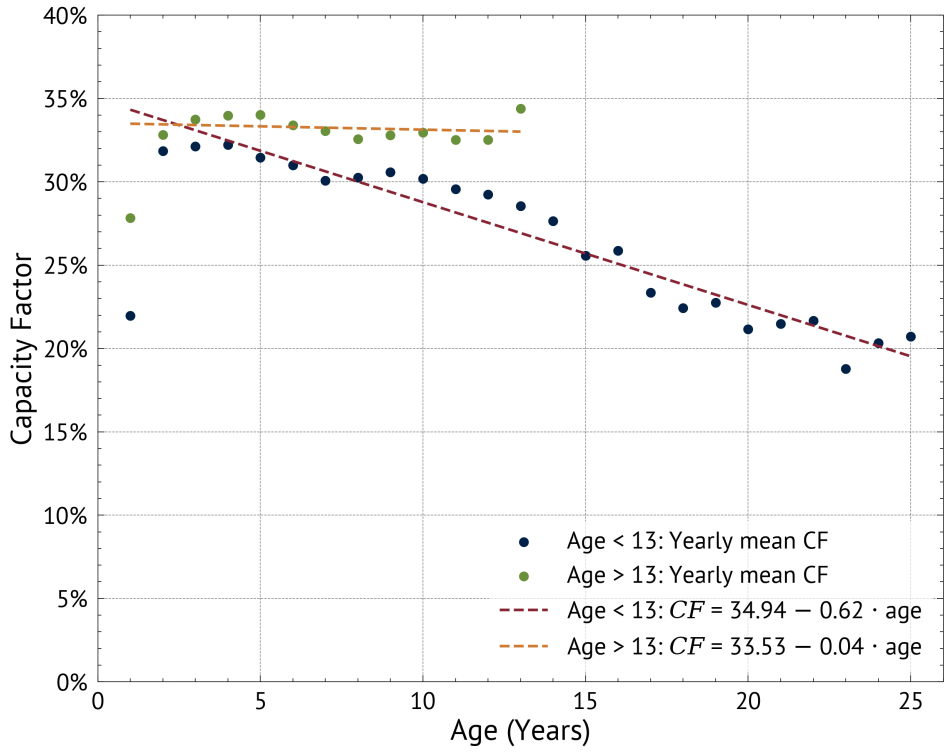


Figure 6.23: Wind farm performance plotted against age for the long-term corrected capacity factors (using the Kelly method and ERA5 reanalysis wind speeds) for wind farms in the US separated by wind farms with a final age (in the dataset) of above or below age 13 years. A linear regression line was fitted to the yearly averaged CF values across all wind farms in each sample set.

commissioning year becomes more recent, the average length of time series decreases and the average degradation rate increases (performance increases). This indicates that as technology, operation and design improves, performance loss is less severe (or improves). However, these wind farms may still experience performance losses as they continue to age. Additionally, there are two outlier years for the US (2007 and 2008) that show a large degradation. The cause for this is unclear, however it could contribute to some increased variability in the US results because these performance declines likely result from factors other than just wind speed variability.

The monthly averaged LTC capacity factors (calculated using the Kelly method and ERA5 reanalysis wind speeds) with the fitted linear regression line can be seen in Figure 6.25. Notably across all countries the degradation rates are similar between monthly and annual averaged CF values. It is clear that some seasonality in the data still exists, especially in the data from US wind farms that are over 15

Table 6.5: Linear regression equation comparison between raw data and LTC data using yearly averaged data.

Country	Raw data (CF)	LTC data (MERRA-2)	LTC data (ERA5)
UK	$CF = 29.39 - 0.23 \cdot age$	$CF = 29.41 - 0.20 \cdot age$	$CF = 30.14 - 0.26 \cdot age$
US	$CF = 36.01 - 0.64 \cdot age$	$CF = 36.95 - 0.69 \cdot age$	$CF = 37.22 - 0.72 \cdot age$
AUS	$CF = 35.62 - 0.36 \cdot age$	$CF = 33.22 - 0.43 \cdot age$	$CF = 32.60 - 0.43 \cdot age$
ALL	$CF = 33.86 - 0.49 \cdot age$	$CF = 34.48 - 0.52 \cdot age$	$CF = 34.91 - 0.57 \cdot age$

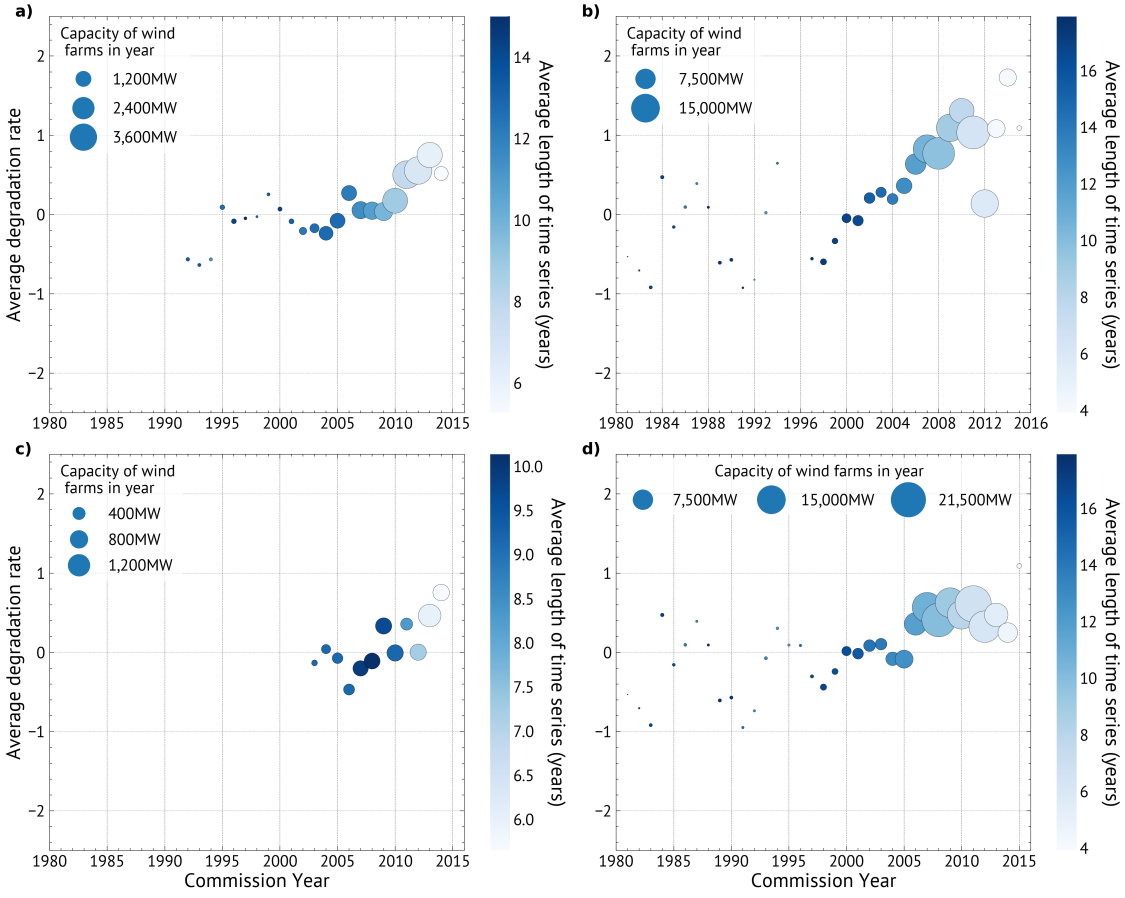


Figure 6.24: Wind farm degradation rates averaged by commission year and calculated using the long-term corrected capacity factors (using the Kelly method and ERA5 reanalysis wind speeds) for wind farms in (a) the UK, (b) US, (c) AUS and (d) ALL. A linear regression line was fitted to the yearly averaged CF values across all wind farms.

years old; however, it is dramatically reduced from the raw data (Figure 6.7). There are two factors at play: firstly, LTC methods are not perfect at removing variability,

and secondly, as the number of wind farms in an averaging period declines, the trend becomes less certain due to outliers that influence averaging.

All LTC methods aim to adjust both low and high wind speed periods towards the trend line; however, since the trend line is unknown it becomes difficult to perform these adjustments accurately or even assess their accuracy (even when using long term wind speed or power output as a proxy). Additionally, in reality, a given wind speed is not guaranteed to always yield the same power output because there are many other factors that affect power output (e.g. wind direction, curtailment, etc.), which makes adjusting the data more difficult. This results in imperfect magnitudes of the correction and the possibility that variability effects not correlated with wind speed could persist. In general, the LTC process is less effective for the US and AUS compared to the UK because of the lower correlation between CF_{obs} and CF_{est} (seen in Figure 6.19), which leads to less accurate data corrections. The increased variability in the US is particularly significant after age 15, at which point the dataset is comprised of the older cohort of wind farms shown in Figure 6.23 that degrade at a faster pace than newer turbines. This could be causing increased variability as the performance decline tends to be more severe, which could lead to more downtime of wind turbines and wind farms or potential post-maintenance performance increases.

As the number of wind farms per averaging period decreases, the mean CF_{LTC} (of all wind farms) is more heavily affected by individual wind farms and fewer outliers are removed, causing potential divergence from the trend line. This is true for all countries, but exacerbated in the US where the relative number of wind farms above the age of 15 drops significantly. However, in the UK (there are insufficient data for AUS above age 15 to compare), variation increases but less severely than the US data. Especially above age 20, when a wind farm is approaching retirement or reconstruction, unexpected, end-of-life trends often appear, which are unable to be captured by LTC analysis because they are not related to the site wind speed. It is likely that the US wind farms are most affected by end-of-life trends, causing more variability in performance. Ages 20–25 have intentionally been included in

this analysis for a few reasons: wind farms are often still operating at this age; the trend still holds true despite more frequent outlier months; and it is important to show some limitations of the LTC process.

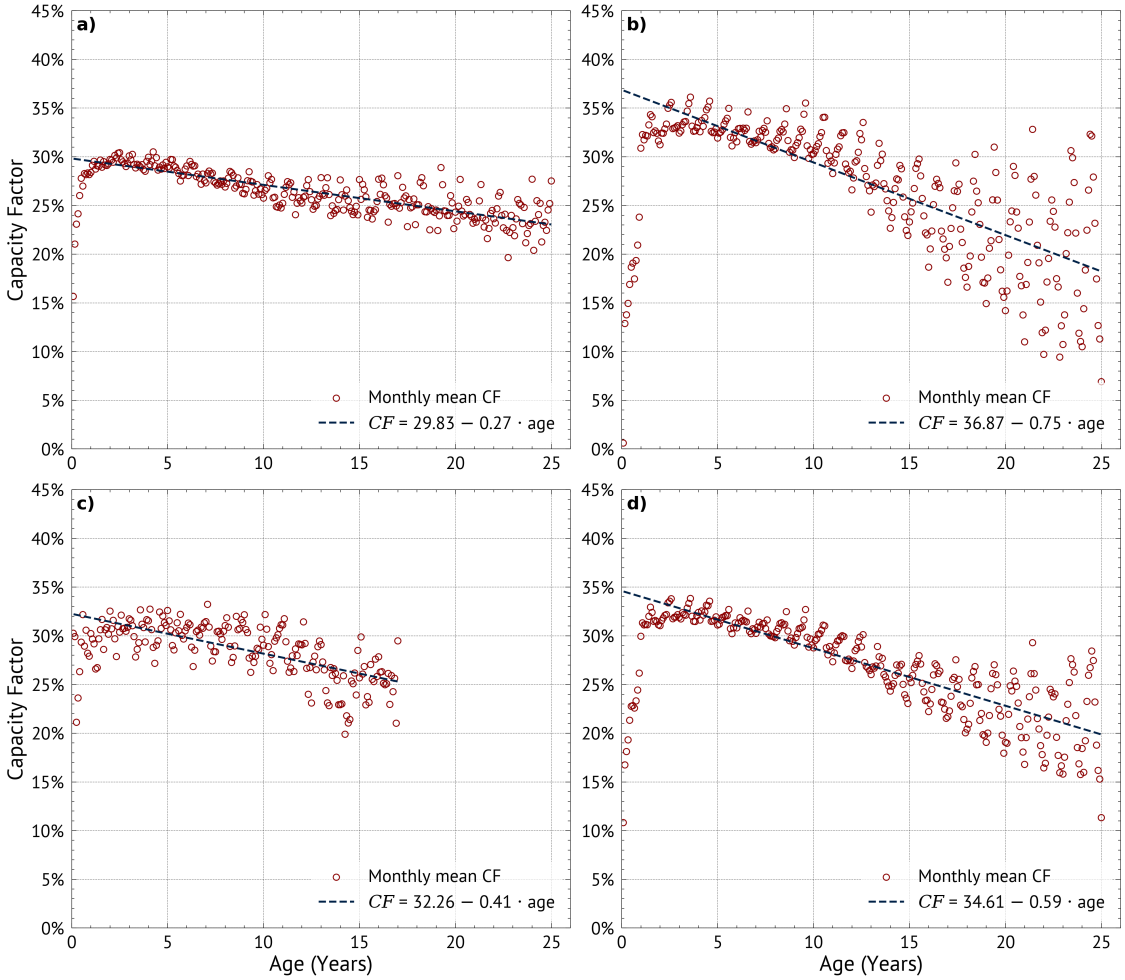


Figure 6.25: Wind farm performance plotted against age for the monthly averaged long-term corrected capacity factors (using Kelly LTC method and ERA5 reanalysis wind speeds) for wind farms in (a) the UK, (b) US, (c) AUS and (d) ALL. A linear regression line was fitted to the monthly averaged CF values for all wind farms.

As discussed in Section 6.3, presenting the results as an aggregated set could skew the degradation rate because there are larger amounts of data present for the earlier ages, and because it is presumed that newer turbines have higher CF due to technology [18], operation and design improvements. Consequently, turbines have been grouped into cohorts according to their age to examine their degradation rates. To do this, the CF_{LTC} for each wind farm is fitted using Equation 6.2.

The regression equation that includes seasonal fluctuations is used because small seasonal effects are still present in the long-term corrected data. The linear trend (regression coefficient β_1 in Equation 6.2) is then extracted as the degradation rate. Aggregating all degradation rates and plotting the data grouped by final turbine age results in Figure 6.26. Figure 6.26 shows only the results for CF_{LTC} produced using ERA5 reanalysis wind speed data. As in Figure 6.5, the result has been fitted with a Student's t distribution to quantify the results. Compared to the raw data, the mean moved towards zero (-0.044, -0.057, 0.048 and -0.048 for the UK, US, AUS and ALL) and in general the standard deviation decreased (0.915, 0.990, 0.372 and 0.967 for the UK, US, AUS and ALL). This indicates that the more extreme negative performance trends were driven by variability in the data, and once corrected using the Kelly method with ERA5 reanalysis wind data, the trends are less negative (except for AUS) with a tighter spread (except for the UK). There are also wind farms with age greater than 15 years that are shown to increase in performance, this likely corresponds with locations where the wind speed is increasing over time. For completeness, the results from the identical process for the CF_{obs} and CF_{LTC} calculated using MERRA-2 data are presented in Table 6.6.

Table 6.6: Comparison of the mean and standard deviation of the aggregated degradation rate histograms fit with a Student's t distribution.

	Raw data		**LTC (ERA5)**		**LTC (MERRA-2)**	
Country	μ	σ	μ	σ	μ	σ
UK	-0.207	0.728	-0.044	0.915	-0.081	1.052
US	-0.158	1.365	-0.057	0.990	-0.070	1.075
AUS	0.104	0.515	0.048	0.372	0.002	0.359
ALL	-0.168	1.094	-0.048	0.967	-0.072	1.137

The MERRA-2 data shows higher average degradation rates than the ERA5 LTC data; however, LTC data in general has degradation rates much closer to zero. What is clear is the shift of older wind farms to increasingly negative degradation values, which is exactly what is expected as wind farms age. This was not clearly

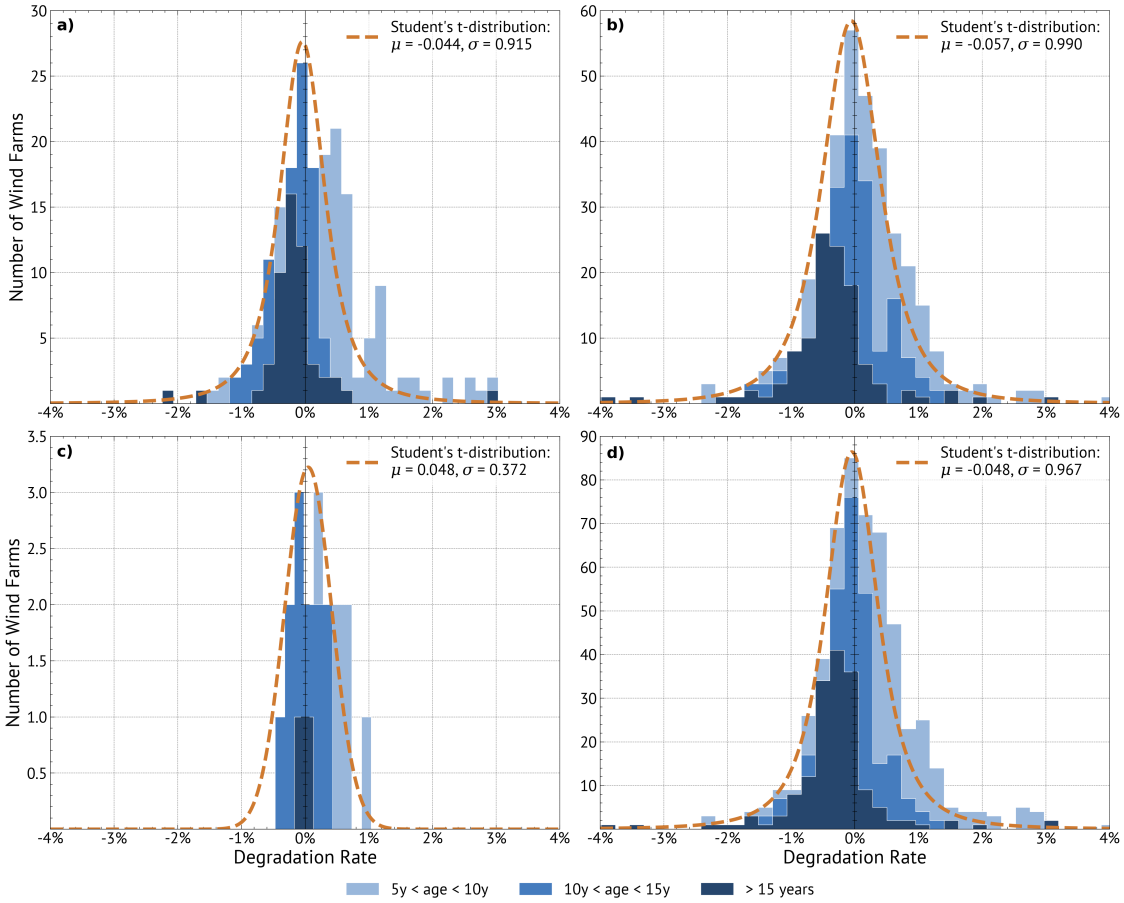


Figure 6.26: Distribution of degradation rates for individual wind farms using the CF_{LTC} data produced using the ERA5 reanalysis wind speed data for (a) the UK, (b) US, (c) AUS and (d) ALL.

shown using the raw data (Figure 6.5). To illustrate this point, the mean and standard deviation of the Student's t distributions fitted to the different age brackets are shown in Table 6.7. There are a couple key takeaways from these data. Firstly, the mean of the distributions decreases as the wind farms age for the LTC data, which is not seen in the raw data. Furthermore, the wind farms from age 5 to 10 years generally increase in performance and it is only after age 10 that performance decline is seen on average. Another takeaway is the standard deviation is lower for ages 10 to 15 than for ages greater than 15, indicating a more varied decline rate as wind farms approach decommissioning.

The lower average degradation rates for turbines under age 10 shown for the LTC data compared to raw data indicate that the aggregated numbers (Figure 6.22)

may be slightly exaggerated due to the averaging effect of having more data for wind farms at an early age. This conflates two possible questions when aggregating all data: is the degradation rate constant with time, and what is the degradation rate implied for a particular age, but for wind farms commissioned at different times? Generally speaking, the LTC method has been able to more realistically differentiate wind farm performance depending on age and has resulted in a more negative degradation rate of wind farms seen in the raw data. This indicates that there are more features than just the long term trend in wind speed that can impact the underlying performance of wind farms (as the long term average wind speed for AUS is slightly decreasing).

Table 6.7: Comparison of the mean and standard deviation of the aggregated degradation rate histograms fit with a Student's t distribution, broken down by age range.

		Raw data		**LTC (ERA5)**		**LTC (MERRA-2)**	
Country	**Wind Farm Age Range**	μ	σ	μ	σ	μ	σ
UK	5y < age < 10y	-0.346	0.817	0.179	0.812	0.055	0.844
UK	10y < age < 15y	-0.101	0.309	-0.173	0.347	-0.154	0.318
UK	age > 15 years	-0.143	0.449	-0.197	0.625	-0.179	0.574
US	5y < age < 10y	-0.127	1.800	0.070	1.152	0.037	1.166
US	10y < age < 15y	-0.104	0.589	-0.052	0.612	-0.061	0.634
US	age > 15 years	-0.311	0.788	-0.377	0.809	-0.327	0.810
AUS	5y < age < 10y	0.225	0.621	0.170	0.395	0.125	0.409
AUS	10y < age < 15y	-0.014	0.373	-0.077	0.320	-0.089	0.262
AUS	age > 15 years	0.075	0.160	-0.042	0.120	-0.049	0.163
ALL	5y < age < 10y	-0.185	1.062	0.106	1.003	0.045	1.042
ALL	10y < age < 15y	-0.101	0.527	-0.093	0.549	-0.097	0.577
ALL	age > 15 years	-0.235	0.697	-0.293	0.756	-0.258	0.741

6.6.5 Capacity-weighted results

The purpose of this research is to define degradation rates of wind farms that can then be used as an input to LCOE models and energy forecasting. Focusing on energy production, the most important wind farms in meeting electricity demand

are the largest wind farms. Consequently, a wind farm capacity-weighted method has been applied to the results using Equation 6.9. This provides insight into whether the largest wind farms are driving degradation or not.

Figure 6.27 shows the equivalent plot to Figure 6.22, however using the weighted average for the yearly means and for the percentile ranges. Compared to the non-weighted data, the regression trend has changed from -0.26 to -0.43 CF percentage points for the UK, -0.72 to -0.70 for the US, -0.43 to -0.33 for AUS, and -0.57 to -0.65 for ALL. The trend has significantly worsened for UK farms. This is especially important because the mean size of installed wind farms has been increasing over time (Figure 6.12) and thus the largest wind farms should be the newest wind farms, which would imply a less severe degradation rate (Figure 6.26).

Table 6.8 details the regression equations for the yearly averaged LTC data and the capacity-weighted yearly averaged LTC data. The biggest impact here is in the UK, where the values increase by 65–80% depending on which reanalysis model is used, indicating that the larger wind farms are degrading at a faster rate. The US degradation rates remain largely the same while the larger AUS wind farms degrade at a lesser rate, with the slope decreasing by 30–35%. These results are especially concerning for the UK as the largest wind farms are degrading at a higher rate.

Table 6.8: Linear regression equation comparison for yearly averaged LTC data and yearly weighted averaged LTC data. Yearly averaged data is labelled "Regular" and the weighted averaged data is labelled "Weighted". The reanalysis model used to produce the LTC data is given in parentheses.

Country	**Regular (MERRA-2)**	**Weighted (MERRA-2)**	**Regular (ERA5)**	**Weighted (ERA5)**
UK	$CF = 29.41 - 0.20 \cdot age$	$CF = 32.49 - 0.36 \cdot age$	$CF = 30.14 - 0.26 \cdot age$	$CF = 33.28 - 0.43 \cdot age$
US	$CF = 36.95 - 0.69 \cdot age$	$CF = 37.06 - 0.65 \cdot age$	$CF = 37.22 - 0.72 \cdot age$	$CF = 37.44 - 0.70 \cdot age$
AUS	$CF = 33.22 - 0.43 \cdot age$	$CF = 32.05 - 0.32 \cdot age$	$CF = 32.60 - 0.43 \cdot age$	$CF = 31.40 - 0.33 \cdot age$
ALL	$CF = 34.48 - 0.52 \cdot age$	$CF = 36.19 - 0.60 \cdot age$	$CF = 34.91 - 0.57 \cdot age$	$CF = 36.67 - 0.65 \cdot age$

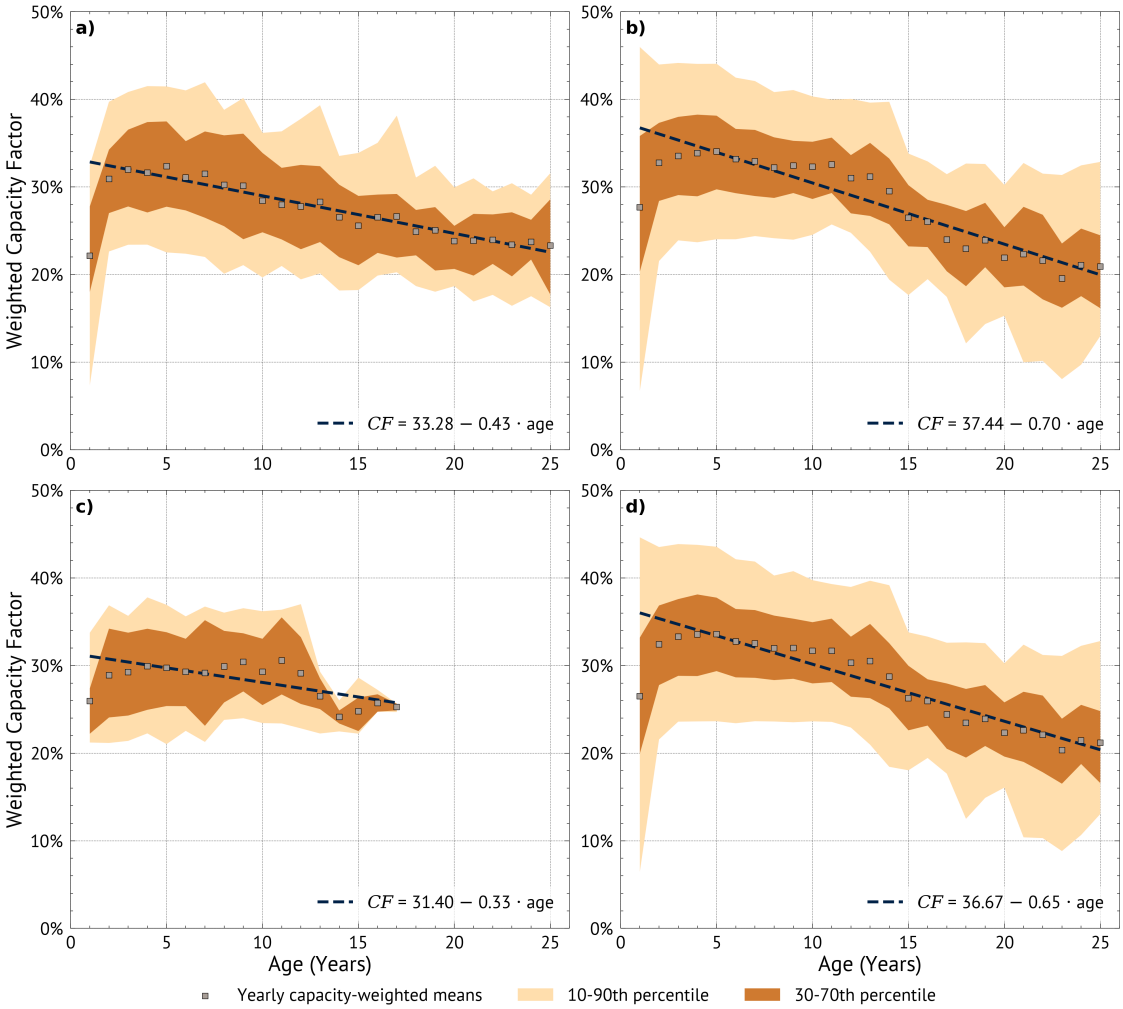


Figure 6.27: Wind farm performance plotted against age for the long-term corrected capacity factors (using the Kelly method and ERA5 reanalysis wind speeds) for wind farms in (a) the UK, (b) US, (c) AUS and (d) ALL. A linear regression line was fitted to the wind farm capacity weighted-average yearly CF values across all wind farms.

Figure 6.28 breaks down the distribution of wind farm sizes for the UK, US and AUS in 10 MW increments. The UK has a larger number of small wind farms, meaning the weighted average will be heavily impacted by a smaller number of large wind farms. Despite this, these wind farms are highly significant, and the change in their performance decline must be unpacked further, especially with the prevalence of offshore wind farms amongst the large farms in the dataset.

The UK dataset includes 19 offshore wind farms, with an average length of time series of 10 years (the longest time series is 17 years) and an average rated

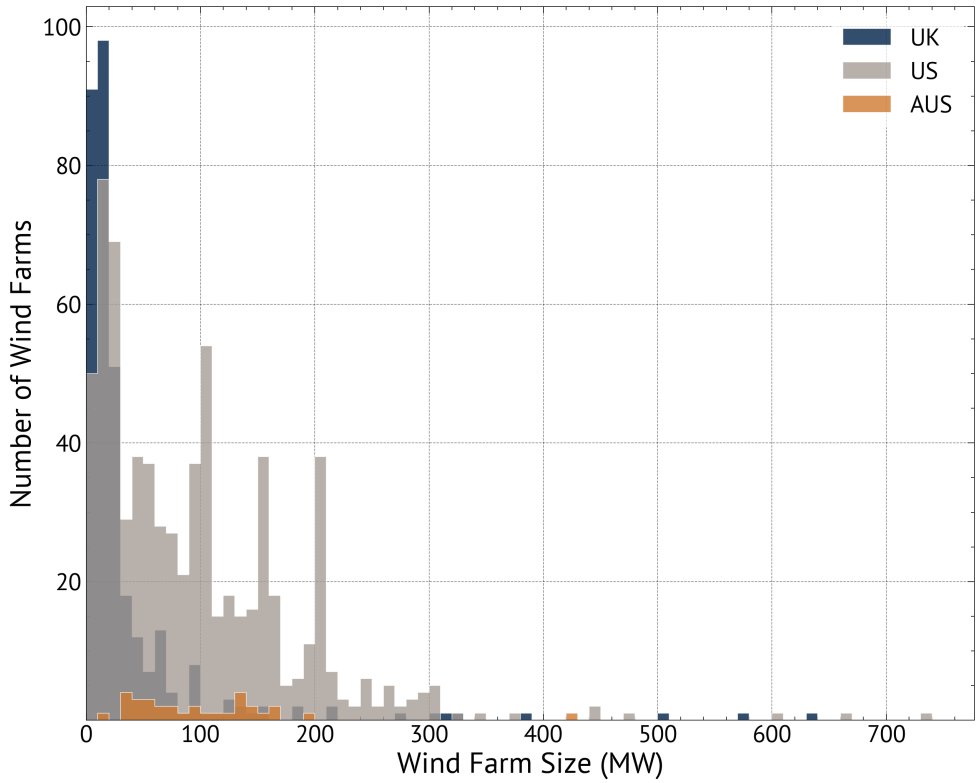


Figure 6.28: The distribution of wind farm rated capacities in 10 MW increments by country (UK, US, and AUS).

capacity of 214 MW (i.e. the majority of large wind farms in Figure 6.28). To understand the difference between onshore and offshore trends, Figure 6.29 shows the onshore/offshore split for monthly averaged UK wind farm CF_{LTC}, and wind farm capacity weighted-averaged UK wind farm CF_{LTC} using the Kelly LTC method and ERA5 data. The weighted-averaged data in plot (b) was calculated separately for the offshore and onshore datasets. What this figure shows is the degradation rate for both onshore and offshore farms increases when the data are capacity weighted, although not to the magnitude shown in Figure 6.27 and Table 6.8. It is the approximately 10% higher average CF for offshore wind farms combined with the predominant amounts of data existing for ages less than 15 that, when more heavily weighted, tilt the performance decline trend to be more negative. This underscores a limitation of aggregating all data together, albeit exaggerated for capacity-weighted data. This would be less noticeable if the offshore dataset extended out to age 25 as the increase in CF would be spread across all ages. However, there is a

clear limitation of combining the onshore and offshore datasets for this analysis, which results in a different performance decline rate. The fundamental driver of this difference is the increase in starting CF_{obs} of the wind farms, which skews the averaging calculation across all wind farms.

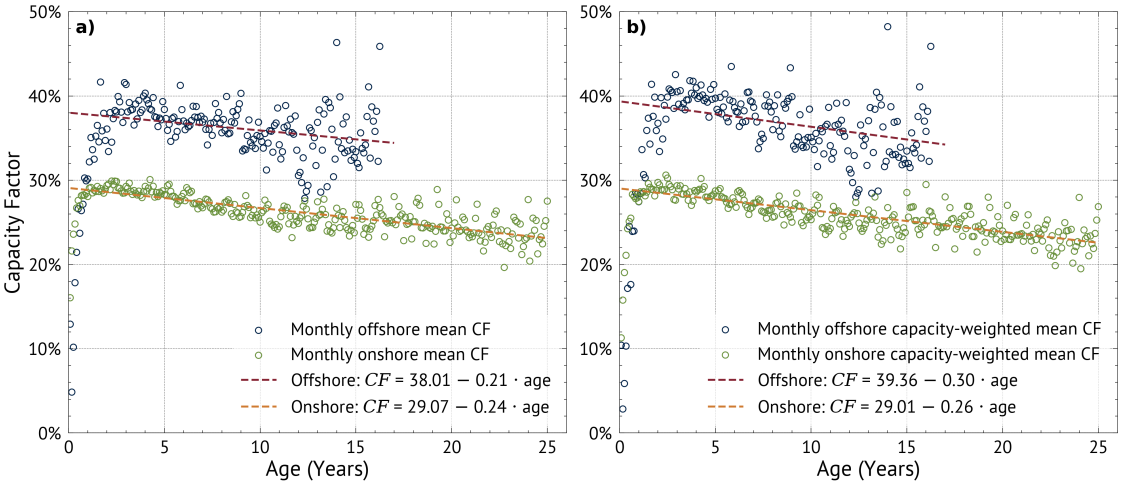


Figure 6.29: Wind farm performance plotted against age for (a) the monthly averaged long-term corrected capacity factors, and (b) monthly wind farm capacity weighted-averaged long term corrected capacity factors (using Kelly LTC method and ERA5 reanalysis wind speeds) for wind farms in the UK. Linear regression lines were fitted to the monthly averaged CF values for (a) all wind farms, and (b) to wind farm capacity weighted-averaged CF values for all wind farms.

6.7 Discussion

The results presented here show clear degradation trends for wind farms in the UK, US and AUS, for both raw data and LTC data. Degradation rates were shown to be larger for long-term corrected data, in the range of 0.03–0.1 CF percentage points more severe performance decline for all countries using the Kelly LTC method and ERA5 data. However, it has been shown that the assumption of a linear degradation trend is questionable becayse degradation rates were not constant; with initial increases (over first 5 or 10 years) prior to commencement of more rapid degradation rates, for some. This indicates that wind speed trends (increasing for the UK and US, decreasing for AUS) had a small but observable effect on the performance trend of wind farms. However, by correcting the observed data to

account for wind speed variation, (robust) estimates of performance were obtained, and the performance decline rate outputs are repeatable.

The estimated performance decline rates varied slightly depending on which reanalysis model was used for wind speed data and whether capacity factors/LTC values were averaged on yearly or monthly values. Table 6.9 provides a summary of the performance decline rates under all scenarios and the percentage per year values (calculated from the average CF) as well.

Table 6.9: Summary of performance degradation rates under all scenarios. CF percentage points per year are given and the percentage per year decline value (calculated from the average CF) is in parentheses. "Weighted" scenarios refer to wind farm capacity-weighted average.

Scenario	**Averaging period**	**UK**	**US**	**AUS**	**ALL**
Raw data	Annual	0.23 (0.82)	0.64 (2.07)	0.36 (1.08)	0.49 (1.63)
Raw data	Annual (weighted)	0.37 (1.31)	0.62 (2.00)	0.26 (0.78)	0.57 (1.89)
Raw data	Monthly	0.23 (0.82)	0.65 (2.10)	0.35 (1.05)	0.52 (1.72)
Raw data	Monthly (weighted)	0.39 (1.39)	0.64 (2.07)	0.24 (0.72)	0.60 (1.99)
LTC (MERRA-2)	Annual	0.20 (0.71)	0.69 (2.23)	0.43 (1.29)	0.52 (1.72)
LTC (MERRA-2)	Annual (weighted)	0.36 (1.28)	0.65 (2.10)	0.32 (0.96)	0.60 (1.99)
LTC (MERRA-2)	Monthly	0.22 (0.78)	0.71 (2.30)	0.41 (1.23)	0.55 (1.82)
LTC (MERRA-2)	Monthly (weighted)	0.39 (1.39)	0.68 (2.20)	0.31 (0.93)	0.63 (2.09)
LTC (ERA5)	Annual	0.26 (0.92)	0.72 (2.33)	0.43 (1.29)	0.57 (1.89)
LTC (ERA5)	Annual (weighted)	0.43 (1.53)	0.70 (2.26)	0.33 (0.99)	0.65 (2.16)
LTC (ERA5)	Monthly	0.27 (0.96)	0.75 (2.42)	0.41 (1.23)	0.59 (1.96)
LTC (ERA5)	Monthly (weighted)	0.44 (1.56)	0.73 (2.36)	0.30 (0.90)	0.68 (2.26)

The LTC process presented provided an improved way to correct CF data and extract underlying performance trends; however, some limitations to the LTC process persisted:

- Increased variability of data above age 15 (perhaps due to removal/non-operation of individual turbines);
- Different decline rates among different populations; and
- Statistical difficulties with averaging data of different sample sizes.

Increased variability of data above the age 15

Figure 6.25 shows the increased in variability of the CF_{LTC} data above age 15. This is likely because of effects other than wind speed variability that are captured in the data (e.g. increased downtime). Thus, it is recommended that to more accurately complete a LTC process, additional data need to be recorded or made publicly available. These data are curtailment and downtime (including partial downtime where a fraction of the turbines are not operating). With these added data points, other causes of variability can be recognised and adjusted for, or omitted accordingly.

Different populations have different decline rates

Figure 6.23 shows that US wind farms that were commissioned multiple years apart exhibited different performance decline rates. When CF_{LTC} was then averaged together over age, these differences caused slightly ambiguous performance trends, as seen by the linear regression fit in Figure 6.22 (b). However, when split by cohort age, the underlying trends were different to the entire US dataset and more closely resembled a linear trend. Splitting the dataset up into defined cohorts should therefore be an important consideration as wind farm datasets mature and this process is repeated.

Statistical difficulties with averaging data of different sample sizes

A problem encountered throughout this analysis was the prevalence of data for wind farms with age less than 10 (see the third column of Figure 5.8). On average, the younger wind farms have a larger mean CF (see Figure 6.3), resulting in a higher CF_{LTC} value when averaging by age (i.e. in Figure 6.22). This then results in an exaggerated degradation rate because of the higher CF_{LTC} for younger wind farms.

This is most clearly demonstrated when offshore wind farms are removed from the dataset in Figure 6.29. Even though they make up only a small number of the wind farms (19 out of 325 UK wind farms), they have an approximately 9% higher CF_{LTC} on average. This results in degradation rates of 0.21 and 0.24 CF percentage points per year for offshore and onshore wind farms respectively when using ERA5 data, compared to 0.27 CF percentage points in Figure 6.25 (a).

This limitation is a function of a wind farm fleet that is continuing to grow at an increased rate and will likely exist if the study is repeated in the near future. It should therefore be acknowledged when averaging CF data by age.

6.7.1 Application

This method has been created to be replicable across multiple countries where wind farm power output data are available. The LTC method has been designed to be effective for data of multiple time resolutions, which is a big improvement on previous methods. The framework consists of the following elements:

1. Choose a reanalysis model to use. This study has used both MERRA-2 and ERA5 models and concluded that ERA5 produced a higher correlation of wind farm power output with recorded values. Any reanalysis model can be used; however, its accuracy will affect the results.

2. Access wind farm generation data and wind farm metadata (location, turbine size, turbine type, number of turbines, etc.).

3. Co-locate reanalysis wind speeds with wind farm location.

4. Average reanalysis wind speed data and wind farm generation data to lowest common time step.

5. Apply LTC method from Section 6.4 to dataset,

6. Fit regression to results and analyse results.

6.8 Conclusions

A robust method has been created to assess the performance trends of wind farms over their lifetime. It has been created to serve as a framework that could be applicable to any country where wind farm power data and reanalysis wind speed are available.

The method is designed to account for the inaccuracies of reanalysis datasets by including a rolling correlation between estimated power output and observed power output, whilst also correcting wind speed variability-affected power outputs. The LTC method has incorporated an algorithmic approach that allows for more careful adjustment of wind farm power output and improves on previous methods.

Wind farm performance declines have been estimated for the UK, US and AUS – 0.26, 0.72 and 0.43 CF percentage points per year respectively using annually averaged long-term corrected CF data. Although the US wind farms exhibit the highest performance decline rate, the average predicted CF from ages 2 to 25 is 27.5% compared to 26.6% and 26.8% for the UK and AUS. This is because the CF for US wind farms at early ages is 5 to 7% higher.

When weighting the performance decline by wind farm power capacity, the degradation rates are 0.43, 0.70 and 0.33 for the UK, US and AUS. These results were replicable across both the MERRA-2 and ERA5 reanalysis datasets (ERA5 presented as the final degradation rates) and both monthly and yearly averaging periods. The weighted results show that large UK wind farms are declining at a faster rate than smaller farms; however, the country-wide averaged results were shown to be exaggerated due to the higher CF (approximately 10% higher on average) and limited data above the age of 15 for offshore wind farms.

The novel LTC method created in this study was able to successfully remove seasonality from individual wind farms, though there were issues aggregating the data of US wind farms above age 15. Previous studies that have used a LTC method have estimated power output from a 'point' wind speed (from reanalysis data). This study instead assumes a wind distribution around the reanalysis wind speed value to more accurately represent the time period being simulated (i.e. one hour or one

month). Furthermore, previous methods have assumed that the reanalysis model is always accurate; however, this thesis accepts that reanalysis wind speeds are not perfectly precise. This has resulted in an algorithmic approach for correcting wind farm power output data based on correlation between predicted power output and actual power output to help avoid correction errors.

This method has been built to handle multiple time resolutions across multiple countries. Previous models fail to accurately correct data at low time resolutions as low power output predictions cause an over exaggerated correction. Previously, LTC studies have been completed for the UK, US, Germany and Denmark, however this is the first study to compare multiple countries using the same method.

This method of analysing wind farm performance over time has been created to serve as a framework for future studies. It may also become increasingly important to consider the effect of curtailment and wind farm management methods; however, this will depend on available data and region specific grid constraints.

What you do makes a difference, and you have to decide what kind of difference you want to make.

— Dr. Jane Goodall

7 Wind Farm Performance Prediction

Contents

7.1 Machine learning model selection 173
7.1.1 XGBoost 177
7.1.2 Random Forest 178
7.2 Model implementation 178
7.3 Results 181
7.3.1 Feature importance 185
7.3.2 Data limitations 188
7.3.3 Modelling results within a country of the dataset 191
7.4 Discussion and conclusions 193

Reanalysis data have previously been used to examine capacity factors to estimate wind resources and potential future wind farm sites. Lu et al. [3] estimated the global potential of wind energy, and Gunturu and co-authors characterised the wind resources in the US [164], Europe [165] and Australia [166]. Becker et al. [167] simulated 100%-renewable wind and solar mixes in the US, and Huber [168] simulated wind and solar outputs across Europe. Staffell and Pfenninger [140] criticise all of these studies due to their lack of effort to validate wind speed data by comparing their results to observed output data.

Staffell and Pfenninger used bias-corrected reanalysis to simulate current wind power output for the UK and Europe. They bias-corrected different areas based on

projected capacity factors using MERRA-2 wind speeds and observed values. They emphasise how reanalysis models need to be validated in the area of use in order to accurately predict wind farm CF, which many studies fail to do. To bias-correct the data, the difference between the simulated wind farm power (using their own wind farm power curve) and the observed power output was calculated, and the reanalysis wind speed was then adjusted to minimise this difference. They then used bias-corrected wind speeds to estimate wind farm power output in a given location. This method assumed that their simulated wind farm power curve is accurate, which may not be true. However, undergoing a validation process is a worthwhile step, even if by using a different validation method.

A limitation of the Staffell and Pfenninger method is the reliance on an accurate wind farm power model, which was also a key limitation of previous LTC methods (as discussed in Chapter 6). Wind farm power prediction is a highly complex, non-linear problem in which outcomes depend on various factors, such as wind speed, turbine size, number of turbines, location, etc. This requires a technique that goes beyond using reanalysis wind speed data (bias corrected or not) and incorporates as much relevant data as possible. A supervised machine learning (ML) model – which aims to forecast/classify a specific outcome of interest [169] – is well suited to the task due to the large number of known inputs and the non-linear trends that need to be captured. A ML model is an expanded regression that enables complex, non-linear relationships to be captured between a wide array of inputs (features), rather than being limited to particular statistic fits.

Throughout this thesis, an extensive wind farm metadata and power output dataset has been built, which is used in this chapter to train machine learning models that can predict wind farm power output. Specifically, this chapter utilises the dataset detailed in Chapter 5 to produce supervised extreme gradient boost (XGBoost) and Random Forest (RF) machine learning models that predict lifetime power output of operational or theoretical wind farms.

7.1 Machine learning model selection

Predicting power output of a wind farm is a highly complex, non-linear task that requires a highly trained, multi-input model. The LTC analysis in Chapter 6 captures some fundamental aspects of predicting power output by utilising wind speed combined with a smoothed power curve, and by quantifying the impact of age; however, limitations remain. Some key questions are:

1. What if the reanalysis wind speed is biased by location?
2. What is the error of reanalysis wind speed data at any given time?
3. Does turbine manufacturer impact performance on average?
4. How important is location for wind farm performance?
5. How does wind speed direction impact power output for each wind farm?
6. How do the size and number of turbines impact the power curve?
7. Does the averaging period of the time series data impact how to best model power output?

Data that capture these aspects must be included to get a better understanding and a more accurate model of wind farm power prediction. To do this, the metadata and wind farm output datasets were combined to produce a dataset with 16 features, which is then used to predict CF. It is important to use a model where the key factors (the features with the biggest influence on predicting CF) can be extracted and understood.

Due to the combination of continuous and categorical variables, and the large number of input variables overall, a machine learning model is well suited to predict CF. Supervised learning is defined as an algorithm that generates a function that maps inputs to specified desired outputs [170]. Supervised machine learning models also provide a computationally efficient way to capture tens or hundreds of features

in a model. Therefore, a supervised machine learning model is desired because the inputs and output of the model are already defined.

A supervised machine learning model, in its simplest form, will identify trends in the input data to predict an output. Because the model will be 'learning' from wind speed trends and biases and thus may incorporate them into predicting CF, questions 1 and 2 listed above will not be directly answered. The model may account for these two questions, but their impact can not be explicitly extracted from the model. Questions 3 to 5 will be analysed by looking at how each of these inputs (turbine manufacturer, wind farm location and wind speed direction) affects the model. Question 6 will be examined by matching output results to wind farm characteristics (wind farm size and number of turbines) and question 7 will be answered by evaluating the accuracy of the model when using different time resolution training data.

Many different supervised machine learning models exist and can broadly be separated into two categories: regression and classification. Classification models are used to separate data into distinct categories, e.g. classifying emails as "junk" or "not junk". Models that are used for classification inlcude logistic regression [171], clustering (e.g. Adaptive Nearest Neighbours [172]), Naive Bayesian [173], Decision Tree Classification [174], Support Vector Machine (SVM) [175], One- and Zero-R [176]. However, classification is not suited to this task, as a continuous variable (CF) is being predicted. In contrast, regression models are used to predict continuous variables and are therefore best-suited to this problem.

Choosing a regression model

Various regression models were considered for this chapter, and multiple were investigated. This section will detail existing regression models and show why the selected models were promoted.

Common supervised machine learning regression models are linear regression, decision tree regression, RF regression [177], and least absolute shrinkage and selection operator (LASSO) regression [178]. Due to the size of the feature set and

non-linear nature of some features (e.g. manufacturer), it was concluded that a multivariate linear regression was unsuitable. Furthermore, LASSO regression is commonly used for variable selection or datasets with high multicollinearity, which is not necessary here. Thus, a decision-tree based method was favoured.

Tree-structured regression is an alternative approach that is not based on assumptions of normality or user-specified model statements, unlike some older methods such as ordinary least squares regression (OLS) [179]. Decision trees therefore allow for the implementation of continuous, discrete and categorical features, such as those that occur in the dataset.

A single decision tree used across an entire dataset will poorly predict the output variable, and is hence known as a 'weak learner'. However, the performance of a set of weak classifiers is generally more accurate than a single classifier given the same quantity of training information [180]. Consequently, this chapter will focus on ensemble learning applications of decision trees, which are methods that generate many classifiers and combine their results [181]. Common ensemble models are bootstrap aggregation (bagging) [182], RF, and boosting [183].

Bagging is a method that generates multiple versions of a predictor (in this case, a decision tree) and uses these to get an aggregated predictor [182]. The process is relatively simple: multiple versions are created by drawing many bootstrap samples[1] from the available data, running a prediction method, and aggregating over all predictions (when predicting a numerical value).

Random Forest is an extension of bagging and also uses randomly selected subsets of features for each sample. RF is an ensemble model that combines randomly selected decision trees that have a subset of attributes. The outputs of these decision trees are combined to reach a single result (in this case, CF).

Boosting is an iterative procedure that combines weak learners sequentially so that each new tree corrects the errors of the previous one. Initially, instances (each tree and its branches) are created with equal weights, but as the process evolves,

[1] A bootstrap sample is made by randomly picking data from the original dataset, where a data point is chosen and returned to the dataset before the next data point is chosen, until the bootstrap sample is the same length as the original dataset (i.e. it can include duplicates).

the weights are modified depending on errors (delta between predicted CF and observed CF). After the error is computed, weights are adjusted down for more accurate instances (and adjusted up for less accurate ones), so that the model puts more weight on improving instances with poor predictions. The importance of each individual learner (decision tree) is weighted by its performance [184] (ability to predict CF). Extreme Gradient Boosting (XGBoost[2]) [185], which is described in Section 7.1.1 below, is a recent implementation of boosting which has been widely recognised in Kaggle machine learning competitions due to its advantages of high efficiency and flexibility [186]. For example, 17 of the 29 Kaggle challenge winners in 2015 and all top-10 finishers in the KDD Cup 2015 used XGBoost [185].

Zhang et al. [186] compared XGBoost and RF algorithms with three baseline machine learning models (SVM, multilayer perceptron (MLP), multivariate adaptive regression splines (MARS)) to predict the relationship between undrained shear stress and various basic soil parameters. XGBoost and RF had superior predictive performance compared to the other models, with XGBoost being slightly more accurate than the RF implementation (approximately 0.5% improvement in mean absolute percentage error (MAPE)). Zhang et al. also highlight that an added benefit of XGBoost is the ability to have insights into feature importance (i.e. gain criterion). The gain is calculated by aggregating each feature's contribution to the model across all trees, with a higher value indicating greater importance. However, feature importance can also be extracted from RF algorithms through a similar mechanism.

Another consideration is how the models handle missing values. XGBoost supports missing values by default, which means branch directions for missing values are learned throughout training. RFs, on the other hand, require values to be filled (e.g. by averaging over the feature) or removed before the model can be run. Nevertheless, due to the similarity of these algorithms and comparable previous performance on similar problems, they will both be investigated for this chapter.

[2]https://github.com/dmlc/xgboost

7.1.1 XGBoost

XGBoost is an implementation of gradient boosting decision trees that is designed for scalability, speed and performance [187]. While the performance of the model cannot be defined by one tree (because it is working to correct the residuals of the previous tree in the sequence), it is useful to look at a decision tree to visually understand the process. Figure 7.1 shows an example decision tree from a reduced feature set model (the full model decision trees are much larger, ranging in depth from 6 to 9 compared to the 4 shown and with more features included). From this tree, key factors used for prediction are the hub height wind speed, the age of the wind farm, and the turbine capacity within the wind farm. The leaf nodes in this figure indicate the margin value that the leaf may contribute to prediction.

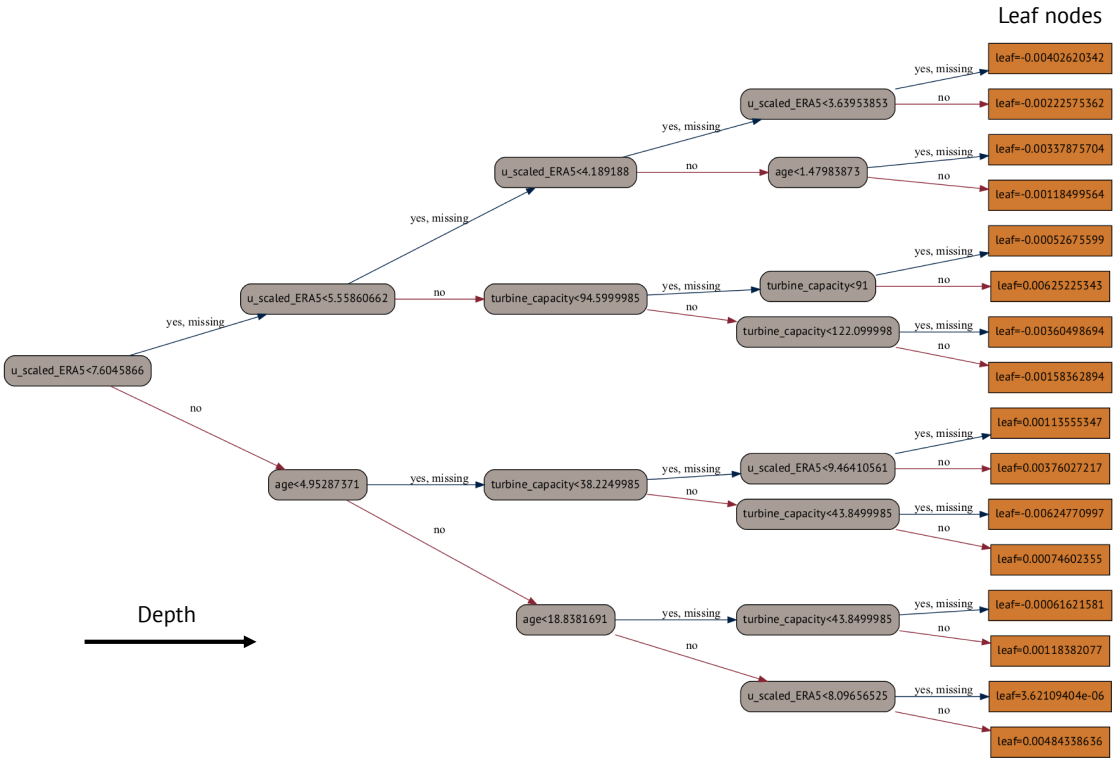


Figure 7.1: An example decision tree for predicting CF produced using XGBoost and a reduced feature list.

7.1.2 Random Forest

RF is highly applicable to a wide range of prediction problems and has few parameters to tune [188]. RFs are suitable for large and small datasets and have been previously been used in many applications, such as in predicting a compound's biological activity [189], ecological prediction [190] and econometrics [191]. RF algorithms use similar trees to XGBoost (Figure 7.1); however, they are created using random selections of features and are not informed by previous decision trees. In this thesis, the *scikit-learn* Python implementation [192] for RF regression is used, and missing values were removed due to them only accounting for less than 1% of the dataset.

7.2 Model implementation

To ensure the models were capturing as many important features as possible, the models had 19 inputs, detailed in Table 7.1. Note the manufacturer, time resolution and location-type inputs are one-hot encoded[3], meaning extra features are created for each unique value (i.e. manufacturer becomes multiple columns: 'manufacturer_Vestas', 'manufacturer_Siemens', 'manufacturer_Enercon', etc.). Furthermore, all ERA5 and MERRA-2 specific columns are used for their respective models, not combined. The output of the model is always CF_{pred}.

XGBoost and RFs have multiple parameters than can be tuned to improve the accuracy of the model, such as the depth of individual decision trees and the number of decision trees (instances). To ensure the most accurate model is being used, a k-fold cross validation (CV) grid search algorithm [193] is used across all the parameters to find the most accurate model. k-fold CV is a resampling procedure that randomly splits the dataset into k groups, or folds; the first fold is treated as a validation set, and the model is fit on the remaining $k-1$ sets [194]. This

[3] A method of converting a column of categorical values into unique columns for each value containing zeros and ones.

Table 7.1: Input features to the CF prediction model. The features that have a data type of categorical are one hot encoded and thus turns into multiple columns (i.e. location_type turns into two columns 'location_type_Onshore' and 'location_type_Offshore'), which contain 0 if it is not in that country and 1 otherwise). Note that ERA5 and MERRA-2 wind speed magnitude, direction and resource data were used in separate models.

Feature name	Description	Data type
age	Wind farm age	Continuous
u_scaled_ERA5	Hub height wind speed magnitude from ERA5	Continuous
angle_u_scaled_ERA5	Hub height wind speed angle from ERA5	Continuous
u_scaled_MERRA2	Hub height wind speed magnitude from MERRA-2	Continuous
angle_u_scaled_ MERRA2	Hub height wind speed angle from MERRA-2	Continuous
turbine_capacity	Individual wind turbine capacity (MW)	Continuous
number_turbines	Number of turbines	Discrete
hub_height	Hub height of turbines	Continuous
blade_diameter	Diameter of the turbine blades	Continuous
lat	Latitude of the wind farm	Continuous
lon	Longitude of the wind farm	Continuous
c_ERA5	Weibull scale factor of the wind distribution from ERA5	Continuous
k_ERA5	Weibull shape factor of the wind distribution from ERA5	Continuous
c_MERRA2	Weibull scale factor of the wind distribution from MERRA-2	Continuous
k_MERRA2	Weibull shape factor of the wind distribution from MERRA-2	Continuous
mean_ws_ERA5	Lifetime mean wind speed from ERA5	Continuous
mean_ws_MERRA2	Lifetime mean wind speed from MERRA-2	Continuous
start_year	Commencement year of wind farm	Discrete
manufacturer	Manufacturer of wind farm	Categorical
time_resolution	Time resolution of the time series data	Categorical
location_type	Whether the wind farm is located onshore or offshore	Categorical
CF_est_ERA5	CF_{est} calculated in Chapter 6 using the ERA5 reanalysis data	Continuous
CF_est_MERRA2	CF_{est} calculated in Chapter 6 using the MERRA-2 reanalysis data	Continuous
mean_CF_est_ERA5	The mean of the CF_{est} time series, calculated using the ERA5 reanalysis data	Continuous
mean_CF_est_ MERRA2	The mean of the CF_{est} time series, calculated using the MERRA-2 reanalysis data	Continuous
day_of_year	Day of the year	Continuous

algorithm sweeps over a provided set of model hyperparameters and returns the most effective set. For this dataset, a k value of 10 was used, as suggested by Kohavi [193] and Gareth et al. [194].

When implementing the model, the entire dataset was combined, shuffled and split into training, testing and validation sets, which were allocated to be 60%, 20% and 20% of the data respectively. The entire dataset refers to CF_{obs} data at every time and age for every wind farm, combined with the metadata for each wind farm. The entire dataset will only change when the model is run using data from the UK, US, AUS or ALL. The model uses the training and testing set to iteratively improve. The validation set remains unseen by the model and is used to conduct a final test of accuracy and to ensure the model has not overfit the training data. The validation accuracy (how close the model can predict CF values in the validation dataset) ultimately reveals how effective the model is. This process is crucial because it is minimising the validation errors that ultimately indicate the strength of the model.

To ensure the model could predict 'external' data, the above process was repeated with a single wind farm removed from the initial dataset (before it was split into test, train, and validation). This is an implementation of the leave-one-out cross-validation (LOOCV) methodology, even though the dataset is not divided into equal subsets due to the length of each wind farm time series dataset. The model is trained on all wind farm data except for the removed wind farm. After the model has been trained and validation accuracy tested, an external validation occurs by predicting the CF for the initially removed wind farm. 1,053 models were run using this method (one for each wind farm in the dataset) and the external accuracy results were compiled. This method not only tests the accuracy of the model, but also acts as a use case for the model – predicting CF values for a wind farm within the countries in the dataset.

For all prediction models, the validation accuracy was measured using the pseudo-R-squared (pseudo-R^2) value for non-linear models defined in Schabenberger and Pierce [195], MAE and relative error (RE) for the predicted CF values compared against the observed CF values. RE is a metric that measures how the mean

absolute error changes as CF increases (i.e. an error of 5% is more significant for a CF value of 10% compared to a CF value of 90%), and is defined as:

$$RE = \frac{\frac{\sum_{i=1}^{N} |CF_{obs,i} - CF_{pred,i}|}{N}}{\frac{\sum_{i=1}^{N} CF_{obs,i}}{N}} = \frac{MAE}{\frac{\sum_{i=1}^{N} CF_{obs,i}}{N}} \tag{7.1}$$

where N is the number of samples in the validation dataset.

The pseudo-R^2 value is not suitable for non-linear regression models [196]; however, it is a common indicator used in machine learning model training and has been included for completeness. However, in all instances the MAE and RE metrics were used to determine the best model, not the pseudo-R^2. The feature importance can also be extracted, based on the gain of each feature. The gain represents the relative contribution of a given feature to the model by assessing the improvement in accuracy brought by the feature to its trees. Effectively, gain is measured by how much the accuracy improves as a result of the feature node being implemented.

7.3 Results

Table 7.2 shows the validation accuracy results in terms of pseudo-R^2, MAE and RE for XGBoost and RF models run using wind farm data from the United Kingdom (UK), United States (US), Australia (AUS) and all three countries combined (ALL). For a given model type, two separate models were run for each country: one using the data produced using the MERRA-2 model and the other using ERA5 data. In all cases, the country-specific model trained using ERA5 data outperformed the model trained using MERRA-2 data. The accuracy of the model also varies for all different datasets. There is a key difference between the UK and US data and the AUS and ALL data, which is due to AUS using exclusively hourly data in the dataset and ALL using a combination of hourly and monthly, which is inherently harder to predict due to a larger variance in outputs. The difference between the UK and US is consistent with the LTC results found in Chapter 6, which showed more inconsistent wind farm output data with respect to wind speed and therefore a more difficult dataset to long-term correct. For these datasets, XGBoost performs

approximately 0.7, 1.1, 0.1, and 0.3 better in absolute CF MAE than RF, and thus the remainder of the presented results will focus on the XGBoost outputs.

Table 7.2: Accuracy results from XGBoost and RF trained models to predict CF of wind farms, using data from wind farms in the UK, US, AUS and all three countries combined. The MAE is displayed in CF units, i.e. a MAE of 1 is a 1% CF MAE.

Country	ML model	Reanalysis Model	Pseudo-R^2	MAE	RE
UK	XGBoost	ERA5	0.893	2.576	0.094
UK	XGBoost	MERRA-2	0.890	2.688	0.099
UK	RF	ERA5	0.845	3.275	0.119
UK	RF	MERRA-2	0.837	3.413	0.124
US	XGBoost	ERA5	0.866	3.139	0.102
US	XGBoost	MERRA-2	0.864	3.183	0.104
US	RF	ERA5	0.780	4.224	0.138
US	RF	MERRA-2	0.748	4.550	0.149
AUS	XGBoost	ERA5	0.770	10.34	0.282
AUS	XGBoost	MERRA-2	0.716	11.66	0.318
AUS	RF	ERA5	0.761	10.48	0.286
AUS	RF	MERRA-2	0.701	11.89	0.325
ALL	XGBoost	ERA5	0.776	9.905	0.273
ALL	XGBoost	MERRA-2	0.725	11.13	0.307
ALL	RF	ERA5	0.763	10.19	0.281
ALL	RF	MERRA-2	0.704	11.52	0.317

Figure 7.2 shows the predicted CF against the validation (observed) CF for XGBoost models trained using wind farm data from the UK, US, AUS and ALL. The one-to-one relationship line is also plotted and the ALL plot has been normalised (equal number of CF_{pred} data points for each country) and colour-coded to show the predictions for the respective countries. This plot shows the differing accuracy between the models trained with different datasets, with the exclusively monthly datasets having a much higher accuracy due to the lower variance wind speed data.

To better understand the impact of data with higher variance (hourly data), the data from AUS wind farms was also used to train models with different time resolutions. The purpose is to show that the model can still train at a high accuracy on AUS wind farm data, and it is the higher time resolution that causes the variance

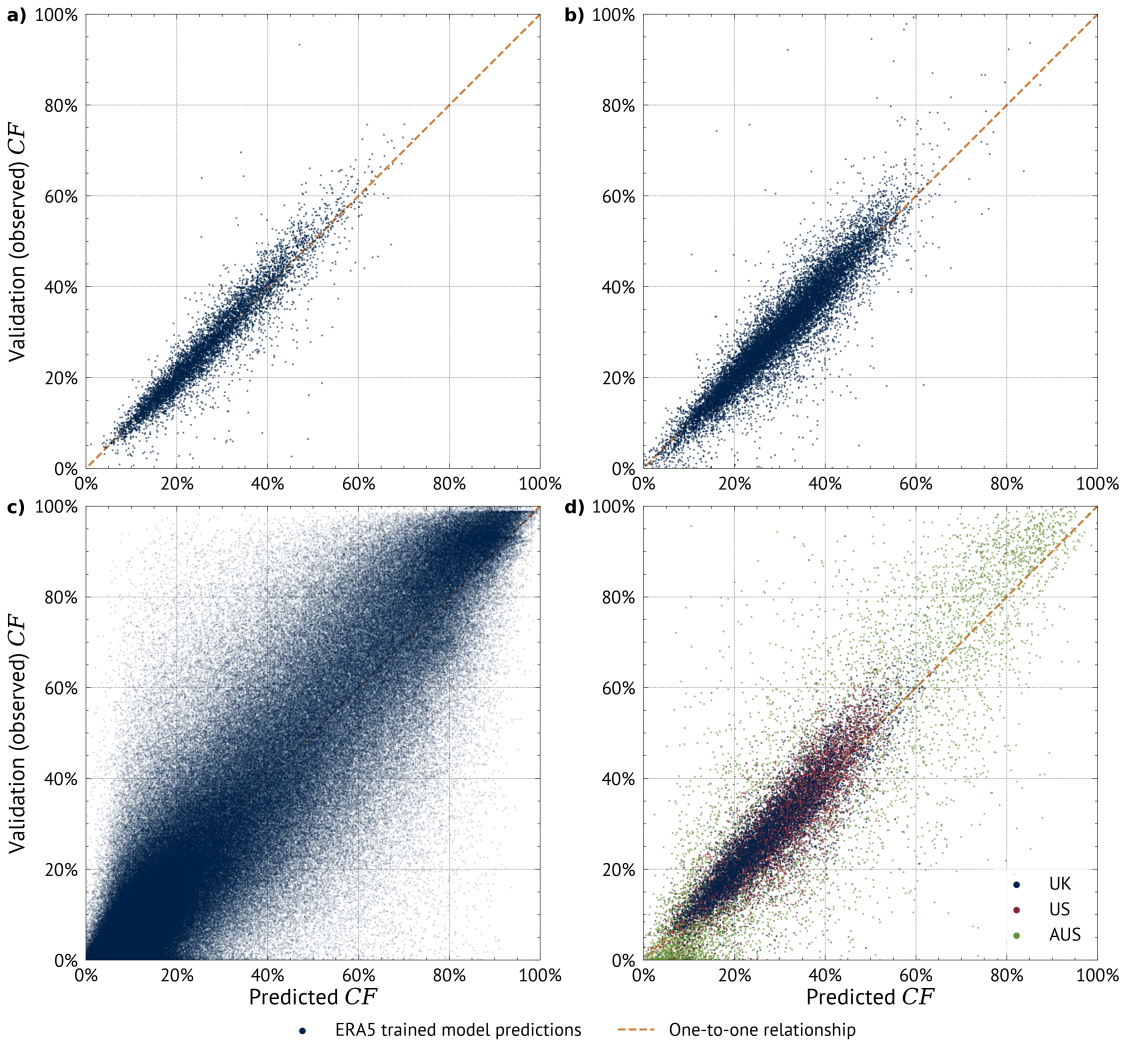


Figure 7.2: The predicted CF from the XGBoost models plotted against the validation (observed) CF values. The models were trained with data from the UK (a), US (b), AUS (c) and ALL (d). All plots show models trained using the ERA5 wind speed features and the ALL plot is displaying an equal number of CF_{pred} data points for each country with the outputs colour coded for the UK (blue), US (red) and AUS (green).

seen in Figure 7.2 (c) and (d). Figure 7.3 shows models (with the same parameters) trained using different time resolution data – hourly (a), daily (b), weekly (c) and monthly (d) – for AUS wind farms. Table 7.3 details the error metrics from these models. When the data resolution changes from hourly to daily, weekly and monthly, CF MAE improves by 5.6, 7.0 and 8.0 respectively. This shows that the AUS model, when using data resolution that matches the UK and US models, displays comparable performance and slightly lower errors.

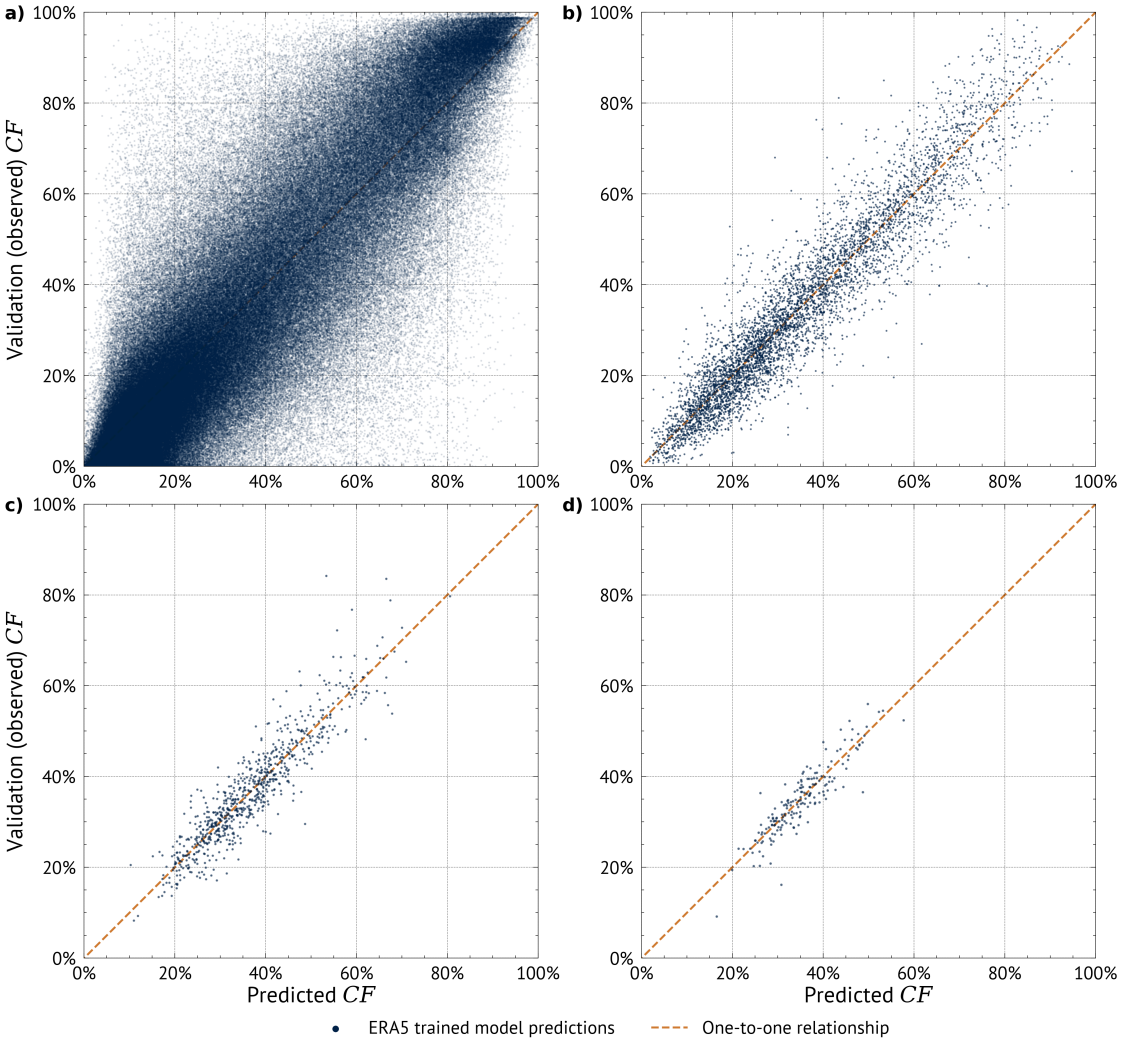


Figure 7.3: The predicted CF from the XGBoost models plotted against the validation (observed) CF values. The models were trained with data from wind farms in AUS and show models trained on hourly data (a), daily data (b), weekly data (c) and monthly data (d). All plots show models trained using the ERA5 wind speed features.

Table 7.3: Accuracy results from XGBoost trained models using data from AUS wind farms averaged to varying time resolution levels. The MAE is displayed in CF units, i.e. a MAE of 1 is a 1% CF MAE.

Data time resolution	**Pseudo-R^2**	**MAE**	**RE**
Hourly	0.770	10.34	0.282
Daily	0.901	4.741	0.133
Weekly	0.864	3.256	0.088
Monthly	0.843	2.286	0.066

7.3.1 Feature importance

An informative output from the XGBoost models is the importance of each feature to the model. This allows insights into which features most strongly affect CF_{pred} in the model. The model is ranking features based on trends in the data and the effects of each feature node on model accuracy. This also means an understanding of the key drivers of the model can be drawn from the results. Figure 7.4 shows the normalised feature importance (only the top features summing to 90% importance are shown) for models trained with data from the UK (a), US (b), AUS (c) and ALL (d).

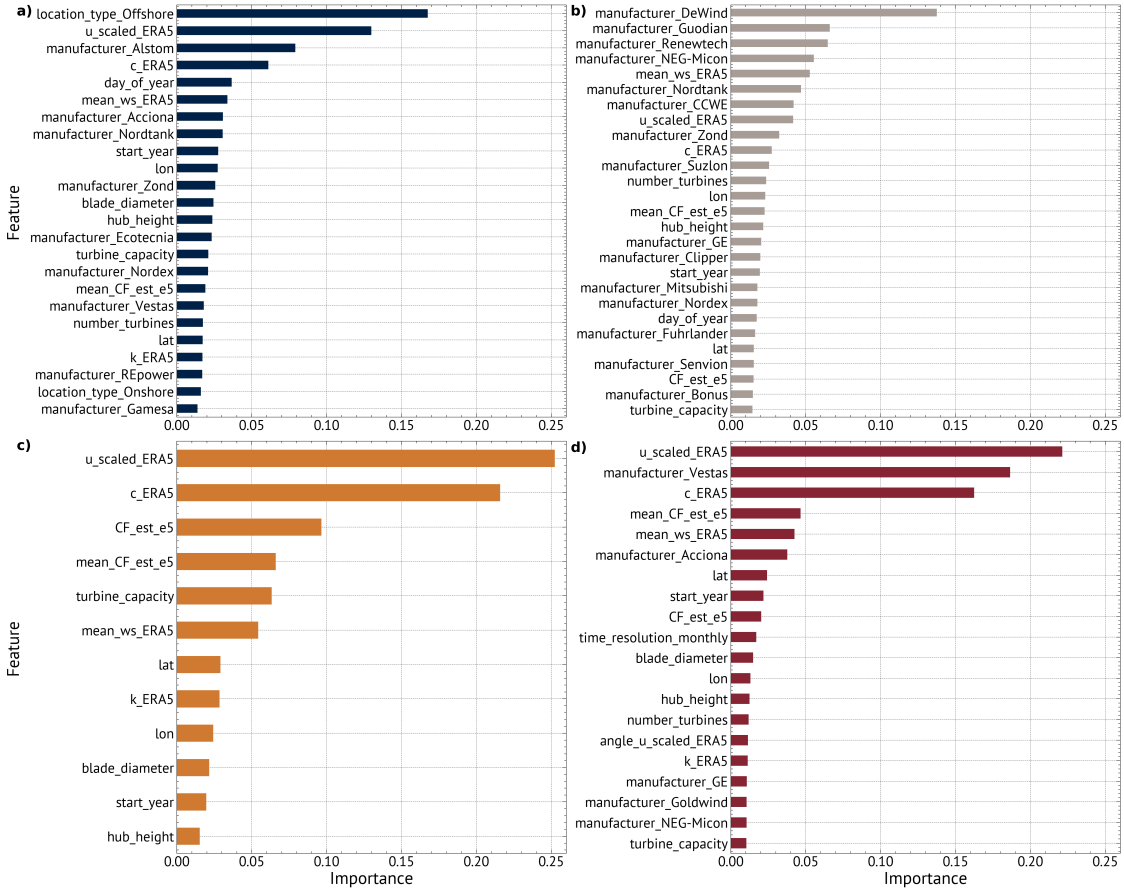


Figure 7.4: Feature importance from the XGBoost models trained with data from wind farms in the UK (a), US (b), AUS (c) and ALL (d). All plots show models trained using the ERA5 wind speed features and feature importance values adding up to 0.9.

Figure 7.4 highlights the differences between models and countries. It also provides an opportunity to understand the important drivers of the prediction model. Across all models, the hub height wind speed features are dominant,

which is expected; however, they are not always the most important features, which is unexpected.

In the UK model, the location type (onshore or offshore) is the most important feature. This is because the mean lifetime CF of offshore wind farms is 34.9% (19 wind farms) compared to 27.1% onshore (306 wind farms). It would therefore make sense that this categorical variable acts as a step change when predicting CF.

In the US model, the manufacturer feature is dominant, especially for manufacturers DeWind, Guodian, Renewtech, NEG-Micon, Nordtank and CCWE. Some key metrics of these wind farm manufacturers (compared to the entire US wind farm dataset) are shown in Table 7.4. These manufacturers are unique in the dataset for a number of reasons. Firstly, they are a subset of the manufacturers with the fewest wind farms in the dataset (for reference the largest three are GE, Vestas and Siemens with 255, 127 and 68 wind farms respectively), making their instances quite unique; therefore if their data stand out in any way, the model will be sensitive to these dissimilarities. The metrics that potentially affect the model are the start year and average turbine size for Renewtech and Nordtank, and the mean lifetime CF for DeWind, Guodian, NEG-Micon and CCWE. The model is likely more sensitive to DeWind because it has five wind farms, compared to Guodian and CCWE's one, and because it has a significantly lower average CF than NEG-Micon.

Table 7.4: Key metrics for the manufacturers deemed important by the XGBoost model trained using US wind farm data.

Manufacturer	**Number of wind farms**	**Average CF**	**Average start year**	**Average wind turbine size**	**Average wind farm size**
DeWind	5	22.4%	2011	2.0 MW	32.0 MW
Guodian	1	24.2%	2012	1.5 MW	9.0 MW
Renewtech	1	27.2%	1986	0.1 MW	5.9 MW
NEG-Micon	11	26.0%	2002	1.21 MW	31.1 MW
Nordtank	4	24.2%	1987	0.3 MW	21.8 MW
CCWE	1	23.1%	2012	3.6 MW	61.2 MW
All	696	30.4%	2007	1.7 MW	94.1 MW

Another interesting feature is that the mean wind farm hub height wind speed is more important than the time-specific hub height wind speed for the US. Figure 7.5 shows the mean ERA5 wind speed at each wind farm, broken down by country. The US distribution shows wind farms with averages much lower than the UK and AUS, including averages around the cut-in wind speed region. While the XGBoost model does not "know" the power curve, it is likely to struggle to make accurate predictions for these wind farms that have such low reanalysis wind speed data. It is likely that this is a reanalysis error; otherwise, it seems implausible for a wind farm to be built in a location with such low wind speeds. However, it is still unexpected that mean reanalysis hub height wind speed is more important than the hub height wind speed. Mean wind speed is less important in the ALL model despite the same wind farms being present in the data, which is most likely due to the additional volume of data from the hourly measurements and to the increase in variance that results.

It is noticeable for the ALL model that Vestas is the most important manufacturer; however, it is not the largest manufacturer in the dataset. There are three dominant manufacturers in the dataset; GE (263 wind farms, 25% of total), Vestas (218 wind farms, 21% of total) and Siemens (127 wind farms, 12% of total). The key difference between the three manufacturers is the start years of the wind farms. Figure 7.6 shows the mean CF for the entire dataset (ALL), and for manufacturers GE, Vestas and Siemens (a), and the cumulative number of wind farms in the data set (b) for manufacturers GE, Vestas and Siemens. Out of the big three manufacturers in the dataset, Vestas is the only one with a significant number of wind farms before the year 2000. This also corresponds to the period of commissioned wind farms with the lowest mean lifetime CF. It is likely that the combination of Vestas being the second largest manufacturer in the dataset and having wind farms across a large range of years and mean CF values, makes it highly important in the ALL model (Figure 7.4 d).

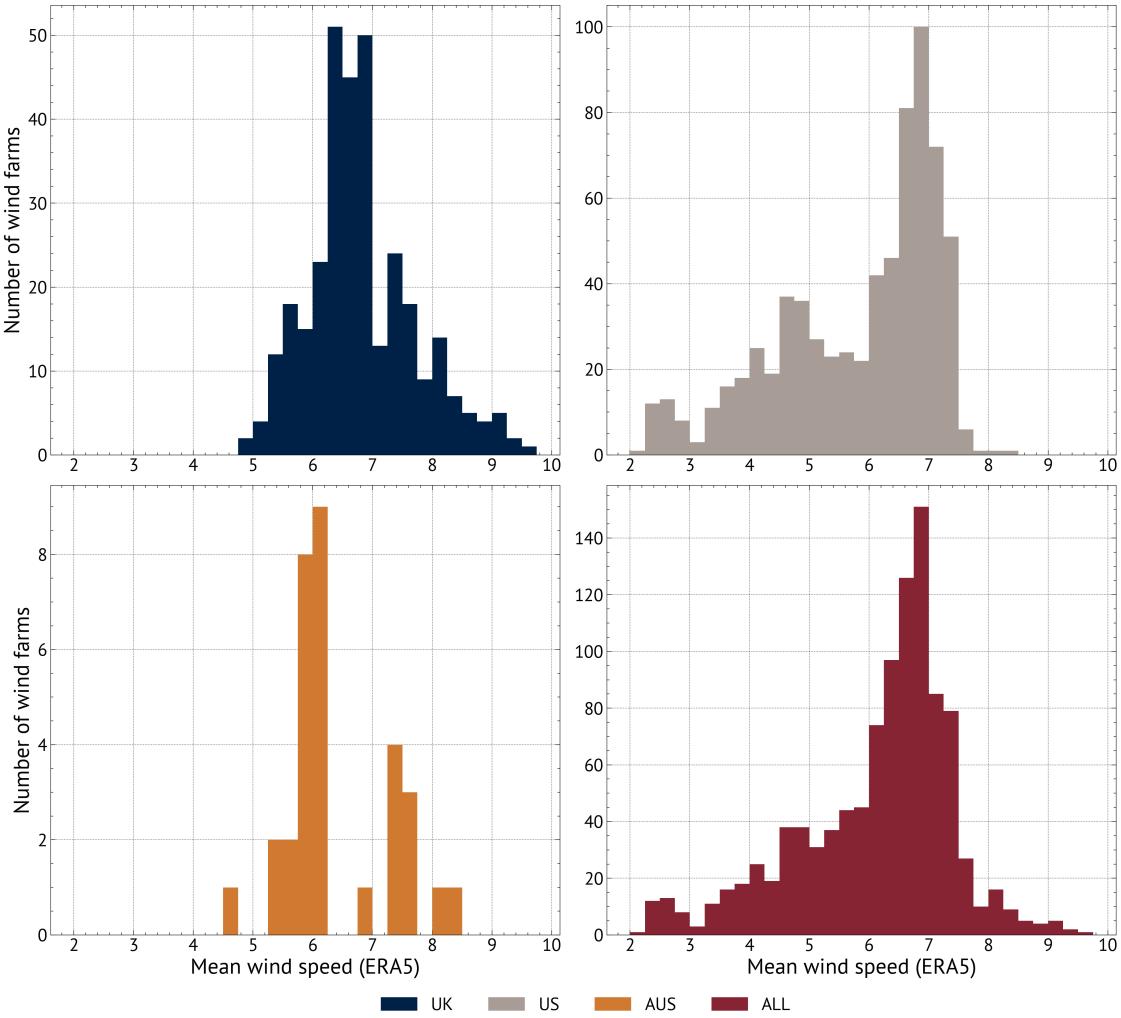


Figure 7.5: Mean ERA5 wind speed at each wind farm in the UK, US, AUS and ALL.

7.3.2 Data limitations

The model presented above is limited by the input data, specifically the use of reanalysis wind speed rather than observed wind speeds. Reanalysis data has location biases [140] and errors, as discussed in Chapter 6. Figure 6.14 shows the divergences between the ERA5 and MERRA-2 models and Figure 6.18 shows the average difference between CF_{obs} and CF_{est}. The limitations associated with using reanalysis data were evident in the CF_{est} results in Chapter 6, and these limitations carry over to the machine learning approach. The method to evaluate CF_{est} described in Section 5.5.5 provides a physics-based approach to estimating CF based on the reanalysis wind speed and can be compared to the supervised

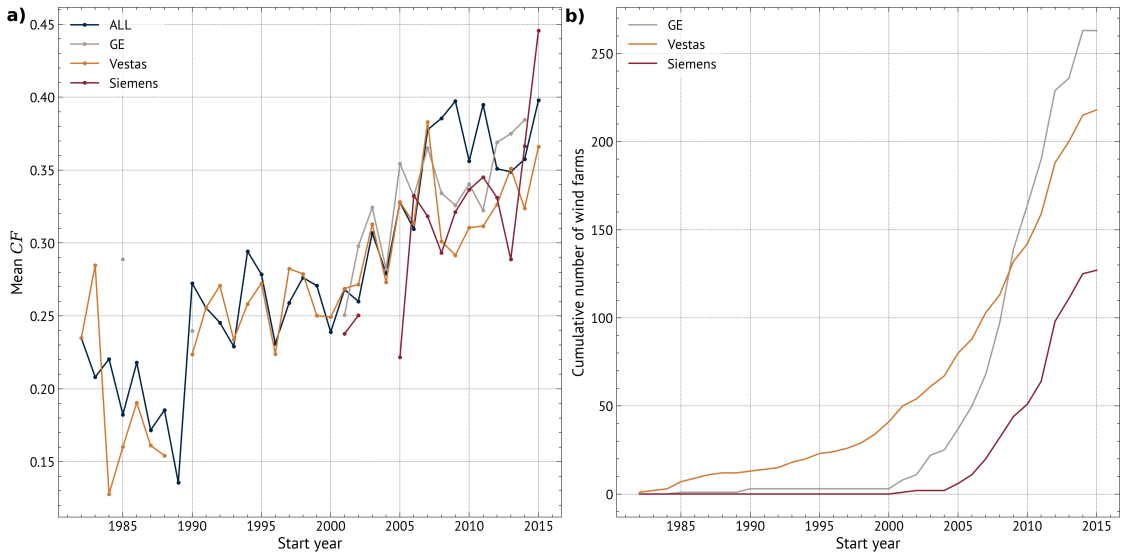


Figure 7.6: Mean CF for the entire dataset (ALL) by commissioning year, and manufacturers GE, Vestas and Siemens (a) and cumulative number of wind farms in the data set (b) for manufacturers GE, Vestas and Siemens.

ML model. To do this $|CF_{obs} - CF_{est}|$ and $|CF_{pred} - CF_{est}|$ can be calculated for all possible data points and compared. While CF_{obs} can be calculated for all data points, CF_{pred} is only calculated for 20% of the dataset (the validation set) - i.e. a limitation of the supervised ML model is the need to train it. To compare these methods fairly, the CF_{est} values are taken for the same data points used for CF_{pred}. Figure 7.7 shows the cumulative density function (CDF) for $|CF_{obs} - CF_{est}|$ and $|CF_{obs} - CF_{pred}|$ across the UK, US, AUS and ALL data and ERA5 (a) and MERRA-2 (b) models. This plot first shows that there is a significant difference between CF_{est} and CF_{obs}, which indicates that the power curve method, when used with reanalysis wind speed, is unable to capture all the trends in observed capacity factor. This occurs because of errors in the reanalysis data, inaccuracy in the wind farm power prediction model, and unexpected real-time events such as maintenance and curtailment. This figure also shows that the XGBoost model is significantly more accurate across both reanalysis datasets and all countries. However, given the limitations of using reanalysis data as the key metric, the model will be fundamentally limited to improve towards to a CF MAE close to 0 until data such as maintenance and curtailment figures are added.

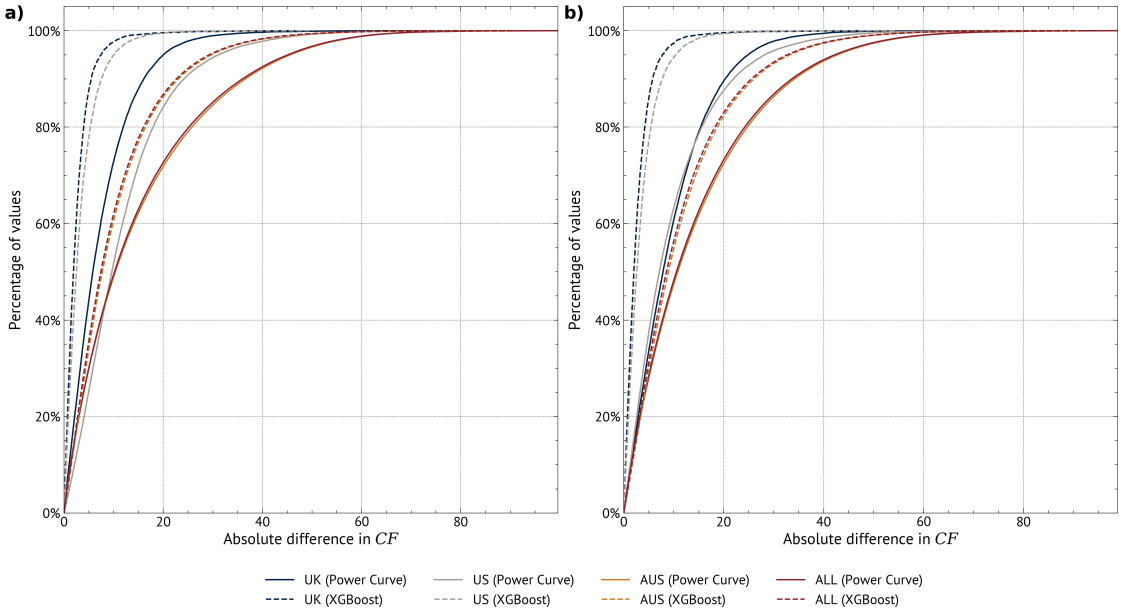


Figure 7.7: The cumulative density functions for $|CF_{obs} - CF_{est}|$ from Chapter 6 (labelled 'Power Curve'), and $|CF_{obs} - CF_{pred}|$ using the XGBoost model for ERA5 data (a) and MERRA-2 data (b).

To quantify the improvement achieved by using the ML model compared to the power curve method, MAE can be calculated for each model and dataset. Table 7.5 shows the MAE for machine learning models trained using UK, US, AUS and ALL data, as well as the MAE for the power curve model for UK, US, AUS and ALL wind farms. This equates to a reduction in MAE of 90%, 90%, 56% and 65% for the UK, US, AUS and ALL respectively.

Table 7.5: Comparison of MAE values for predicting CF_{obs} using the ML model and the power curve model.

Country	MAE (ML)	MAE (power curve)	MAE change	Percentage reduction
UK	2.576	24.69	22.11	90%
US	3.139	30.54	27.40	90%
AUS	10.34	23.62	13.28	56%
ALL	9.905	28.52	18.62	65%

7.3.3 Modelling results within a country of the dataset

As an additional test of the validity of the model, a LOOCV test, was completed where a single wind farm was left out of the training data to simulate modelling a wind farm that could exist within the domain of the dataset (i.e. within the countries in the dataset). This was done for all wind farms across models trained using wind farm data from the UK, US, AUS and ALL. The accuracy (calculated using MAE values) of these validation tests for each wind farm and model can be seen in Figure 7.8, and the mean (μ) and median (M) of the distributions are shown in the legend. The results indicate that for a wind farm in the UK and US, CF can, on average, be predicted within 4–5% of CF, and for AUS within 13%. On average results are positively skewed (i.e. $M > \mu$), except for the AUS model, which indicates the average accuracy is impacted by a small number of wind farms where MAE is large. The ALL dataset shows μ and M values similar to UK and US because this calculation is completed on a per wind farm basis, not a per data point basis (as in Figure 7.2), so it is not distorted by the mix of data periods and small number of AUS wind farms.

Using this analysis, the impact of metadata features can be inspected by identifying which wind farms have the lowest CF MAE when left out of the trained model. Figures 7.9 and 7.10 show the impact of wind farm size and number of turbines respectively by plotting them against MAE data extracted from the LOOCV analysis. Figure 7.9 shows an improvement in average MAE values as size of wind farm increases; however, there are fewer wind farms at larger farm size values (greater than 100 MW). There are also many wind farms with a MAE less than 5 and wind farm size less than 100 MW. The same trend exists for the US data, Figure 7.9 (b) shows the MAE values decreasing in general as wind farm size increases. There is no such pattern in the AUS data, which could be due to the much lower number of wind farms.

There is a lot more uncertainty in Figure 7.10, especially in the US data. This is primarily driven by wind farms with a large number of small turbines (less than 1 MW). These wind farms were often constructed prior to the year

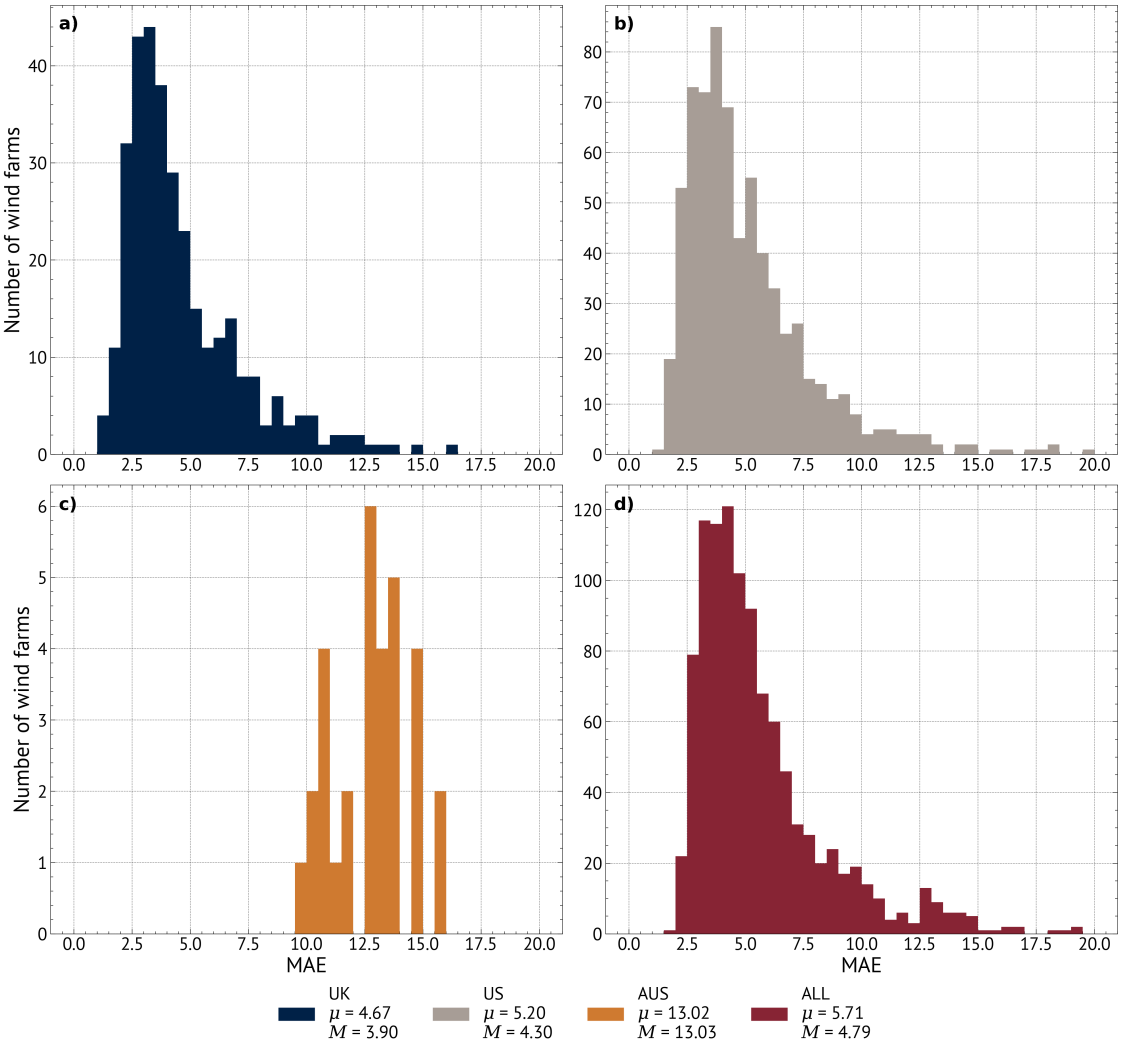


Figure 7.8: The MAE for predicting CF for each wind farm that was removed from the training dataset for (a) the UK, (b) US, (c) AUS and (d) ALL. Each wind farm is modelled using the ERA5 wind speed features.

2005 and with inferior technology (e.g. fixed pitch wind turbines), which makes predicting CF more difficult for a model that is trained with data from wind farms predominantly constructed after 2005.

Figure 7.11 shows the wind farm size in MW (a) and number of turbines (b) against MAE for the ALL model, with the scatter point colour indicating the start year of the wind farm. There are two key takeaways from this: a high number of turbines for the older wind farms does not result in a large wind farm capacity; and the older wind farms generally drive the MAE up in plot (b). The older wind farms in Figure 7.11 (a) are congregated at smaller wind farm sizes but do not exhibit

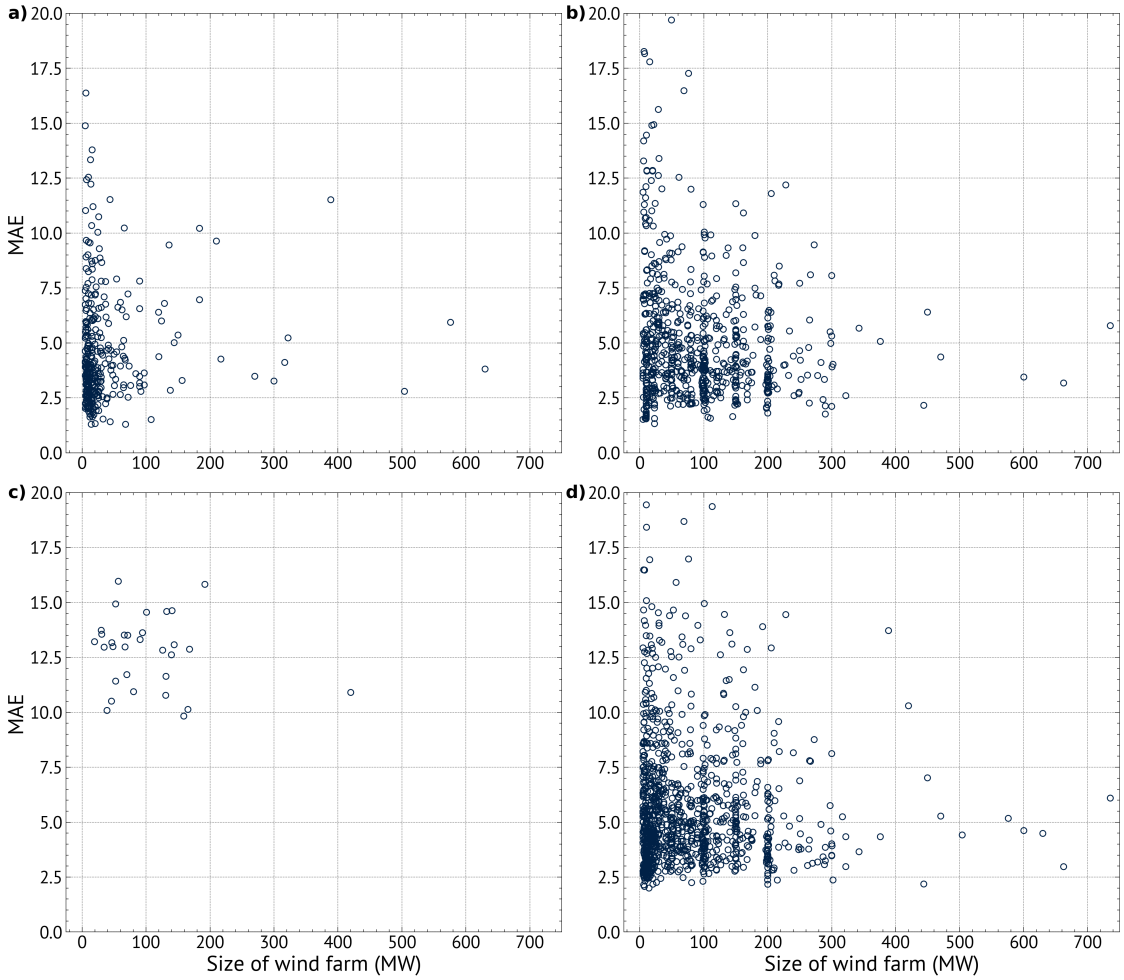


Figure 7.9: The size of wind farm in MW plotted against the MAE of CF_{pred} for the LOOCV analysis split by models trained with wind farm data from the UK (a), US (b), AUS (c) and ALL (d).

larger MAE than other wind farms of similar size. These plots therefore indicate that without the presence of the wind farms with a large number of turbines but low wind farm size, the model will be able to use the number of turbines variable more effectively, but it will still have a problem predicting CF for wind farms with a relatively small rated capacity.

7.4 Discussion and conclusions

Using reanalysis data results in clear limitations for hindcast predicting wind farm performance. Reanalysis models do not always reflect observed wind speeds, which makes performance prediction more difficult. In the simplest sense, the model is

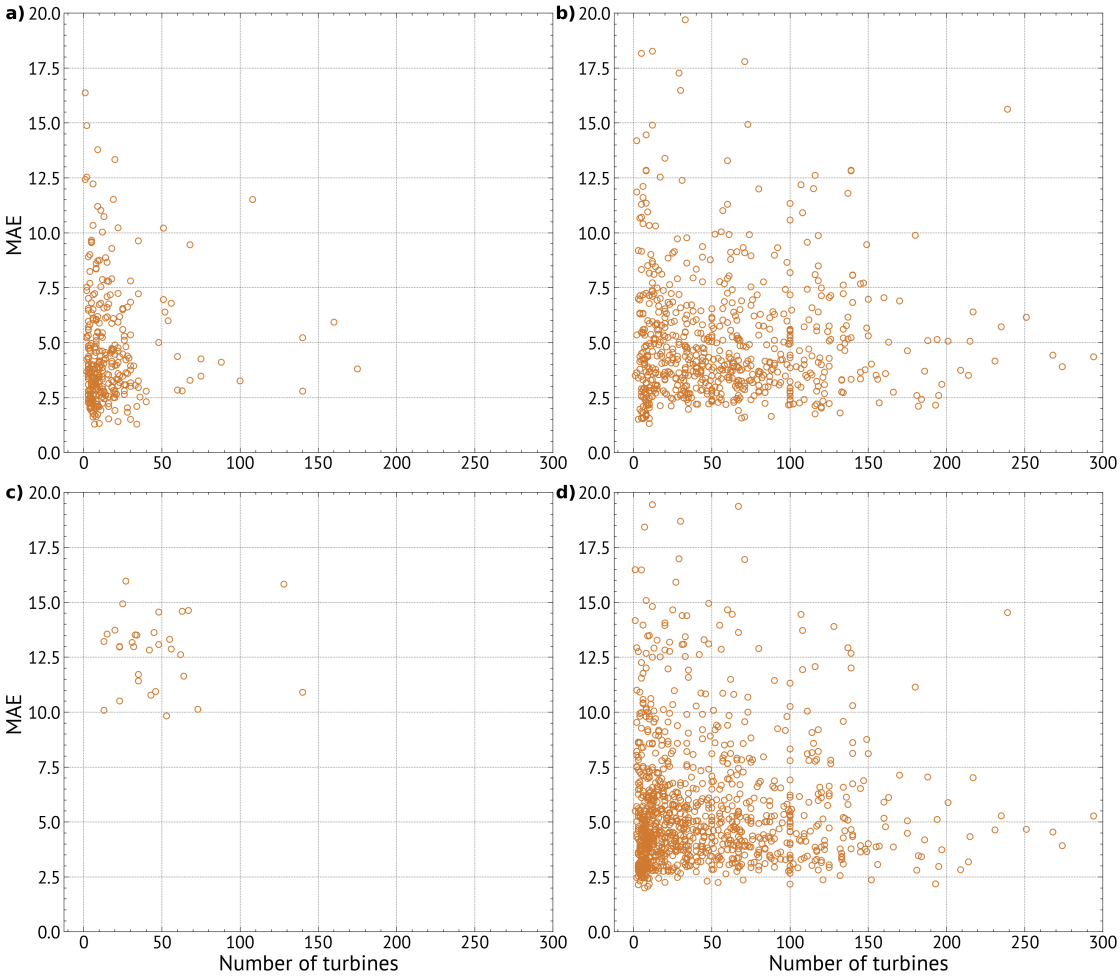


Figure 7.10: The number of turbines in a wind farm plotted against the CF MAE for the LOOCV analysis split by models trained with wind farm data from the UK (a), US (b), AUS (c) and ALL (d).

looking for trends in the data, so the hope is that the model will be able to 'back calculate' the wind farm power curve. Any erroneous reanalysis data (especially when hub height wind speed is a highly important feature) make this trend much more difficult to find. Nevertheless, it was shown that a trained supervised XGBoost machine learning model was able to significantly outperform a wind farm power curve estimate of CF. The limitation of the machine learning model compared to the wind farm power curve approach is the need to use a large proportion of the dataset (80% in this case) to train the model. However, given that the scope of this work is to produce a prediction model for future wind farms, this is not a serious constraint. To produce forward looking predictions, reanalysis data would need to

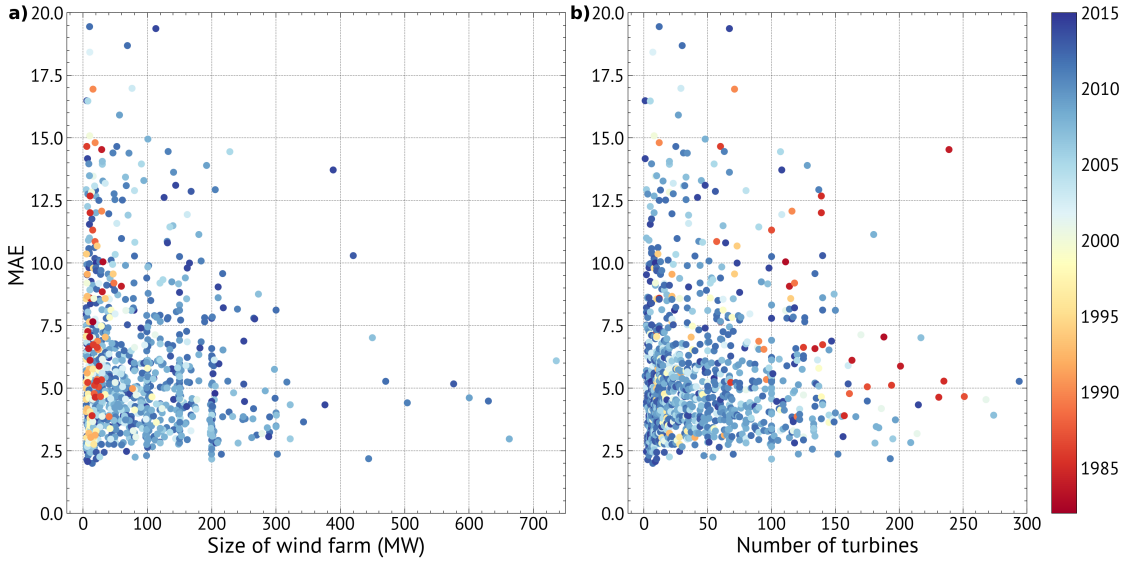


Figure 7.11: The size of wind farm in MW plotted against the *CF* MAE (a), and the number of turbines in a wind farm plotted against the MAE (b) for the LOOCV analysis split by models trained with wind farm data from ALL. The colour of the scatter points indicate the start year of the wind farm.

be available in real time (highly unlikely) or a different set of wind speeds would need to be used (e.g., weather station wind speed data). If the data were available, this would be an important expansion on the work presented in this chapter.

Throughout the results section, the key questions proposed at the beginning of the chapter were answered (other than questions 1 and 2, as discussed above).

Does turbine manufacturer impact performance on average?

Figure 7.4 shows that for different countries, the turbine manufacturer for wind farms had varying importance. The US model (Figure 7.4 (b)) is highly affected by manufacturers with a small number of wind farms in the sample, as explained in Section 7.3.1. In the other models, the manufacturers with large numbers of wind farms in the dataset show up as important features, coming after the wind-speed-related variables. However, due to the large number of variables into which the manufacturer categorical variable gets split (one for each unique manufacturer in the dataset), some manufacturers have little importance. Some manufacturers have little importance because they exhibit expected performance for given wind speeds,

very small samples need to be included to capture outliers that would otherwise not be described, and very large samples need to be included to capture overall trends.

Outside of the wind-speed-related features, manufacturer is, on average, the next most important variable, as shown in Figure 7.4.

How important is location for wind farm performance?

Similar to turbine manufacturer, location factors well below wind speed in its importance, as shown by the 'lat' and 'lon' features in Figure 7.4. The model in which location was the most important is the model trained on ALL wind farm data. This makes sense because the location now spreads across countries located around the world. In Figure 7.4 (d), the latitude and longitude features are two of the top eleven most important features. It would be expected that as more countries were added to the dataset and the latitudes and longitudes were more diverse, the more important location would be in predicting power output.

How does wind speed direction impact power output for each wind farm?

Adding the angle of the wind speed to the model improved its accuracy, but only very marginally (less than 0.5% improvement in MAE). In Figure 7.4, angle only shows up in the top 90% in the ALL model, and it is low on the importance score. In all other cases it did not contribute to the top 90% importance. This indicates that the angle of the wind speed when taken from reanalysis datasets, does not significantly influence power output prediction.

How do the size and number of turbines impact the power curve?

There is an indication that the size of the wind farm has a material impact on the ability to predict CF. Figure 7.9 shows that as size of wind farm increases, the average MAE decreases. However the number of data points also decreases as size of wind farm increases, which adds uncertainty. Nevertheless, this figure clearly shows the variance in MAE values for wind farms with a small rated capacity (less than 50 MW). This is an encouraging result because the larger wind farms are much more important, in terms of their impact on the grid.

There is a less clear relationship between the number of turbines in a wind farm and the accuracy of the model. Figure 7.10 shows the number of turbines in a wind farm plotted against MAE, and there is a high amount of variance in the results. Figure 7.11 relates some of the variance in these results to the wind farms with earlier start dates, indicating that this metric may be a better predictor for wind farms constructed more recently (after 2005).

The impact of the size of wind farm and number of turbines shown in Figures 7.9, 7.10 and 7.11 indicates the difficulty of predicting CF for a wind farm with a low rated capacity. These figures also show that these metrics are not always correlated, making manual adjustment difficult and promoting a ML method that can find the combined trends of these variables.

Does the averaging period of the time series data impact how to best model power output?

The accuracy results presented in Table 7.2 clearly indicate that the averaging period of the time series data impacts how to model power output. This is expected because of the higher variability of hourly data when compared to monthly data. Therefore, the time resolution of data inputted into the model is the time resolution of CF that should be predicted from the model. The model cannot predict what is outside the input data in the same way that interpolating from a power curve can.

Those who contemplate the beauty of the earth find reserves of strength that will endure as long as life lasts. There is something infinitely healing in the repeated refrains of nature – the assurance that dawn comes after night, and spring after winter.

— Rachel Carson [197]

8 Conclusion

Contents

8.1 Conclusions 200
8.1.1 Wind turbine roughness 200
8.1.2 Long-term correction and wind farm performance decline 201
8.1.3 Wind farm power prediction 202
8.2 Future work 202
8.2.1 Roughness data 202
8.2.2 Wind farm data and reporting standards 202
8.3 Contributions 203

This thesis has investigated the lifetime performance of wind turbines and wind farms. Changes in wind turbine performance have been analysed through the effect of roughness on the airfoil surface, which alters the aerodynamic lift and drag properties of the airfoil. Through this analysis, control strategies to mitigate performance loss have been investigated numerically, and a framework has been created to account for the impact of roughness on the reduction of lift to drag ratio relative to the clean airfoil data. Wind farms have been analysed through the examination of their output data in combination with reanalysis wind speed data. This process led to the quantification of wind farm degradation rates. The data used for this process was also utilised to predict wind farm power output through the implementation of a supervised machine learning model.

The conclusions of this thesis are summarised in three parts: the effect of roughness on wind turbines and how to predict it, summarising Chapters 2 to 3; long-term correction and wind farm performance decline, summarising Chapters 5 to 6; and finally wind farm power prediction using reanalysis data, summarising Chapter 7. This chapter concludes with suggestions for future work.

8.1 Conclusions

8.1.1 Wind turbine roughness

Performance loss of wind turbines due to blade roughness was quantified through the use of experimental airfoil data and an in-house designed rotor. Energy yield reduction due to roughness was shown to be in the range of 2.9–8.6%. A control method that allowed a rough-bladed rotor to optimise its performance (based on its roughened aerodynamical performance), rather than operating with a clean-turbine control strategy, was created to mitigate performance losses. When implemented, this method mitigated up to 1.0% of performance losses. Turbine roughness is inevitable over the lifetime of a wind farm; however, by accounting for its effect on blade aerodynamics, a more optimal control strategy can be used. The implementation of this strategy (and therefore the improved performance of rough-bladed turbines) is possible because on current operational turbines, control strategies do not adapt over time. With knowledge of the level of blade roughness, the operating point can be adjusted to increase turbine power.

While the recovered power is not large, the solution is only a software adjustment and does not require halting turbines. The amount of power recovered will depend on the rotor used and level of roughness.

A roughness evolution parameter was then created through the aggregation of existing experimental airfoil roughness datasets that allows researchers to convert clean airfoil data to roughened data. This was completed by finding trends in lift and drag data when experimental datasets were combined, before converting them into the roughness evolution parameter. This addresses a key limitation of this research area: access to experimental roughened airfoil data. A framework was

created to improve on the roughness evolution parameter in the future through the addition of future experimental datasets. Airfoil roughness was then able to be used as input to a wind turbine model to quantify decline of wind turbine performance and AEP due to increasing roughness.

8.1.2 Long-term correction and wind farm performance decline

Individual wind turbine performance decline is different from wind farm performance decline which was analysed through the implementation of a long-term correction algorithm. LTC analyses are used to extract the underlying performance trend of a wind farm by removing variable effects – mainly wind speed variation. They do this by using reanalysis wind speed data to correct CF data over the lifetime of the wind farm.

The fundamental challenge of completing a LTC analysis is correcting the data when the modelled CF is significantly different from the observed CF. This can be due to reanalysis wind speed errors, downtime or curtailment. A novel LTC method has been designed to improve on previous methods by accounting for these periods and improving the correction calculation. This approach is not found in previous LTC applications and is a novel contribution.

Through this process, wind farm performance declines in the UK, US, Australia, and all three countries combined were able to be quantified – 0.27, 0.75, 0.41 and 0.59 CF percentage points (per year of age) respectively. LTC degradation rates were higher than the raw data, which can only partially be attributed to long-term wind speed change. Furthermore, limitations of the linear trend were shown when datasets of varying age and generation wind farms are combined. When offshore data was removed from the larger dataset, they were found to have a performance decline rate of 0.21 CF percentage points. Through the process of isolating the offshore farms dataset, a limitation of averaging across cohorts with significant differences in mean lifetime CF was shown, and this will need to be accounted for in future studies.

This was the first study to be completed across different time resolution data, with the Australian dataset being hourly and the UK and US, monthly. The LTC process was completed across three countries, and a framework is provided for future research.

8.1.3 Wind farm power prediction

The dataset amassed throughout Chapters 5 to 6, containing both wind farm power output and wind farm metadata, was used to train a supervised machine learning model to predict CF. Two supervised machine learning techniques were considered: extreme gradient boosting (XGBoost) and Random Forest. Four models were trained, using wind farm data from the UK, US, Australia, and all three countries combined. XGBoost performed marginally more accurately across all models with MAE values (when predicting observed CF) of 2.713, 3.397, 10.45 and 10.04 for the UK, US, AUS and ALL respectively.

Using additional wind farm metadata to predict CF, the machine learning model was found to have a 56–90% improvement in CF mean absolute error value when compared to predicting power output using a wind farm power curve.

8.2 Future work

8.2.1 Roughness data

The roughness evolution parameter is built on the foundation of experimental airfoil roughness studies. It is limited by the small number of datasets that have been publicly published. Further studies are needed that detail the effect of roughness when it progressively increases in a consistent way over time. The framework is set up in a way to support additional data and future research, which is encouraged to improve the accuracy and applicability of the parameter.

8.2.2 Wind farm data and reporting standards

One of the key limitations of Chapters 5 to 7 is access to key data metrics such as downtime and curtailment. These metrics are simple to understand and possible to

analyse. Curtailment is not always as simple to define as downtime but an estimate is possible. Simple data inputs such as these would allow the LTC algorithm and supervised machine learning models to be significantly improved. Furthermore, it is expected that these events will only increase in the future due to the increase in installed capacity and penetration of variable renewable energy in electricity grids.

The other key limitation is the accuracy of reanalysis models. If recorded wind speed data for each wind farm location were available, it would significantly improve the LTC algorithm and supervised machine learning model. However, these datasets are less likely to become publicly available than downtime and curtailment data due to their complexity and limited impact on electricity markets.

8.3 Contributions

Roughened turbine control strategy

A method to 'retune' a turbine as its blades become increasingly roughened has been proposed. This control strategy is entirely driven by operational software and can provide a way to produce more power out of wind turbines as the aerodynamic surfaces degrade.

Roughness evolution parameter

A roughness evolution parameter has been defined using existing experimental airfoil roughness data. This parameter has never previously been defined and enables further study into wind turbine roughness by allowing roughened airfoil data to be created from clean airfoil data.

LTC algorithm

A novel long-term correction algorithm has been presented that improves upon previous methods, which are limited due to over-simplifying wind farm power output calculations, not accounting for errors in the reanalysis wind speed data, and not recognizing their unsuitability for higher time resolution datasets. The novel algorithm aims to correct data based off long term trends, but includes

mitigation measures to ensure over correction does not occur and periods where the reanalysis wind speed may be erroneous are accounted for. It also accounts for wind speed variability when calculating estimated CF for a wind farm. This is the first method to be validated over higher time resolution data and quantifies degradation rates for UK, US and Australian wind farms.

Supervised machine learning model and dataset

Wind farm power output data, reanalysis wind speed data, and wind farm metadata have been combined to produce a capacity factor prediction supervised machine learning model. This is a new method for predicting wind farm power output and is made possible by having a unique wind farm meta dataset. This dataset is very valuable and should be continued to be expanded as more wind farms come online.

Appendices

Thought is the wind, knowledge the sail, and mankind the vessel.

— Augustus Hare

A Wind Farm Metadata

Table A.1: Wind farm metadata variables used in this study and their description.

Metadata name	Description
wind_farm_id	Identifier used by local region authority
wind_farm_name	Wind farm name
turbine_capacity_mw	Individual turbine capacity (MW)
number_of_turbines	Number of turbines in farm
total_capacity	Capacity of wind farm
hub_height	Hub height of turbines
blade_diameter	Diameter of the turbine rotor
model_name	Wind turbine model where listed in the public domain
model_name_filled_in	Missing data from model_name filled in with most suitable wind turbine model
power_curve	Power curve used for calculating CF_{est}. Can differ from model_name_filled_in depending on availability of power curves
manufacturer	Wind turbine manufacturer
mean_CF	Average capacity factor based on data in the dataset
extension_date	Date the wind farm was extended/repowered (if applicable)
extension_number_turbines	Number of turbines in the extension/repowering (if applicable)
extension_turbine_size	Size of turbines in the extension/repowering (if applicable)
start_date	Commissioning date of the wind farm
start_year	Commissioning year of the wind farm
lat	Latitude of the wind farm
lon	Longitude of the wind farm
country	Country of the wind farm
location_type	Either 'onshore' or 'offshore' depending on the location of the wind farm
retired	Whether or not the wind farm has been retired. 'Y' for yes and 'N' for no
c_MERRA2	Weibull scale factor of the lifetime wind distribution from MERRA-2
k_MERRA2	Weibull shape factor of the lifetime wind distribution from MERRA-2
c_ERA5	Weibull scale factor of the lifetime wind distribution from ERA5
k_ERA5	Weibull shape factor of the lifetime wind distribution from ERA5
l_MERRA2	Weibull location parameter of the lifetime wind distribution from MERRA-2
l_ERA5	Weibull location parameter of the lifetime wind distribution from ERA5
mean_ws_ERA5	Lifetime mean wind speed from ERA5
mean_ws_MERRA2	Lifetime mean wind speed from MERRA-2
time_resolution	Time resolution of the time series data

Everybody is smarter than anybody.

— Carl Sandburg

References

[1] Monique Hoogwijk, Bert de Vries, and Wim Turkenburg. "Assessment of the global and regional geographical, technical and economic potential of onshore wind energy". In: *Energy Economics* 26 (2004), pp. 889–919.

[2] Cristina L. Archer and Mark Z. Jacobson. "Evaluation of global wind power". In: *Journal of Geophysical Research* 110.D12 (2005).

[3] Xi Lu, Michael B. McElroy, and Juha Kiviluoma. "Global Potential for Wind-Generated Electricity". In: *Proceedings of the National Academy of Sciences of the United States of America* 106.27 (2009), pp. 10933–10938.

[4] Paul Denholm, Gerald L Kulcinski, and Tracey Holloway. "Emissions and energy efficiency assessment of baseload wind energy systems". In: *Environmental Science & Technology* 39.6 (2005), pp. 1903–1911.

[5] OA Jaramillo, R Saldaña, and U Miranda. "Wind power potential of Baja California Sur, Mexico". In: *Renewable Energy* 29.13 (2004), pp. 2087–2100.

[6] Lin Lu, Hongxing Yang, and John Burnett. "Investigation on wind power potential on Hong Kong islands—An analysis of wind power and wind turbine characteristics". In: *Renewable Energy* 27.1 (2002), pp. 1–12.

[7] Nicolas Boccard. "Capacity Factor of Wind Power: Realized Values vs. Estimates". In: *Energy Policy* 37 (2009), pp. 2679–2688.

[8] Akira Inoue, Rion Takahashi, Toshiaki Murata, Junji Tamura, Mamoru Kimura, Moto-O Futami, and Kazumasa Ide. "A calculation method of the total efficiency of wind generators". In: *Electrical Engineering in Japan* 157.3 (2006), pp. 52–62.

[9] Alfred J. Cavallo. "High-Capacity Factor Wind Energy Systems". In: *Journal of Solar Energy Engineering* 117.2 (May 1995), pp. 137–143.

[10] Linyue Gao, Ramsankar Veerakumar, Yang Liu, and Hui Hu. "Quantification of the 3D shapes of the ice structures accreted on a wind turbine airfoil model". In: *Journal of Visualization* 22 (2019), pp. 661–667.

[11] Rebecca Milligan. *Wind-power – now cheaper than gas*. Accessed on 20/12/19. URL: `https://energysavingtrust.org.uk/blog/wind-power-%E2%80%93-now-cheaper-gas`.

[12] *World's Most Powerful Offshore Wind Turbine: Haliade-X 12 MW: GE Renewable Energy*. Date accessed 15/02/2020. URL: `https://www.ge.com/renewableenergy/wind-energy/offshore-wind/haliade-x-offshore-turbine`.

[13] Maria Isabel Blanco. "The economics of wind energy". In: *Renewable and Sustainable Energy Reviews* 13 (2009), pp. 1372–1382.

[14] Philip Heptonstall, Robert Gross, Philip Greenacre, and Tim Cockerill. "The cost of offshore wind: Understanding the past and projecting the future". In: *Energy Policy* 41 (2012), pp. 815–821.

[15] Peter Blasco, Jose Palacios, and Sven Schmitz. "Effect of icing roughness on wind turbine power production". In: *Wind Energy* 20.4 (2017), pp. 601–617.

[16] G Pechlivanoglou, S Fuehr, CN Nayeri, and CO Paschereit. "The effect of distributed roughness on the power performance of wind turbines". In: *Turbo Expo: Power for Land, Sea, and Air*. Vol. 44007. 2010, pp. 845–855.

[17] Matthias Schramm, Hamid Rahimi, Bernhard Stoevesandt, and Kim Tangager. "The Influence of Eroded Blades on Wind Turbine Performance Using Numerical Simulations". In: *Energies* 10.9 (2017), p. 1420.

[18] Iain Staffell and Richard Green. "How does wind farm performance decline with age?" In: *Renewable Energy* 66 (2014), pp. 775–786.

[19] Jon Olauson, Per Edström, and Jesper Rydén. "Wind turbine performance decline in Sweden". In: *Wind Energy* 20.12 (2017), pp. 2049–2053.

[20] Sonja Germer and Axel Kleidon. "Have wind turbines in Germany generated electricity as would be expected from the prevailing wind conditions in 2000-2014?" In: *PLOS one* 14.2 (2019), pp. 1–16.

[21] Sofia D Hamilton, Dev Millstein, Mark Bolinger, Ryan Wiser, and Seongeun Jeong. "How does wind project performance change with age in the United States?" In: *Joule* 4.5 (2020), pp. 1004–1020.

[22] Anastasia Ioannou, Andrew Angus, and Feargal Brennan. "Stochastic prediction of offshore wind farm LCOE through an integrated cost model". In: *Energy Procedia* 107 (2017), pp. 383–389.

[23] Anastasia Ioannou, Andrew Angus, and Feargal Brennan. "Risk-based methods for sustainable energy system planning: A review". In: *Renewable and Sustainable Energy Reviews* 74 (2017), pp. 602–615.

[24] Carl Sagan and Ann Druyan. *Pale blue dot: A vision of the human future in space*. Ballantine Books, 2011.

[25] P Haines and J Luers. "Aerodynamic penalties of heavy rain on landing airplanes". In: *Journal of Aircraft* 20.2 (1983), pp. 111–119.

[26] James R Valentine and Rand A Decker. "Tracking of raindrops in flow over an airfoil". In: *Journal of Aircraft* 32.1 (1995), pp. 100–105.

[27] R John Hansman and Martitia F Barsotti. "Surface wetting effects on a laminar flow airfoil in simulated heavy rain". In: *Journal of Aircraft* 22.12 (1985), pp. 1049–1053.

[28] Nihad Dukhan, Kenneth J De Witt, KC Masiulaniec, and G James Van Fossen. "Experimental Frossling numbers for ice-roughened NACA 0012 airfoils". In: *Journal of Aircraft* 40.6 (2003), pp. 1161–1167.

[29] Walter O Valarezo, Frank T Lynch, and Robert J McGhee. "Aerodynamic performance effects due to small leading-edge ice (roughness) on wings and tails". In: *Journal of Aircraft* 30.6 (1993), pp. 807–812.

[30] Sam Lee and Michael B Bragg. "Experimental investigation of simulated large-droplet ice shapes on airfoil aerodynamics". In: *Journal of Aircraft* 36.5 (1999), pp. 844–850.

[31] Ece Sagol, Marcelo Reggio, and Adrian Ilinca. "Issues concerning roughness on wind turbine blades". In: *Renewable and Sustainable Energy Reviews* 23 (2013), pp. 514–525.

[32] F Bañuelos-Ruedas, C Angeles-Camacho, and S Rios-Marcuello. "Analysis and validation of the methodology used in the extrapolation of wind speed data at different heights". In: *Renewable and Sustainable Energy Reviews* 14.8 (2010), pp. 2383–2391.

[33] Marshall L. Buhl. Wind Turbine Airfoils, NREL – National Renewable Energy Laboratory, NWTC Information Portal, date accessed 01/10/2019. URL: `https://wind.nrel.gov/airfoils/Documents/AirfoilDocuments.html`.

[34] Gustave P. Corten and Herman F. Veldkamp. "Insects can halve wind-turbine power". In: *Nature* 412.6842 (2001), pp. 41–42.

[35] G Petrone, C de Nicola, D Quagliarella, J Witteveen, and G Iaccarino. "Wind turbine performance analysis under uncertainty". In: *49th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition* (2011).

[36] Ali Doosttalab et al. "Near-wall turbulence and its modulation to the outer flow: From canonical ZPG to wind turbine airfoils". PhD thesis. Texas Tech University, 2018.

[37] G.P Corten. "Flow separation on wind turbine blades". PhD thesis. University of Utrecht, 2001.

[38] Timo Laakso, H Holttinen, G Ronsten, L Tallhaug, R Horbaty, I Baring-Gould, A Lacroix, E Peltola, and B Tammelin. "State-of-the-art of wind energy in cold climates". In: *IEA Annex XIX* 24 (2003), p. 53.

[39] S Barber, Y Wang, S Jafari, N Chokani, and Reza S Abhari. "The impact of ice formation on wind turbine performance and aerodynamics". In: *Journal of Solar Energy Engineering* 133.1 (2011). 011007.

[40] Øyvind Byrkjedal and K Vindteknikk. "Estimating wind power production loss due to icing". In: *13th International Workshop on Atmospheric Icing of Structures* (2009).

[41] Matthew C Homola, Muhammad S Virk, Per J Nicklasson, and Per A Sundsbø. "Performance losses due to ice accretion for a 5 MW wind turbine". In: *Wind Energy* 15.3 (2012), pp. 379–389.

[42] Yiqiang Han and Jose Palacios. "Analytical and experimental determination of airfoil performance degradation due to ice accretion". In: *4th AIAA Atmospheric and Space Environments Conference* (2012), p. 2794.

[43] William Jasinski, Shawn Noe, Michael Selig, Michael Bragg, William Jasinski, Shawn Noe, Michael Selig, and Michael Bragg. "Wind turbine performance under icing conditions". In: *35th Aerospace Sciences Meeting and Exhibit* (1997), p. 977.

[44] David Switchenko, Wagdi Habashi, Thomas Reid, Isik Ozcer, and Guido Baruzzi. "FENSAP-ICE simulation of complex wind turbine icing events, and comparison to observed performance data". In: *32nd ASME Wind Energy Symposium* (2014), p. 1399.

[45] Michael B Bragg, Andy P Broeren, and Leia A Blumenthal. "Iced-airfoil aerodynamics". In: *Progress in Aerospace Sciences* 41.5 (2005), pp. 323–362.

[46] *XTurb-PSU*. date accessed 12/12/2019. URL: `https://www.rotoraero.psu.edu/xturb-psu/`.

[47] Robert D Corrigan and Richard D DeMiglio. "Effect of precipitation on wind turbine performance". In: *Report No. NASA-TM-86986* (1985).

[48] David C Maniaci, Carsten Westergaard, Alan Hsieh, and Joshua A Paquette. "Uncertainty quantification of leading edge erosion impacts on wind turbine performance". In: *Journal of Physics: Conference Series* 1618.5 (2020), p. 052082.

[49] Alessandro Corsini, Alessio Castorrini, Enrico Morei, Franco Rispoli, Fabrizio Sciulli, and Paolo Venturini. "Modeling of rain drop erosion in a multi-MW wind turbine". In: *ASME Paper No. GT2015-42174* (2015).

[50] Junlei Chen, Jihui Wang, and Aiqing Ni. "A review on rain erosion protection of wind turbine blades". In: *Journal of Coatings Technology and Research* 16.1 (2019), pp. 15–24.

[51] N Dalili, Afsaneh Edrisy, and Rupp Carriveau. "A review of surface engineering issues critical to wind turbine performance". In: *Renewable and Sustainable Energy Reviews* 13.2 (2009), pp. 428–438.

[52] Amrit Shankar Verma, Zhiyu Jiang, Zhengru Ren, Marco Caboni, Hans Verhoef, Harald van der Mijle-Meijer, Saullo GP Castro, and Julie JE Teuwen. "A probabilistic long-term framework for site-specific erosion analysis of wind turbine blades: A case study of 31 Dutch sites". In: *Wind Energy* 24.11 (2021), pp. 1315–1336.

[53] Mohammed G Khalfallah and Aboelyazied M Koliub. "Effect of dust on the performance of wind turbines". In: *Desalination* 209.1-3 (2007), pp. 209–220.

[54] Walt Musial and Sandy Butterfield. *Future for offshore wind energy in the United States*. Tech. rep. National Renewable Energy Lab., Golden, CO (US), 2004.

[55] Mark Hugh Keegan, DH Nash, and MM Stack. "On erosion issues associated with the leading edge of wind turbine blades". In: *Journal of Physics D: Applied Physics* 46.38 (2013), p. 383001.

[56] Xiao Hui Dong, Tie Jun Yuan, and Ru Hong Ma. "Corrosion mechanism on offshore wind turbine blade in salt fog environment". In: *Applied Mechanics and Materials* 432 (2013), pp. 258–262.

[57] J.M. Janiszewska, R.R. Ramsay, M.J. Hoffmann, and G.M. Gregorek. "Effects of grit roughness and pitch oscillations on the S814 airfoil". In: *National Renewable Energy Laboratory* (1996).

[58] Agrim Sareen, Chinmay A Sapre, and Michael S Selig. "Effects of leading edge erosion on wind turbine blade performance". In: *Wind Energy* 17.10 (2014), pp. 1531–1542.

[59] Xabier Munduate and Esteban Ferrer. "CFD Predictions of Transition and Distributed Roughness Over a Wind Turbine Airfoil". In: *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition* (2009), p. 269.

[60] Mark Drela. "XFOIL: An analysis and design system for low Reynolds number airfoils". In: Springer, 1989, pp. 1–12.

[61] Bonnie Jonkman and Jason Jonkman. "FAST v8. 16.00". In: *National Renewable Energy Laboratory* (2016).

[62] J. Jonkman, S. Butterfield, W. Musial, and G. Scott. "Definition of a 5-MW Reference Wind Turbine for Offshore System Development". In: *Technical Report NREL/TP-500-38060* (2009).

[63] Martin O.L.. Hansen. *Aerodynamics of wind turbines*. 2nd ed. Wiley, 2009.

[64] W Rankine. "On the mechanical properties of propellers". In: *Transactions of the Institute of Naval Architects* 6 (1865), pp. 13–30.

[65] R Froude. "On the part played in propulsion by differences of fluid pressure". In: *Transactions of the Institute of Naval Architects* 30 (1889), pp. 390–405.

[66] F Lanchester. "A contribution to the theory of propulsion and the screw propeller". In: *Transactions of the Institute of Naval Architects* 57 (1915), pp. 98–116.

[67] A Betz. "Das Maximum der theoretisch moglichen Ausnutzung des Windes durch Windmotoren". In: *Zeitschrift fur das gesamte Turbinenwesen* 26 (1920), pp. 307–309.

[68] H. Glauert. *The elements of aerofoil and airscrew theory*. Cambridge University Press, 1926.

[69] J. F. Manwell, J. G. McGowan, and Anthony L. Rogers. *Wind energy explained: Theory, design and application*. 2nd ed. Wiley, 2009.

[70] Tony Burton, Nick Jenkins, David Sharpe, and Ervin Bossanyi. *Wind energy handbook*. Chichester: John Wiley & Sons, 2011.

[71] S Andrew Ning. "A simple solution method for the blade element momentum equations with guaranteed convergence". In: *Wind Energy* 17.9 (2014), pp. 1327–1345.

[72] R. P. Brent. "An algorithm with guaranteed convergence for finding a zero of a function". In: *The Computer Journal* 14.4 (1971), pp. 422–425.

[73] Jacob Aho, Andrew Buckspan, Jason Laks, Paul Fleming, Yunho Jeong, Fiona Dunne, Matthew Churchfield, Lucy Pao, and Kathryn Johnson. "A tutorial of wind turbine control for supporting grid frequency through active power control". In: *2012 American Control Conference (ACC)* (2012), pp. 3120–3131.

[74] Ludwig Prandtl and Albert Betz. *Vier Abhandlungen zur Hydrodynamik und Aerodynamik*. Göttingen: Universitätsverlag Göttingen, 1927.

[75] M L Buhl. "New Empirical Relationship between Thrust Coefficient and Induction Factor for the Turbulent Windmill State". In: *Technical Report NREL/TP-500-36834* (Aug. 2005).

[76] Jack Kelly, Christopher Vogel, and Richard Willden. "Impact and mitigation of blade surface roughness effects on wind turbine performance". In: *Wind Energy* 25.4 (2022), pp. 660–677.

[77] Jack Kelly, Richard Willden, and Christopher Vogel. "Parameterising the Impact of Roughness Evolution on Wind Turbine Performance". In: *Wind* 2.2 (June 2022), pp. 415–428. URL: http://dx.doi.org/10.3390/wind2020022.

[78] C. R. Vogel, R. H. J. Willden, and G. T. Houlsby. "Blade element momentum theory for a tidal turbine". In: *Ocean Engineering* 169 (2018), pp. 215–226.

[79] Martin O L Hansen. *Aerodynamics of wind turbines*. London: Routledge, 2015.

[80] Christian Bak, Alexander M Forsting, and Niels N Sorensen. "The influence of leading edge roughness, rotor control and wind climate on the loss in energy production". In: *Journal of Physics* 1618.5 (2020), p. 052050.

[81] Charlotte Bay Hasager, Jakob Ilsted Bech, and Christian Bak. "On leading edge erosion at turbine blades due to heavy rain and the application of erosion safe mode control to extend life". In: *3rd ACI Operations & Maintenance for Offshore Wind* (2019).

[82] Anna-Maria Tilg, Charlotte Bay Hasager, Hans-Jürgen Kirtzel, and Poul Hummelshøj. "Brief communication: Nowcasting of precipitation for leading-edge-erosion-safe mode". In: *Wind Energy Science* 5.3 (2020), pp. 977–981.

[83] Charlotte Bay Hasager, Flemming Vejen, Witold Robert Skrzypiński, and Anna-Maria Tilg. "Rain Erosion Load and Its Effect on Leading-Edge Lifetime and Potential of Erosion-Safe Mode at Wind Turbines in the North Sea and Baltic Sea". In: *Energies* 14.7 (2021), p. 1959.

[84] Marshall L. Buhl. *Ohio State University Wind Tunnel Tests*. NREL Wind Research, date accessed 01/10/2019. URL: https://wind.nrel.gov/airfoils/OSU_data/data/.

[85] M.J. Hoffmann, R. Reuss Ramsay, and G.m. Gregorek. "Effects of grit roughness and pitch oscillations on the NACA 4415 airfoil". In: *National Renewable Energy Laboratory* (1996).

[86] James L Tangler and Dan M Somers. *NREL airfoil families for HAWTs*. Tech. rep. National Renewable Energy Lab., Golden, CO (United States), 1995.

[87] R Reuss Ramsay, MJ Hoffmann, and GM Gregorek. *Effects of grit roughness and pitch oscillations on the S801 airfoil*. Tech. rep. NREL/TP-442-7818. National Renewable Energy Lab., Golden, CO (United States); The Ohio State Univ., Columbus, OH (United States), 1996.

[88] R Reuss Ramsay, MJ Hoffmann, and GM Gregorek. *Effects of grit roughness and pitch oscillations on the S810 airfoil*. Tech. rep. NREL/TP-442-7816. National Renewable Energy Lab., Golden, CO (United States); The Ohio State Univ., Columbus, OH (United States), 1996.

[89] Vestas. *V120-2.2 MW™*. Date accessed 20/09/2019. URL: https://www.vestas.com/en/products/2-mw-platform/v120-22_mw#!.

[90] *Siemens Gamesa Onshore Wind Turbine SG 2.6-114*. Date accessed 20/09/2019. URL: https://www.siemensgamesa.com/en-int/products-and-services/onshore/wind-turbine-sg-2-6-114.

[91] James F Manwell, Jon G McGowan, and Anthony L Rogers. *Wind energy explained: theory, design and application*. Chichester: John Wiley & Sons, 2010.

[92] Eduardo José Novaes Menezes, Alex Maurıcio Araújo, and Nadège Sophie Bouchonneau da Silva. "A review on wind turbine control and its associated methods". In: *Journal of Cleaner Production* 174 (2018), pp. 945–953.

[93] M Fréchet. "Sur la loi de probabilité de l'écart maximum". In: *Ann. Soc. Polon. Math.* 6 (Jan. 1927), pp. 93–116.

[94] Wind Data, The Crown Estate - Marine Data Exchange, date accessed 08/08/2019. URL: `http://www.marinedataexchange.co.uk/wind-data.aspx`.

[95] Jon Wieringa. "Updating the Davenport roughness classification". In: *Journal of Wind Engineering and Industrial Aerodynamics* 41.1-3 (1992), pp. 357–368.

[96] Terry Pratchett. *A Hat Full of Sky:(Discworld Novel 32)*. Random house, 2010.

[97] Christian Bak, Frederik Zahle, Robert Bitsche, Taeseong Kim, Anders Yde, Lars Christian Henriksen, Morten Hartvig Hansen, Jose Pedro Albergaria Amaral Blasques, Mac Gaunaa, and Anand Natarajan. "The DTU 10-MW reference wind turbine". In: *Danish Wind Power Research* (2013).

[98] Thomas Welte. "Deterioration and maintenance models for components in hydropower plants". In: *Fakultet for Ingeniørvitenskap og Teknologi* (2008).

[99] Gustavo de Novaes Pires Leite, Alex Maurıcio Araújo, and Pedro André Carvalho Rosas. "Prognostic techniques applied to maintenance of wind turbines: a concise and specific review". In: *Renewable and Sustainable Energy Reviews* 81 (2018), pp. 1917–1925.

[100] Fausto Pedro Garcıa Márquez, Jesús Marıa Pinar Pérez, Alberto Pliego Marugán, and Mayorkinos Papaelias. "Identification of critical components of wind turbines using FTA over the time". In: *Renewable Energy* 87 (2016), pp. 869–883.

[101] Lovisa Brandrud and Julie Krøgenes. "Aerodynamic Performance of the NREL S826 Airfoil in Icing Conditions". MA thesis. NTNU, 2017.

[102] Richard Hann, R Jason Hearst, Lars Roar Sætran, and Tania Bracchi. "Experimental and numerical icing penalties of an S826 airfoil at low Reynolds numbers". In: *Aerospace* 7.4 (2020), p. 46.

[103] Magnus K Vinnes and R Jason Hearst. "Aerodynamics of an airfoil with leading-edge icing". In: *Wind Energy* 24.8 (2020), pp. 795–811.

[104] Fabrizio De Gregorio, Antonio Ragni, M Airoldi, and GP Romano. "PIV investigation on airfoil with ice accretions and resulting performance degradation". In: *ICIASF 2001 Record, 19th International Congress on Instrumentation in Aerospace Simulation Facilities (Cat. No. 01CH37215)* (2001), pp. 94–105.

[105] Adriána Hudecz, Holger Koss, and Martin OL Hansen. "Ice accretion on wind turbine blades". In: *Proceedings of the 15th International Workshop on Atmospheric Icing of Structures (IWAIS XV), St. John's, NL, Canada* (2013), pp. 8–13.

[106] Clement Hochart, Guy Fortin, Jean Perron, and Adrian Ilinca. "Wind turbine performance under icing conditions". In: *Wind Energy* 11.4 (2008), pp. 319–333.

[107] Yiqiang Han, Jose Palacios, and Sven Schmitz. "Scaled ice accretion experiments on a rotating wind turbine blade". In: *Journal of Wind Engineering and Industrial Aerodynamics* 109 (2012), pp. 55–67.

[108] Linyue Gao, Yang Liu, and Hui Hu. "An experimental investigation on the dynamic ice accretion process over the surface of a wind turbine blade model". In: *9th AIAA atmospheric and space environments conference* (2017), p. 3582.

[109] *Leading Edge Erosion*. Sandia National Laboratories. Date accessed 01/05/2020. URL: `https://energy.sandia.gov/programs/renewable-energy/wind-power/blade-reliability/leading-edge-erosion/`.

[110] Beatriz Mendez and Xabier Munduate. "Study of distributed roughness effect over wind turbine airfoils performance using CFD". In: *33rd Wind Energy Symposium* (2015), p. 0994.

[111] Atmosphere to Electrons (A2e). 2017. lees/report.z01.00. Maintained by A2e Data Archive and Portal for U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy. DOI: 10.21947/1373097. Accessed: 01/05/2020.

[112] Mikko Tiihonen, Tuomas Jokela, Lasse Makkonen, and Geert-Jan Bluemink. "VTT icing wind tunnel 2.0". In: *Winterwind International Wind Energy Conference* (2016).

[113] Richard Hann, Kasper Borup, Artur Zolich, Kim Sorensen, Håvard Vestad, Martin Steinert, and Tor Johansen. *Experimental investigations of an icing protection system for UAVs*. Tech. rep. SAE Technical Paper, 2019.

[114] William Wright. "User's manual for LEWICE version 3.2". In: *NASA CR* 214255 (2008), p. 2008.

[115] Peter Michael Blasco. "An experimental and computational approach to iced wind turbine aerodynamics". MA thesis. Pennsylvania State University, 2015.

[116] P. Berrisford, D.P. Dee, P. Poli, R. Brugge, Mark Fielding, Manuel Fuentes, P.W. Kållberg, S. Kobayashi, S. Uppala, and Adrian Simmons. "The ERA-Interim archive Version 2.0". In: 1 (Nov. 2011), p. 23. URL: `https://www.ecmwf.int/node/8174`.

[117] Ronald Gelaro, Will McCarty, Max J Suárez, Ricardo Todling, Andrea Molod, Lawrence Takacs, Cynthia A Randles, Anton Darmenov, Michael G Bosilovich, Rolf Reichle, et al. "The modern-era retrospective analysis for research and applications, version 2 (MERRA-2)". In: *Journal of Climate* 30.14 (2017), pp. 5419–5454.

[118] Mark Decker, Michael A Brunke, Zhuo Wang, Koichi Sakaguchi, Xubin Zeng, and Michael G Bosilovich. "Evaluation of the reanalysis products from GSFC, NCEP, and ECMWF using flux tower observations". In: *Journal of Climate* 25.6 (2012), pp. 1916–1944.

[119] Jaume Ramon, Llorenç Lledó, Verónica Torralba, Albert Soret, and Francisco J Doblas-Reyes. "What global reanalysis best represents near-surface winds?" In: *Quarterly Journal of the Royal Meteorological Society* 145.724 (2019), pp. 3236–3251.

[120] Copernicus Climate Change Service (C3S). *ERA5: Fifth generation of ECMWF atmospheric reanalyses of the global climate.* Accessed on 09/02/20. 2017. URL: https://cds.climate.copernicus.eu/cdsapp#!/home.

[121] Dick P Dee, Sakari M Uppala, AJ Simmons, Paul Berrisford, P Poli, S Kobayashi, U Andrae, MA Balmaseda, G Balsamo, d P Bauer, et al. "The ERA-Interim reanalysis: Configuration and performance of the data assimilation system". In: *Quarterly Journal of the Royal Meteorological Society* 137.656 (2011), pp. 553–597.

[122] Shinya Kobayashi, Yukinari Ota, Yayoi Harada, Ayataka Ebita, Masami Moriya, Hirokatsu Onoda, Kazutoshi Onogi, Hirotaka Kamahori, Chiaki Kobayashi, Hirokazu Endo, et al. "The JRA-55 reanalysis: General specifications and basic characteristics". In: *Journal of the Meteorological Society of Japan. Ser. II* 93.1 (2015), pp. 5–48.

[123] Eugenia Kalnay, Masao Kanamitsu, Robert Kistler, William Collins, Dennis Deaven, Lev Gandin, Mark Iredell, Suranjana Saha, Glenn White, John Woollen, et al. "The NCEP/NCAR 40-year reanalysis project". In: *Bulletin of the American Meteorological Society* 77.3 (1996), pp. 437–472.

[124] Péter Kiss, László Varga, and Imre M Jánosi. "Comparison of wind power estimates from the ECMWF reanalyses with direct turbine measurements". In: *Journal of Renewable and Sustainable Energy* 1.3 (2009), p. 033105.

[125] Sakari M Uppala, PW Kållberg, AJ Simmons, U Andrae, V Da Costa Bechtold, M Fiorino, JK Gibson, J Haseler, A Hernandez, GA Kelly, et al. "The ERA-40 re-analysis". In: *Quarterly Journal of the Royal Meteorological Society* 131.612 (2005), pp. 2961–3012.

[126] S Hawkins, D Eager, and GP Harrison. "Characterising the reliability of production from future British offshore wind fleets". In: *IET Conference on Renewable Power Generation (RPG 2011)* (2011).

[127] Marek Kubik, David Brayshaw, and Phil Coker. "Reanalysis: an improved data set for simulating wind generation". In: *WREF 2012* (2012).

[128] Jon Olauson and Mikael Bergkvist. "Modelling the Swedish wind power production using MERRA reanalysis data". In: *Renewable Energy* 76 (2015), pp. 717–725.

[129] Abhnil A Prasad, Robert A Taylor, and Merlinde Kay. "Assessment of solar and wind resource synergy in Australia". In: *Applied Energy* 190 (2017), pp. 354–367.

[130] Gorm B Andresen, Anders A Søndergaard, and Martin Greiner. "Validation of Danish wind time series from a new global renewable energy atlas for energy system analysis". In: *Energy* 93 (2015), pp. 1074–1088.

[131] *Electricity Data Model Monthly Archive*. Accessed on 21/02/20. URL: https://www.nemweb.com.au/#mms-data-model.

[132] *Ofgem Renewables and CHP Register*. Accessed on 21/02/20. URL: https://www.renewablesandchp.ofgem.gov.uk/.

[133] *Form EIA-923 detailed data with previous form data (EIA-906/920)*. Accessed on 21/02/20. URL: https://www.eia.gov/electricity/data/eia923/.

[134] Gordon Hughes. "The performance of wind farms in the United Kingdom and Denmark". In: *Renewable Energy Foundation* (2012), p. 48.

[135] Christopher J Crabtree, Donatella Zappalá, and Simon I Hogg. "Wind energy: UK experiences and offshore operational challenges". In: *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy* 229.7 (2015), pp. 727–746.

[136] Aleksis Xenophon and David Hill. "Open grid model of Australia's National Electricity Market allowing backtesting against historic data". In: *Scientific Data* 5 (2018), p. 180203.

[137] Jerome H Friedman. "Greedy function approximation: a gradient boosting machine". In: *Annals of Statistics* (2001), pp. 1189–1232.

[138] Gustav Hellmann. *Über die Bewegung der Luft in den untersten Schichten der Atmosphäre*. Kgl. Akademie der Wissenschaften G. Reimer, 1919.

[139] AS Monin and AM Obukhov. "Dimensionless characteristics of turbulence in the surface layer". In: *Akad. Nauk SSSR, Geofiz. Inst., Tr* 24 (1954), pp. 163–187.

[140] Iain Staffell and Stefan Pfenninger. "Using bias-corrected reanalysis to simulate current and future wind power output". In: *Energy* 114 (2016), pp. 1224–1239.

[141] *Electricity Capacity Assessment*. Accessed on 02/06/21. URL: `https://www.ofgem.gov.uk/ofgem-publications/40203/electricity-capacity-assessment-2012.pdf`.

[142] Paritosh Bhattacharya and Rakhi Bhattacharjee. "A study on Weibull distribution for estimating the parameters". In: *Journal of Applied Quantitative Methods* 5.2 (2010), pp. 234–241.

[143] Emil Artin. *The gamma function*. Courier Dover Publications, 2015.

[144] Christophe Leys, Christophe Ley, Olivier Klein, Philippe Bernard, and Laurent Licata. "Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median". In: *Journal of Experimental Social Psychology* 49.4 (2013), pp. 764–766.

[145] Peter J Huber. *Robust statistics*. Vol. 523. New York: John Wiley & Sons, 2004.

[146] Peter J Rousseeuw and Christophe Croux. "Alternatives to the median absolute deviation". In: *Journal of the American Statistical Association* 88.424 (1993), pp. 1273–1283.

[147] Matti Niclas Scheu, Athanasios Kolios, Tim Fischer, and Feargal Brennan. "Influence of statistical uncertainty of component reliability estimations on offshore wind farm availability". In: *Reliability Engineering & System Safety* 168 (2017), pp. 28–39.

[148] Aiwina Heng, Sheng Zhang, Andy CC Tan, and Joseph Mathew. "Rotating machinery prognostics: State of the art, challenges and opportunities". In: *Mechanical Systems and Signal Processing* 23.3 (2009), pp. 724–739.

[149] Peter Tavner. "How are we going to make offshore wind farms more reliable?" In: *Supergen Wind* (2011).

[150] Frank J Brooks. "GE gas turbine performance characteristics". In: *GE Power Systems, Schenectady, NY* (2000).

[151] Michele M Rienecker, Max J Suarez, Ronald Gelaro, Ricardo Todling, Julio Bacmeister, Emily Liu, Michael G Bosilovich, Siegfried D Schubert, Lawrence Takacs, Gi-Kong Kim, et al. "MERRA: NASA's modern-era retrospective analysis for research and applications". In: *Journal of Climate* 24.14 (2011), pp. 3624–3648.

[152] *EMD ConWx Mesoscale Data*. Accessed on 23/09/21. URL: `https://www.emd-international.com/files/windpro2.9/EMDConWx_MesoScale_data.pdf`.

[153] Christoph Bollmeyer, JD Keller, Christian Ohlwein, S Wahl, Susanne Crewell, P Friederichs, A Hense, J Keune, S Kneifel, and I Pscheidt. "Towards a high-resolution regional reanalysis for the European CORDEX domain". In: *Quarterly Journal of the Royal Meteorological Society* 141.686 (2015), pp. 1–15.

[154] Iain Staffell. "Wind turbine power curves". In: *Imperial College London* 2 (2012).

[155] *UK Wind Speed Index*. Accessed on 30/05/21. URL: `http://web.archive.org/web/20080223103132/www.garradhassan.com/services/ghindex/`.

[156] Wolf-Gerrit Früh. "Long-term wind resource and uncertainty estimation using wind records from Scotland as example". In: *Renewable Energy* 50 (2013), pp. 1014–1026.

[157] Per Nielsen. *New Danish Wind energy index calculation*. Accessed on 01/06/21. URL: `http://www.vindstat.dk/PDF_sider/Windindex_DK_new_short.pdf`.

[158] SC Pryor and RJ Barthelmie. "Long-term trends in near-surface flow over the Baltic". In: *International Journal of Climatology: A Journal of the Royal Meteorological Society* 23.3 (2003), pp. 271–289.

[159] SC Pryor, Rebecca Jane Barthelmie, and JT Schoof. "Inter-annual variability of wind indices across Europe". In: *Wind Energy* 9.1-2 (2006), pp. 27–38.

[160] Morten Lybech Thøgersen, Maurizio Motta, Thomas Sørensen, and Per Nielsen. "Measure-correlate-predict methods: Case studies and software implementation". In: *European Wind Energy Conference & Exhibition* (2007), p. 10.

[161] Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. "Pearson correlation coefficient". In: Springer, 2009, pp. 1–4.

[162] David G Kleinbaum, Lawrence L Kupper, Azhar Nizam, and Eli S Rosenberg. *Applied regression analysis and other multivariable methods*. Cengage Learning, 2013.

[163] Graham Sinden. "Characteristics of the UK wind resource: Long-term patterns and relationship to electricity demand". In: *Energy Policy* 35.1 (2007), pp. 112–127.

[164] Udaya Bhaskar Gunturu and C Adam Schlosser. "Characterization of wind power resource in the United States". In: *Atmospheric Chemistry and Physics* 12 (2012), pp. 9687–9702.

[165] Alexandra Cosseron, Udaya Bhaskar Gunturu, and C Adam Schlosser. "Characterization of the wind power resource in Europe and its intermittency". In: *Energy Procedia* 40 (2013), pp. 58–66.

[166] Willow Hallgren, Udaya Bhaskar Gunturu, and Adam Schlosser. "The potential wind power resource in Australia: a new perspective". In: *PLOS one* 9.7 (2014), e99608.

[167] Sarah Becker, Bethany A Frew, Gorm B Andresen, Timo Zeyer, Stefan Schramm, Martin Greiner, and Mark Z Jacobson. "Features of a fully renewable US electricity system: Optimized mixes of wind and solar PV and transmission grid extensions". In: *Energy* 72 (2014), pp. 443–458.

[168] Matthias Huber, Desislava Dimkova, and Thomas Hamacher. "Integration of wind and solar power in Europe: Assessment of flexibility requirements". In: *Energy* 69 (2014), pp. 236–246.

[169] Tammy Jiang, Jaimie L Gradus, and Anthony J Rosellini. "Supervised machine learning: a brief primer". In: *Behavior Therapy* 51.5 (2020), pp. 675–687.

[170] Vladimir Nasteski. "An overview of the supervised machine learning methods". In: *Horizons* 4 (2017), pp. 51–62.

[171] David G Kleinbaum, K Dietz, M Gail, Mitchel Klein, and Mitchell Klein. *Logistic regression*. Springer, 2002.

[172] Trevor Hastie and Robert Tibshirani. "Discriminant adaptive nearest neighbor classification and regression". In: *Advances in Neural Information Processing Systems* 8 (1995).

[173] Richard O Duda, Peter E Hart, et al. *Pattern classification and scene analysis*. Vol. 3. Wiley New York, 1973.

[174] S Rasoul Safavian and David Landgrebe. "A survey of decision tree classifier methodology". In: *IEEE Transactions on Systems, Man, and Cybernetics* 21.3 (1991), pp. 660–674.

[175] William S Noble. "What is a support vector machine?" In: *Nature Biotechnology* 24.12 (2006), pp. 1565–1567.

[176] Rishabh Choudhary and Hemant Kumar Gianey. "Comprehensive review on supervised machine learning algorithms". In: *2017 International Conference on Machine Learning and Data Science (MLDS)* (2017), pp. 37–43.

[177] Leo Breiman. "Random forests". In: *Machine Learning* 45.1 (2001), pp. 5–32.

[178] J Ranstam and JA Cook. "LASSO regression". In: *Journal of British Surgery* 105.10 (2018), pp. 1348–1348.

[179] Clifton D Sutton. "Classification and regression trees, bagging, and boosting". In: *Handbook of Statistics* 24 (2005), pp. 303–329.

[180] Naphaporn Sirikulviriya and Sukree Sinthupinyo. "Integration of rules from a random forest". In: *International Conference on Information and Electronics Engineering* 6 (2011), pp. 194–198.

[181] Thais Mayumi Oshiro, Pedro Santoro Perez, and José Augusto Baranauskas. "How many trees in a random forest?" In: *International Workshop on Machine Learning and Data Mining in Pattern Recognition* (2012), pp. 154–168.

[182] Leo Breiman. "Bagging predictors". In: *Machine Learning* 24.2 (1996), pp. 123–140.

[183] Yoav Freund, Robert E Schapire, et al. "Experiments with a new boosting algorithm". In: *International Conference on Machine Learning* 96 (1996), pp. 148–156.

[184] Gang Wang, Jinxing Hao, Jian Ma, and Hongbing Jiang. "A comparative assessment of ensemble learning for credit scoring". In: *Expert Systems with Applications* 38.1 (2011), pp. 223–230.

[185] Tianqi Chen and Carlos Guestrin. "Xgboost: A scalable tree boosting system". In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2016), pp. 785–794.

[186] Wengang Zhang, Chongzhi Wu, Haiyi Zhong, Yongqin Li, and Lin Wang. "Prediction of undrained shear strength using extreme gradient boosting and random forest based on Bayesian optimization". In: *Geoscience Frontiers* 12.1 (2021), pp. 469–477.

[187] Tianqi Chen, Tong He, Michael Benesty, Vadim Khotilovich, Yuan Tang, Hyunsu Cho, Kailong Chen, et al. "Xgboost: extreme gradient boosting". In: *R package version 0.4-2* 1.4 (2015), pp. 1–4.

[188] Gérard Biau and Erwan Scornet. "A random forest guided tour". In: *Test* 25.2 (2016), pp. 197–227.

[189] Vladimir Svetnik, Andy Liaw, Christopher Tong, J Christopher Culberson, Robert P Sheridan, and Bradley P Feuston. "Random forest: a classification and regression tool for compound classification and QSAR modeling". In: *Journal of Chemical Information and Computer Sciences* 43.6 (2003), pp. 1947–1958.

[190] Anantha M Prasad, Louis R Iverson, and Andy Liaw. "Newer classification and regression tree techniques: bagging and random forests for ecological prediction". In: *Ecosystems* 9.2 (2006), pp. 181–199.

[191] Hal R Varian. "Big data: New tricks for econometrics". In: *Journal of Economic Perspectives* 28.2 (2014), pp. 3–28.

[192] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. "Scikit-learn: Machine Learning in Python". In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.

[193] Ron Kohavi. "A study of cross-validation and bootstrap for accuracy estimation and model selection". In: *Ijcai* 14.2 (1995), pp. 1137–1145.

[194] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An introduction to statistical learning*. Vol. 112. Springer, 2013, pp. 181–184.

[195] Oliver Schabenberger and Francis J Pierce. *Contemporary statistical models for the plant and soil sciences*. CRC Press, 2001.

[196] Andrej-Nikolai Spiess and Natalie Neumeyer. "An evaluation of R2 as an inadequate measure for nonlinear models in pharmacological and biochemical research: a Monte Carlo approach". In: *BMC pharmacology* 10.1 (2010), pp. 1–11.

[197] Rachel Carson. "Silent spring". In: Routledge, 2015, pp. 150–155.