




Investigating the Effect of Atmospheric Stratification on Offshore Wind Turbine Performance

How Does Thermal Stability Truly Dictate Energy Yield?

Dimitrios Gkoufas

Investigating the Effect of Atmospheric Stratification on Offshore Wind Turbine Performance

How Does Thermal Stability Truly Dictate
Energy Yield?

by

Dimitrios Gkoufas

Chairman of Thesis Committee:	Professor Dr. SJ (Simon) Watson
Thesis Committee Member:	Professor Dr. S.P. (Sara) Porchetta
Company Supervisor:	J. (Justin) Burstein
Faculties:	Faculty of Aerospace Engineering, Delft Faculty of Electrical Engineering, Mathematics and Computer Science, Delft

Cover:	Rampion Offshore Wind Farm, East Quay, Newhaven, UK by Nicholas Doherty
Style:	TU Delft Report Style

Summary

Offshore wind energy is expanding rapidly to meet global decarbonisation targets, and realising this transition, together with its socio-economic benefits, depends on large capital investments that are only committed when projects can demonstrate a reliable energy yield. Accurate energy yield assessments are therefore central: they secure competitive financing and limit financial risk. Yet standard power performance models typically assume idealised, neutral atmospheric conditions, while real offshore environments feature dynamic atmospheric boundary layers in which thermodynamic stability alters the wind profile and turbulence intensity. This mismatch introduces uncertainty into yield predictions under real-world stratification.

This research investigates how atmospheric stability influences the aerodynamic power performance of a modern offshore wind turbine, and evaluates the extent to which these thermodynamic effects can be decoupled from interconnected kinematic characteristics.

To isolate these aerodynamic impacts, the study processed continuous operational and meteorological data from a commercial wind farm located in the North Sea. The campaign used a forward-looking nacelle-mounted Light Detection and Ranging (Lidar) system to characterise the incoming wind field, alongside thermodynamic sensors to measure temperature, pressure, and humidity. The analysis was conducted within a wake-free sector, applying instrument availability and operational filters to ensure clean, undisturbed free-stream data. The dataset was then categorised into stable, neutral, and unstable regimes using kinematic proxies such as wind shear and Turbulence Intensity (TI), along with thermodynamic metrics, primarily the Bulk Richardson Number (Ri_b). Standardised empirical power curves were then constructed to quantify performance deviations between regimes.

The majority of the analysis relied on the Bulk Richardson Number classification. Because it directly incorporates thermodynamic information, it provided the most physically consistent basis for classification among the methods tested. Wind shear served as a viable low-cost alternative proxy, while Turbulence Intensity was less reliable due to its high volatility. In terms of aerodynamic performance, the data showed clear deviations in the partial-load operational region: unstable, convective conditions were consistently associated with overperformance, yielding an energy surplus of approximately +1.80% relative to the manufacturer's baseline, while stable, highly stratified conditions corresponded to a near-symmetric deficit of about -1.81%. In the rated power region, the turbine's active pitch control absorbed these stability-driven differences, bringing power anomalies close to zero. Seasonal variations in atmospheric stability, along with other aerodynamic phenomena, were also examined to provide a more complete evaluation of the turbine's performance across the year.

The research concludes that atmospheric stability significantly affects offshore wind turbine efficiency, and that standard neutral baseline power curves tend to overestimate energy production in stable offshore environments. The most plausible explanation is that unstable conditions promote strong convective mixing, producing a more uniform wind profile that improves kinetic energy extraction across the rotor swept area. To reduce financial risk, project developers should integrate stability-weighted probability distributions into pre-construction Energy Yield Assessments rather than relying on generic wind distributions. This methodology might also be suitable for resource assessment planning by deploying floating Lidar systems that capture local thermodynamic data. Future research should use vertically profiling Lidars to implement Rotor Equivalent Wind Speed (REWS) corrections and apply machine learning algorithms for continuous, multivariate stability classification.

Contents

Summary	i
Nomenclature	vii
1 Introduction	1
1.1 Background and Motivation	1
1.2 Problem Statement	2
1.3 Research Objectives and Questions	3
1.4 Research Scope and Limitations	3
1.5 Thesis Outline	4
2 Literature Review	5
2.1 Fundamentals of the Offshore Atmospheric Boundary Layer	5
2.1.1 The Vertical Wind Profile	6
2.1.2 Turbulence Intensity (TI)	6
2.1.3 Wind Veer	7
2.2 Atmospheric Stability and Stratification	7
2.2.1 The Physics of Buoyancy	8
2.2.2 Monin-Obukhov Similarity Theory (MOST)	9
2.2.3 Stability Classification Metrics	9
2.3 Wind Turbine Aerodynamics and Power Conversion	10
2.3.1 Aerodynamic Power Extraction and Efficiency Constraints	10
2.3.2 Environmental Dependencies	11
2.3.3 The Power Curve	11
2.4 Remote Sensing and Measurement Technology	12
2.4.1 Working Principles of the WindCube Nacelle Lidar	12
2.4.2 Nacelle-Mounted Lidar Challenges	12
2.5 State of the Art: Stability Impacts on Power Performance	13
2.5.1 Impact on Energy Yield: The Stability Conflict	13
2.5.2 Rotor Equivalent Wind Speed (REWS)	14
3 Methodology & Computational Framework	15
3.1 Site Characterisation and Measurement Campaign	16
3.1.1 Offshore Site and Layout	17
3.1.2 Instrumentation and Data Streams	17
3.2 Data Integration and Synchronisation	19
3.2.1 Temporal Alignment	19
3.3 Data Filtering and Quality Control	20
3.3.1 Environmental and Directional Filtering	20
3.3.2 Lidar and Meteorological Data Quality Control	22
3.3.3 Filtering Turbine Underperformance and Curtailment	23
3.4 Atmospheric Stability Classification	23
3.4.1 Wind Shear and Turbulence Intensity Proxies	23
3.4.2 Thermodynamic Metrics	24
3.5 Power Performance Assessment	26
3.5.1 Environmental Corrections and Normalisation	26
3.5.2 Power Curve Construction	28
3.5.3 Annual Energy Production (AEP) Estimation	28
4 Results	30
4.1 Characteristics of the Finalized Dataset	30

4.1.1	Data Retention and Attrition	30
4.1.2	Baseline Climatology	31
4.1.3	Lidar and Power Curve Scatter Validation	32
4.2	Atmospheric Stability Characterization	33
4.2.1	Classification via Kinematic Proxies	33
4.2.2	Classification via Thermodynamic Metrics	34
4.2.3	Comparative Evaluation	35
4.3	Impact of Stability on Power Performance	37
4.3.1	Environmental Normalisation and Correction	37
4.3.2	Baseline and Stability-Classified Power Curves	39
4.3.3	Quantification of Performance Deviations	40
4.3.4	Indicative Socio-Economic Implications	43
4.4	Advanced Performance Dynamics	44
4.4.1	Seasonal Variations in Stability and Performance	44
4.4.2	Upstream Flow and Blockage Effects	46
4.4.3	Isolating Kinematic Turbulence from Thermodynamic Drivers	49
5	Discussion	51
5.1	Methodological Constraints and Uncertainties	51
5.2	Industry Applications and Yield Optimisation	52
6	Conclusions & Recommendations	54
6.1	Conclusions	54
6.1.1	What are the critical data processing requirements to develop a framework that extracts atmospheric and operational variables from nacelle-mounted Lidar, meteorological instruments and SCADA datasets?	54
6.1.2	How can offshore atmospheric stability regimes be accurately classified utilising a combination of kinematic proxies and bulk thermodynamic metrics?	55
6.1.3	To what extent do the overall energy yield and power performance deviate from theoretical baselines across stable, neutral, and unstable atmospheric conditions?	55
6.2	Main Conclusion	55
6.3	Recommendations for Future Research	56
6.3.1	Enhanced Instrumentation and REWS Implementation	56
6.3.2	Advanced Multivariate Classification via Machine Learning	56
6.3.3	Comprehensive Uncertainty Quantification	57
6.3.4	Pre-Construction Application via Floating Remote Sensing	57
6.3.5	Geographic and Technological Transferability	57
	Acknowledgments	58
A	Data Validation	59
B	REWS Correction	61
C	Turbulence Intensity Normalization	64
	References	67

List of Figures

1.1	Evolution of offshore wind installed capacity (2010–2024). Continental analysis by Ruiz-Leo, Gil García, and Fernández-Guillamón (2025).	2
2.1	Visualization of the wind speed profile.	5
2.2	Comparison of turbulence intensity (TI) measurements from different LiDAR configurations against meteorological mast instruments.	7
2.3	Vertical gradient of potential temperature for primary atmospheric stability criteria.	8
2.4	Conceptual visualization of wind speed profiles under stable, neutral, and unstable atmospheric stratification.	9
2.5	Wind turbine power curve and operating regions.	11
2.6	WindCube Nacelle Lidar system illustrating the configuration of the four distinct Lines of Sight (LOS0, LOS1, LOS2, and LOS3) used for wind field reconstruction, Vaisala (2023).	12
3.1	Scheme of the proposed methodology, detailing the data pipeline from raw ingestion to final power deviation quantification.	16
3.2	Offshore wind farm cluster characteristics.	17
3.3	Primary environmental monitoring instrumentation deployed on the target turbine during the measurement campaign.	18
3.4	Cross-correlation plot of Lidar vs. SCADA wind speeds across selected distance gates, indicating peak correlation at zero lag.	19
3.5	Sectors excluded due to wakes of neighbouring and operating WTGs.	21
3.6	Scatter plot showing relative wind direction vs. yaw position with neighboring turbines marked by dashed red lines.	21
3.7	Radial speed turbulence intensity along the four line-of-sight vs. turbine yaw position for (a) the raw dataset and (b) the filtered dataset between 218° and 304.9° following artefact removal and quality control. Yaw position is averaged into 10° bins.	22
3.8	Diagram showing the rotor swept area divided into horizontal segments (A_i) to illustrate the REWS calculation.	26
3.9	Rayleigh distribution, probability density function (a) & cumulative distribution function (b), Ren (2025).	29
4.1	Scatter power curve illustrating the finalized retained data (Clean) against the cumulative excluded data points (Noise).	31
4.2	Baseline environmental climatology of the finalized measurement dataset.	32
4.3	Unclassified 10-minute scatter power curves of the filtered dataset, comparing remote sensing (a) and conventional anemometry (b).	32
4.4	Behaviour of kinematic stability proxies across the operational wind speed spectrum. Error bars indicate ± 1 standard deviation from the mean.	33
4.5	Comparison of thermodynamic metrics against hub-height wind speed. Includes filtered data for normal turbine operations.	35
4.6	Distribution of Turbulence Intensity (a) and Wind Shear Exponent (b) across the three atmospheric stability regimes defined by the Bulk Richardson Number. Boxes indicate the interquartile range (IQR), horizontal lines represent the median, and whiskers extend to the extremes of the distribution ($1.5 \times \text{IQR}$).	36
4.7	Confusion matrices evaluating the predictive performance of kinematic proxies against the Ri_b (BRN) baseline. Values represent the row-normalised recall percentages.	37
4.8	Analysis of site air density variance and its seasonal impact on wind speed normalisation.	37
4.9	Validation of the Turbulence Intensity normalisation process, illustrating the algorithm's effect on the transition region and the magnitude of power adjustments based on measured TI.	38

4.10	Baseline measured power curve and power coefficient (C_p) curve derived from the filtered dataset.	39
4.11	Comparison of wind turbine power curves disaggregated by atmospheric stability regime. Shaded areas represent ± 1 standard deviation.	40
4.12	Comparison of absolute power anomalies (ΔP) from the manufacturer baseline across four atmospheric classification methods: (a) Wind Shear, (b) Turbulence Intensity, (c) Bulk Richardson Number, and (d) Obukhov length. Shaded envelopes denote the 95% Confidence Interval. The vertical dashed line indicates the normalised rated wind speed.	41
4.13	Extrapolated Annual Energy Production (AEP) by stability regime across standard annual average wind speeds. All values are normalised against the manufacturer's theoretical AEP at $V_{ave} = 8$ m/s.	42
4.14	Seasonal stability characterization and underlying thermodynamic drivers at the offshore site.	44
4.15	Seasonal comparison of Annual Energy Production (AEP) across stability classes. Values are normalised against the theoretical yield derived from the OEM reference power curve.	45
4.16	Normalised upstream wind speed profiles across varying distance gates, categorized by stability regime. The panels represent: (a) partial load region, (b) transition region, (c) above-rated wind speed region, and (d) all available data. Profiles are normalised against the 407 m reference gate.	46
4.17	Theoretical vertical shear profiles by stability regime, overlaid with Lidar beam heights and rotor span.	48
4.18	Correlation coefficient between the Lidar and nacelle SCADA hub-height wind speeds at zero time lag, plotted against Lidar distance gate for all 20 gates (57–657 m).	48
4.19	Power deviation as a fraction of the theoretical OEM yield, mapped across Normalised Wind Speed and Turbulence Intensity. Cool colours (blue) indicate overperformance; warm colours (red) indicate underperformance.	49
A.1	Validation and cross-check analysis between nacelle anemometer and LiDAR measurements at the 407 m distance gate. Background data points represent the raw 10-minute records, while the highlighted subset shows the remaining data after applying the filtering criteria.	59
A.2	Wind speed validation between SCADA and LiDAR measurements, segmented by atmospheric stability regime. Atmospheric stability was classified using the Bulk Richardson Number (BRN). Each panel displays the filtered operational dataset alongside its corresponding linear regression fit and coefficient of determination (R^2).	60
B.1	Geometric visualization of the nacelle-mounted Lidar beams. At a horizontal measurement distance of $L = 2.5D$, the $\pm 5^\circ$ beam elevation angles result in a vertical offset (Δz) from the hub height (H). Both physical measurement heights (z_{low} and z_{high}) fall within the turbine's rotor swept area.	61
B.2	Impact analysis of the implemented Wind Shear (REWS) correction. (Left) Scatter comparison between the standard Hub Height Wind Speed and the Rotor Equivalent Wind Speed (both normalized to the 13 m/s rated wind speed). (Right) Histogram detailing the distribution of the absolute wind shear correction impact ($\Delta V = V_{eq} - V_{hub}$) in m/s, demonstrating the mean shift applied to the dataset's velocity profile.	63
C.1	First-approach approximation for the initial zero-turbulence power curve parameters IEC (2017).	65

List of Tables

3.1	Classification thresholds for atmospheric stability based on kinematic proxies.	24
3.2	Classification thresholds for atmospheric stability based on thermodynamic metrics.	25
4.1	Data Filtering Summary	31
4.2	Percentage distribution of 10-minute data among atmospheric stability regimes derived from kinematic proxies.	34
4.3	Percentage distribution of 10-minute data among atmospheric stability regimes derived from thermodynamic metrics.	35
4.4	General statistics of the filtered, free-stream dataset used for power curve construction.	39
4.5	Percentage deviation of extrapolated AEP by stability regime relative to the OEM reference at $V_{ave} = 8$ m/s, illustrating the sensitivity of yield results to different normalisation and kinematic correction methodologies.	42
4.6	Indicative revenue sensitivity by stability regime (€/yr per GWh of baseline annual yield) at $V_{ave} = 8$ m/s, assuming a wholesale price of €82/MWh. The adopted air-density and TI correction is shown in bold.	43
4.7	Seasonal distribution of atmospheric stability regimes and their corresponding mean wind shear (α).	46
4.8	Lidar measurement bias analysis at the 2.5D reference gate.	48

Nomenclature

Abbreviations

Abbreviation	Definition
ABL	Atmospheric Boundary Layer
AEP	Annual Energy Production
BRN	Bulk Richardson Number
CI	Confidence Interval
CNR	Carrier-to-Noise Ratio
CW	Continuous-Wave Lidar
ERA5	ECMWF Reanalysis v5
EYA	Energy Yield Assessment
FLiDAR	Floating Lidar Systems
GMM	Gaussian Mixture Model
HWS	Hub-Height Wind Speed
IEC	International Electrotechnical Commission
IQR	Interquartile Range
LCOE	Levelized Cost of Energy
LES	Large Eddy Simulations
Lidar	Light Detection and Ranging
LOS	Line of Sight
MOST	Monin-Obukhov Similarity Theory
NTF	Nacelle Transfer Function
OEM	Original Equipment Manufacturer
REWS	Rotor Equivalent Wind Speed
RH	Relative Humidity
RWS	Radial Wind Speed
SCADA	Supervisory Control and Data Acquisition
SEM	Standard Error of the Mean
SST	Sea Surface Temperature
TI	Turbulence Intensity
WRF	Weather Research & Forecasting
WTG	Wind Turbine Generator

Symbols

Symbol	Definition	Unit
A	Total rotor swept area	[m ²]
A_i	Area of the i -th horizontal segment	[m ²]
$B_{10 \text{ min}}$	Air pressure corrected to hub height averaged over 10 min	[Pa]
C_p	Power coefficient	[-]
D_n	Rotor diameter of the neighbouring, wake-producing turbine	[m]
$F(V)$	Rayleigh cumulative probability function	[-]
g	Gravitational acceleration	[m/s ²]
$g(z)$	Area of a circular segment integrated up to height z	[m ²]
H	Turbine hub height	[m]
I_{measured}	Measured turbulence intensity	[-]
I_{ref}	Reference turbulence intensity	[-]
L	Monin-Obukhov length	[m]
L_n	Absolute distance from the Lidar to the neighbouring tower	[m]
N	Total number of bins	[-]
N_h	Total number of hours in one year (8760)	[-]
N_i	Total number of 10-minute datasets in bin i	[-]
P	Aerodynamic power extracted by the turbine	[kW]
$P_{10 \text{ min}}$	Raw measured active power averaged over 10 min	[kW]
$P_{\text{available}}$	Total kinetic power available in the free-stream wind	[kW]
P_i	Normalized and averaged power output in bin i	[kW]
$P_{I,\text{norm}}$	TI-normalized power output	[kW]
P_r	Rated power	[kW]
P_{sat}	Saturation vapour pressure	[Pa]
P_{sim}	Simulated power	[kW]
P_{vap}	Actual vapour pressure	[Pa]
P_w	Vapour pressure	[Pa]
R	Rotor radius	[m]
R_0	Gas constant of dry air	[J/(kg·K)]
R_b	Adjusted measurement distance of the Lidar	[m]
R_{conf}	Selected Lidar measurement range (distance gate)	[m]
R_w	Gas constant of water vapour	[J/(kg·K)]
Ri_b	Bulk Richardson Number	[-]
T	Absolute temperature	[K]
T_a	Ambient air temperature	[K]
T_s	Sea surface temperature	[K]
$T_{10 \text{ min}}$	Absolute air temperature averaged over 10 min	[K]
$U(z)$	Wind speed at height z	[m/s]
$\bar{U}$	Mean longitudinal horizontal wind speed	[m/s]
U_{ref}	Reference wind speed	[m/s]
u_*	Friction velocity	[m/s]
V	Free-stream wind speed	[m/s]
$\bar{V}$	Mean lateral horizontal wind speed	[m/s]
V_{ave}	Annual average wind speed at hub height	[m/s]
$V_{\text{cut-in}}$	Cut-in wind speed	[m/s]
$V_{\text{cut-out}}$	Cut-out wind speed	[m/s]
V_{eq}	Rotor equivalent wind speed	[m/s]
V_i	Measured / averaged wind speed in bin i	[m/s]
V_{norm}	Normalized wind speed	[-]
V_r	Rated wind speed	[m/s]
w	Water-air mixing ratio	[-]

Symbol	Definition	Unit
w'	Vertical wind speed fluctuation	[m/s]
z	Height	[m]
z_0	Aerodynamic surface roughness length	[m]
z_m	Geometric mean measurement height	[m]
z_{ref}	Reference height	[m]
α	Wind shear exponent	[-]
β_{max}	Largest horizontal angle between any two beam positions	[°]
ΔP	Power deviation	[kW]
$\Delta \bar{U}$	Difference in longitudinal mean horizontal wind speed	[m/s]
$\Delta \bar{V}$	Difference in lateral mean horizontal wind speed	[m/s]
Δy	Vertical integration displacement ($z - H$)	[m]
Δz	Vertical height interval	[m]
$\Delta \bar{\theta}_v$	Difference in virtual potential temperature	[K]
ζ	Dimensionless stability parameter (z/L)	[-]
θ_{wake}	Excluded wind direction sector due to wakes	[°]
$\bar{\theta}_v$	Mean virtual potential temperature	[K]
θ'_v	Virtual potential temperature fluctuation	[K]
κ	von Kármán constant	[-]
ρ	Air density	[kg/m ³]
ρ_0	Reference air density	[kg/m ³]
$\rho_{10 \text{ min}}$	Derived 10-minute averaged air density	[kg/m ³]
σ_u	Standard deviation of the horizontal wind speed	[m/s]
Φ	Relative humidity	[-]

Introduction

The deployment of multi-megawatt offshore wind turbines is accelerating to meet global decarbonisation targets. However, as these large structures push further into the marine atmospheric environment, the thermodynamic behaviour of the boundary layer introduces notable deviations in expected energy yield. This thesis investigates the relationship between offshore atmospheric stability, localised boundary layer aerodynamics, and the resulting power performance of a modern offshore wind turbine. By processing remote sensing data to isolate these environmental drivers, this research aims to bridge the gap between meteorological physics and aerodynamic power conversion.

1.1. Background and Motivation

Offshore wind energy has emerged as a cornerstone of the global transition towards sustainable and decarbonised power generation (Figure 1.1). Driven by the urgent need to mitigate climate change and the advantage of capturing stronger, more consistent wind resources at sea, the industry is currently scaling quickly Martynowicz, Ślimak, and Kumsa (2026). Modern offshore wind farms are evolving from supplementary power sources towards becoming base-load contributors to national energy grids, as emerging projects and innovations in variable power management continue to bridge this gap UNFCCC (n.d.). This expansion brings socio-economic responsibilities, as entire regional economies and grid balancing authorities increasingly rely on the dependable, predictable performance of these large offshore installations.

To sustain the financial viability of these assets and support the wider energy transition, accurately predicting their lifetime energy yield is important. The development of a modern offshore wind farm requires multi-billion-euro capital investments, estimated at roughly €3–€4 million per MW (e.g., €30+ million for an 8 MW unit) Delfos (2025), meaning that even marginal deviations in expected power output can substantially alter a project's financial risk profile. Precise energy yield assessments are directly tied to securing competitive project financing and determining the true investment potential for industry stakeholders. Accurate yield prediction also ensures grid reliability by allowing operators to balance supply and demand efficiently, a systems-level capability that ultimately lowers the Levelized Cost of Energy (LCOE) and delivers tangible socio-economic benefits to end consumers.




Figure 1.1: Evolution of offshore wind installed capacity (2010–2024). Continental analysis by Ruiz-Leo, Gil García, and Fernández-Guillamón (2025).

Achieving this precision in energy yield prediction requires a fundamental understanding of the complex physical environment in which these turbines operate: the Atmospheric Boundary Layer (ABL). Unlike the idealised, neutral conditions often assumed in standard power performance models, the real offshore environment is highly dynamic and governed by varying atmospheric stability regimes. The thermodynamic interactions between the sea surface and the air mass above it alter the shape of the vertical wind profile. These interactions define the flow characteristics that impact turbine performance, dictating the wind shear, veer, and turbulence intensity that the rotor encounters. Understanding how these physical stratifications alter the kinetic energy available to the turbine is therefore the first step in explaining the observed discrepancies between expected yield and actual power generation.

1.2. Problem Statement

Isolating the distinct impact of atmospheric stability on wind turbine power performance presents a considerable analytical challenge due to the strong interconnection of meteorological variables. Atmospheric stability fundamentally dictates the vertical exchange of momentum and heat within the boundary layer, which in turn modulates both the vertical wind shear profile and the Turbulence Intensity (TI). For instance, a stable atmosphere naturally suppresses turbulent mixing (yielding low TI) and at the same time generates high vertical wind shear; these characteristics almost always fluctuate together. As a result, there is no straightforward methodology to decouple these heavily correlated metrics and quantify the isolated effect of thermodynamic stability on the aerodynamic efficiency of the rotor.

This inherent inability to easily disentangle stability from shear and turbulence has led to an unresolved conflict in the current literature. Because researchers often assess varying combinations of these intertwined proxies, typically under undisturbed flow conditions, studies reach contradicting conclusions on how stable atmospheric regimes impact energy yield. On one hand, research by Dörenkämper et al. (2014) indicates that stable stratification leads to a reduction in harvested energy, up to a 15% deficit, when compared to unstable and neutral conditions. By contrast, findings from St. Martin et al. (2016) and Berg (2008) report the opposite, observing that the enhanced wind shear associated with stable atmospheres yields notably higher power output at wind speeds approaching rated power. This contradiction points to a gap in the industry's analytical framework: without a clear method to isolate and quantify the specific physical drivers behind these power deviations, accurately predicting operational energy yield in complex offshore environments remains uncertain.

1.3. Research Objectives and Questions

The primary objective of this research is to isolate and quantify the specific impacts of atmospheric stability on offshore wind turbine power performance, independent of interconnected flow characteristics. Building upon the analytical challenges identified in the current literature, this study seeks to decouple the underlying thermodynamic stability of the marine boundary layer from site-specific kinematic variables. By isolating these driving physical mechanisms, this research aims to provide a clearer, more accurate representation of how the atmosphere dictates energy yield. The primary research question guiding this thesis is therefore formulated as follows:

Main Research Question

How does atmospheric stability influence the aerodynamic power performance of a modern offshore wind turbine, and to what extent can these thermodynamic effects be decoupled from interconnected kinematic characteristics?

To answer this question, the research methodology is structured around three foundational sub-questions. These supporting questions dictate the progression of the analytical pipeline, starting with the ingestion and filtering of raw operational data and culminating in the construction of normalised power curves. Addressing these specific technical milestones is necessary to ensure the validity of the final performance assessment:

- What are the critical data processing requirements to develop a framework that extracts atmospheric and operational variables from nacelle-mounted Lidar, meteorological instruments and SCADA datasets?
- How can offshore atmospheric stability regimes be accurately classified utilising a combination of kinematic proxies and bulk thermodynamic metrics?
- To what extent do the overall energy yield and power performance deviate from theoretical baselines across stable, neutral, and unstable atmospheric conditions?

Using a dataset that combines high-resolution nacelle-mounted Lidar, meteorological instruments and operational SCADA measurements, this thesis offers a unique, data-driven perspective on offshore aerodynamic performance. Given the conflicting observations in current literature, this research proceeds under the hypothesis that neither opposing perspective is entirely incorrect; rather, the perceived contradictions stem from confounding variables, such as site topography (onshore, offshore, or coastal) and instrumentation limitations. It is hypothesised that by applying environmental and operational filtering, the impact of atmospheric stratification will emerge as a distinct, quantifiable deviation from standard manufacturers' baselines. Decoupling these intertwined variables advances the scientific understanding of turbine-atmosphere interactions, while providing insights into offshore yield optimisation and wind resource assessments across the industry.

1.4. Research Scope and Limitations

The scope of this research is defined by the spatial boundaries of the measurement campaign and the specific instrumentation utilised to capture the offshore boundary layer. This study focuses on operational data collected from a commercial offshore wind farm located in the North Sea. The performance assessment is bounded by the data acquired from a Vaisala WindCube nacelle-mounted Lidar, thermodynamic transition-piece sensors, and the target turbine's internal Supervisory Control and Data Acquisition (SCADA) system. Methodologically, the scope is restricted to categorising atmospheric stability through a combination of empirical kinematic proxies and theoretical bulk thermodynamic metrics. By processing and filtering this synchronised dataset, the research isolates the free-stream flow to construct normalised power curves that reflect the turbine's aerodynamic response to these distinct environmental regimes.

While the methodology provides a robust framework for this performance assessment, several instrumental limitations inherent to the measurement campaign must be acknowledged. First, the campaign lacks high-frequency sonic anemometers, preventing the direct eddy-covariance measurement of turbulent heat fluxes required to execute the core calculations of Monin-Obukhov Similarity Theory (MOST). Second, the reliance on a single nacelle-mounted remote sensing device introduces measurement uncertainties, specifically regarding geometric beam spread and horizontal flow homogeneity, which are evaluated during the analysis. Finally, while the data anal-

ysis examines aerodynamic phenomena such as the upstream blockage effect, wake interactions generated by neighbouring wind turbines are filtered out and remain beyond the scope of the performance assessment.

1.5. Thesis Outline

To address the research questions and objectives, this thesis is structured into the following chapters:

Chapter 2: Literature Review establishes the theoretical foundation of the research. It explores the physics of the offshore Atmospheric Boundary Layer, aerodynamic power conversion principles, and the measurement constraints of pulsed Doppler Lidar. It culminates in a review of the current academic conflict regarding stability-driven power deviations.

Chapter 3: Methodology and Computational Framework details the analytical pipeline. This includes site characterisation, temporal synchronisation, directional wake filtering, and the mathematical framework used to classify atmospheric stability and normalise turbine power output according to environmental standards.

Chapter 4: Results presents the findings of the data analysis. It covers the characteristics of the finalised dataset, evaluates the atmospheric stability classifications, quantifies the aerodynamic power deviations and expected Annual Energy Production (AEP), and explores advanced performance dynamics such as seasonal variations and upstream blockage effects.

Chapter 5: Discussion contextualises the research findings within the broader industry. It addresses the inherent methodological constraints and uncertainties of the measurement campaign, and explores the practical applications of the results for refining energy yield assessments and optimising future offshore campaigns.

Chapter 6: Conclusions and Recommendations synthesizes the primary findings of the research, providing direct answers to the research questions. It outlines the broader implications for offshore wind farm yield assessments and suggests targeted avenues for future research to further optimise remote sensing and power prediction methodologies.

2

Literature Review

This chapter establishes the theoretical foundation and reviews the current state-of-the-art research on the aerodynamic impact of atmospheric stability on offshore wind turbine performance. Because accurately quantifying these impacts requires synthesising meteorology, aerodynamics, and remote sensing, this review is structured to bridge these distinct disciplines. The chapter begins by defining the fundamental kinematic physics of the offshore Atmospheric Boundary Layer (ABL) and the thermodynamic principles governing atmospheric stratification. The theoretical mechanics of wind turbine power conversion are then detailed, followed by a critical examination of nacelle-mounted Lidar technology and its spatial measurement limitations. Finally, the review ends with an analysis of contemporary literature on energy yield deviations under varying stability regimes, setting out the prevailing academic conflicts and methodological challenges that this thesis seeks to resolve.

2.1. Fundamentals of the Offshore Atmospheric Boundary Layer

While the atmosphere is stratified into multiple macroscopic layers, the main area of interest for wind energy is the Atmospheric Boundary Layer (ABL), which constitutes the lowest portion of the troposphere. This layer is distinguished by the frictional drag exerted by the Earth's surface, creating a vertical wind speed gradient that extends from zero velocity at the aerodynamic surface boundary to the geostrophic wind speed at the top of the ABL.




Figure 2.1: Visualization of the wind speed profile.

In the offshore environment, the specific shape and behaviour of this vertical profile are dictated by the aerodynamic roughness of the sea surface and the thermodynamic stability resulting from air-sea temperature differen-

tials. These interactions define the key flow characteristics that govern turbine performance, which are detailed in the following subsections: the vertical wind profile, turbulence intensity, and wind veer.

2.1.1. The Vertical Wind Profile

The vertical wind speed gradient is mathematically quantified by the wind shear exponent (α), which defines the rate at which wind velocity increases with altitude. To model this profile, two primary mathematical approaches are utilised in literature: the Logarithmic Law (Log Law) and the Power Law. The Log Law is derived from fundamental boundary layer physics Monin and Obukhov (1954) and is expressed as:

$$U(z) = \frac{u_*}{\kappa} \ln \left(\frac{z}{z_0} \right) \quad (2.1)$$

where u_* is the friction velocity, κ is the von Kármán constant, z is the evaluation height, and z_0 is the aerodynamic surface roughness length.

Alternatively, the Power Law, which is credited most widely to Hellmann (1916), offers an empirical relationship based on a known reference height:

$$\frac{U(z)}{U_{ref}} = \left(\frac{z}{z_{ref}} \right)^\alpha \quad (2.2)$$

where U_{ref} is the measured wind speed at a specific reference height z_{ref} . While the Log Law provides a robust theoretical foundation, the Power Law is predominantly adopted in offshore wind energy assessments due to its simplicity and alignment with IEC standards. Its primary advantage lies in eliminating the reliance on the roughness length (z_0), a parameter that is notoriously difficult to estimate accurately over varying sea states. The Power Law is therefore widely accepted for extrapolating wind speeds to modern hub heights, providing a practical balance between accuracy and computational efficiency.

The shear exponent (α) in this relationship serves as a critical indicator of the atmospheric state. In general terms, high shear values indicate stable stratification with suppressed mixing and high vertical gradient of wind velocity, whereas low shear values suggest unstable conditions with significant vertical mixing and reduction of velocity gradient; thus a flattened profile is expected. This contrast in profile shape is illustrated in Figure 2.1a, and the physical drivers behind these regimes will be further analysed in Section 2.2.

2.1.2. Turbulence Intensity (TI)

Turbulence in the Atmospheric Boundary Layer originates primarily from two distinct physical mechanisms: mechanical and thermal production. Mechanical turbulence is generated by the frictional wind shear as the air flows over the aerodynamic roughness of the sea surface, with its intensity scaling with the magnitude of the mean wind speed (Stull 1988). As detailed by Stull (1988), this shear production converts kinetic energy from the mean flow into turbulent fluctuations.

Thermal turbulence, on the other hand, is driven by buoyancy forces resulting from temperature differentials between the sea surface and the air mass. Following Monin-Obukhov Similarity Theory (Monin and Obukhov 1954), this buoyancy becomes the dominant driver of boundary layer mixing during unstable atmospheric conditions, a relationship empirically validated by Businger et al. (1971).

To quantify these fluctuations, Turbulence Intensity (TI) acts as a normalised measure of the wind speed variance relative to the mean flow over a specified period. It is mathematically defined as the ratio of the standard deviation of the horizontal wind speed (σ_u) to the mean wind speed ($\bar{U}$), typically calculated over a 10-minute interval:

$$TI = \frac{\sigma_u}{\bar{U}} \quad (2.3)$$

Analogous to the wind shear exponent discussed in Section 2.1.1, the magnitude of TI can serve as a distinct proxy for classifying atmospheric stability. In general terms, high TI values correspond to unstable atmospheric conditions driven by thermal buoyancy and mixing, while low TI values indicate stable stratification where turbulent exchange is suppressed.

Despite its utility, accurately quantifying TI using remote sensing technology presents considerable methodological challenges compared to conventional met masts. Standard cup or sonic anemometers capture high-frequency

turbulent fluctuations at a discrete, localised point (Figure 2.2). In contrast, Lidars measure wind speed by averaging Doppler shifts over a spatial probe volume (often tens of metres in length) along the beam path. This spatial averaging acts as a low-pass filter, attenuating the high-frequency components of the turbulence spectrum and leading to systemic underestimation of the true variance Borraccino and Wagner (n.d.). This measurement process can also be skewed by the operational movements of the wind turbine itself. Dynamic nacelle sway and pitch motions induced by aeroelastic tower deflection introduce non-atmospheric velocity components into the line-of-sight measurements, artificially inflating the apparent variance Rozenn Wagner and Davoust (2013). The combination of volume averaging damping the true variance and nacelle motion adding artificial noise often leads to a complex distortion of the true boundary layer TI. Specific scanning configurations or advanced post-processing algorithms are therefore required to accurately reconstruct the turbulence spectrum Peña, Yankova, and Mallini (n.d.).




Figure 2.2: Comparison of turbulence intensity (TI) measurements from different LiDAR configurations against meteorological mast instruments.

2.1.3. Wind Veer

Wind veer is defined as the change in wind direction with increasing altitude across the vertical profile. This directional shear is primarily driven by the balance between the Coriolis force and frictional surface drag, resulting in a rotation of the wind vector known as the Ekman spiral. In the Northern Hemisphere, this phenomenon typically manifests as a clockwise turning of the wind (veering) as altitude increases from the surface toward the free atmosphere, where the flow eventually aligns with the geostrophic wind direction.

The magnitude of this directional shear is not constant but is strongly modulated by the atmospheric stability regime. Empirical studies, such as those conducted by Zhang, Walker, and Markfort (2025), have reported significant deviations in stable atmospheres, demonstrating that wind veer can exceed 30° across the boundary layer under strongly stratified conditions. This variation is important for modern multi-megawatt offshore rotors; strong veer introduces a continuous misalignment between the blade chord and the inflow wind vector. Because the wind direction changes with height, the turbine cannot physically yaw to align perfectly with the entire wind field at once. This localised yaw misalignment alters the aerodynamic angle of attack at different blade elevations, reducing the total kinetic energy extraction and suppressing available power. To mathematically account for this loss, industry guidelines such as the Annex Q of the IEC (2017) introduce specific correction methodologies, using the Rotor Equivalent Wind Speed (REWS) to integrate these varying directional vectors into a single, representative inflow parameter.

2.2. Atmospheric Stability and Stratification

A central component of this research involves categorising the offshore wind resource into distinct stability regimes to isolate their impact on turbine performance. Because atmospheric stability governs the vertical exchange of momentum and heat within the boundary layer, a sound understanding of the underlying theory is essential. This section establishes the theoretical framework for stratification and details the specific metrics used to quantify these regimes.

2.2.1. The Physics of Buoyancy

The physical concept of atmospheric stability is rooted in the buoyancy forces that either suppress or encourage the vertical motion of air parcels. To quantify this, the concept of adiabatic lapse rates is utilised, defined as the rate at which a parcel of air cools as it rises and expands without exchanging heat with its surroundings (Wallace and Hobbs 2006). As detailed by Stull (1988), evaluating these forces is simplified by using potential temperature (θ). Potential temperature represents the temperature an air parcel would possess if brought adiabatically to a standard reference pressure (usually 1000 hPa). As shown in Figure 2.3, by evaluating the vertical gradient of potential temperature ($\partial\theta/\partial z$), one can determine whether a vertically displaced air parcel will return to its original position (stable) or continue to accelerate away from it (unstable).

Based on these thermodynamic interactions, the atmosphere is generally classified into three primary stability regimes: Stable, Neutral, and Unstable. While some studies Araújo Da Silva et al. (2022) and Anantharaman et al. (2025) expand this classification to include finer gradations such as "very stable," "near neutral," or "very unstable," the tripartite division remains the standard convention in wind energy literature for evaluating aerodynamic performance.




Figure 2.3: Vertical gradient of potential temperature for primary atmospheric stability criteria.

Stable Stratification

Stable conditions occur when the potential temperature increases with height ($\partial\theta/\partial z > 0$), typically arising when a warm air mass advects over a colder sea surface. In this regime, any vertically displaced air parcel becomes cooler and denser than its surrounding environment, creating a negative buoyancy force that pushes it back to its original elevation. This suppression of vertical motion dampens turbulent kinetic energy and decouples the atmospheric layers, frequently resulting in high wind shear and pronounced directional veer.

Unstable Stratification

By contrast, unstable stratification is defined by a decrease in potential temperature with height ($\partial\theta/\partial z < 0$), a phenomenon commonly occurring when cold continental air flows over warmer waters. Here, a rising air parcel remains warmer and less dense than the surrounding environment, generating positive buoyancy that accelerates its upward motion. This regime promotes strong vertical convective mixing and thermal turbulence, which homogenises the wind speed profile and typically results in low shear and veer.

Neutral Stratification

Neutral conditions represent the intermediate boundary state where the potential temperature remains constant with height ($\partial\theta/\partial z \approx 0$). In this scenario, a displaced air parcel maintains the same density as its environment and experiences zero net buoyancy force. This regime is conventionally assumed in standard power performance guidelines, such as the IEC (2022) standards, and acts as the baseline for OEM reference power curves. It is characterised by wind profiles dominated purely by mechanical friction (surface roughness) rather than thermal buoyancy effects.

These three thermodynamic regimes manifest as distinctly different velocity profiles. As depicted conceptually in Figure 2.4, stable atmospheres produce a steep, highly sheared profile, whereas the thermal mixing of unsta-

ble atmospheres results in a flatter, more uniform velocity curve, with neutral conditions resting as the baseline between these two.




Figure 2.4: Conceptual visualization of wind speed profiles under stable, neutral, and unstable atmospheric stratification.

2.2.2. Monin-Obukhov Similarity Theory (MOST)

Monin-Obukhov Similarity Theory (MOST) constitutes the fundamental framework for describing the non-linear vertical profiles of wind speed and temperature within the atmospheric surface layer. The theory assumes that the mean flow and turbulence characteristics are governed by a specific set of surface layer scales, principally the friction velocity (u_*) and the surface turbulent heat flux.

Central to this framework is the Obukhov length (L), a scaling parameter that represents the theoretical height at which the production of turbulence by buoyancy (thermal) forces exactly equals that produced by shear (mechanical) forces. This length scale is mathematically defined as:

$$L = -\frac{u_*^3 \bar{\theta}_v}{\kappa g (w' \theta'_v)} \quad (2.4)$$

where κ is the von Kármán constant, g is the gravitational acceleration, $\bar{\theta}_v$ is the mean virtual potential temperature, and $\overline{w' \theta'_v}$ represents the kinematic surface heat flux.

While MOST is widely regarded as the most rigorous methodology for characterising stability, direct calculation of the Obukhov length (L) necessitates high-frequency heat flux data often absent in industrial campaigns. Therefore, precise derivation is typically restricted to high-fidelity environments, such as Large Eddy Simulations (LES) Cao et al. (2025), Souaiby and Porté-Agel (2025) or experimental setups employing sonic anemometers to capture the required turbulent fluctuations via eddy covariance Bardal et al. (2018), St. Martin et al. (2016).

The application of MOST offshore is also strictly valid only within the surface layer, typically the lowest 10% of the ABL, the absolute height of which varies considerably with the active stability conditions. This poses distinct theoretical challenges for accurate flow extrapolation up to the massive rotor heights of modern multi-megawatt turbines, which frequently extend well beyond this surface limit.

2.2.3. Stability Classification Metrics

Because the direct application of MOST is often limited by operational data availability, researchers have adopted alternative metrics that utilise widely available bulk meteorological data to approximate stability conditions without necessitating direct, high-frequency flux observations.

The most prominent alternative metric in experimental offshore setups is the Bulk Richardson Number (Ri_b), which relates the buoyancy production of turbulence directly to the mechanical shear production. This dimensionless parameter approximates the theoretical gradient Richardson number using discrete, finite differences across available measurement heights, generally defined as:

$$Ri_b = \frac{g}{\bar{\theta}_v} \frac{\Delta \bar{\theta}_v z_m}{(\Delta \bar{U})^2} \ln \left(\frac{z_2}{z_1} \right) \quad (2.5)$$

where z_m is the geometric mean height, and $\Delta\bar{\theta}_v$ and $\Delta\bar{U}$ represent the differences in virtual potential temperature and mean horizontal wind speed between heights z_1 and z_2 . In simplified offshore applications where only sea surface temperature (T_s) and air temperature (T_a) are available, this expression can be reduced to relate the temperature difference directly to the wind speed at a reference height z .

This metric serves a dual purpose in contemporary literature: it acts either as a standalone classifier for stratification or as a mathematical bridge to estimate the Obukhov length via empirical correlations. Studies such as Pérez et al. (2023) and Bardal et al. (2018) use Ri_b to derive L by applying stability correction functions derived from Businger et al. (1971), Arya (1988), and Dyer (1974). These empirical relations map the bulk Richardson number to the dimensionless stability parameter $\zeta = z/L$, allowing for the theoretical reconstruction of L even in the absence of sonic anemometry. Equation 2.6 below illustrates the widely adopted empirical mapping formulated by Dyer (1974):

$$\frac{z}{L} = \begin{cases} Ri_b, & Ri_b \leq 0 \\ \frac{Ri_b}{1-5Ri_b}, & 0 < Ri_b < 0.2 \end{cases} \quad (2.6)$$

A third, simplified category of classification, bypasses these thermodynamic derivations entirely, categorising stability regimes based strictly on isolated kinematic proxies such as wind shear, wind veer, and Turbulence Intensity (TI). While computationally simple and reliant solely on standard anemometry or remote sensing, this method fundamentally overlooks the thermal drivers of the ABL. Nonetheless, it remains a common pragmatic approach in the industry when accurate temperature gradient data is completely absent.

2.3. Wind Turbine Aerodynamics and Power Conversion

To fully understand how atmospheric stability impacts offshore wind generation, it is first necessary to establish the theoretical mechanics of how a wind turbine converts kinetic energy into electrical power. This section bridges the gap between the meteorological boundary layer conditions discussed previously and the machine's aerodynamic performance. The following subsections outline the fundamental power equation, the environmental variables that influence energy extraction, and the standardised aerodynamic regions of the turbine's operational power curve.

2.3.1. Aerodynamic Power Extraction and Efficiency Constraints

The extraction of kinetic energy from the wind is governed by the fundamental aerodynamic power equation. Rather than extracting the entirety of the wind's energy, the total power captured by the turbine is a fraction of the available kinetic energy flowing through the rotor swept area, defined by the power coefficient. These foundational relationships are expressed as:

$$P = \frac{1}{2}\rho AV^3 C_p \quad \text{and} \quad C_p = \frac{P}{P_{\text{available}}} \quad (2.7)$$

where:

P	is the actual aerodynamic power extracted by the turbine;	
ρ	is the air density;	
A	is the rotor swept area;	
V	is the free-stream wind speed;	The effi-
C_p	is the power coefficient;	
$P_{\text{available}}$	is the total kinetic power available in the free-stream wind ($\frac{1}{2}\rho AV^3$).	

ciency of this mechanical energy extraction is bounded by strict physical limits, most notably the Betz limit. Formulated by physicist Albert Betz (1920), this theoretical law dictates that an ideal wind turbine can extract a maximum of 59.3% of the kinetic energy from the wind. The power coefficient (C_p), which represents the actual aerodynamic efficiency of a specific turbine design, will therefore always remain well below this 59.3% theoretical ceiling, owing to inherent mechanical inefficiencies and aerodynamic phenomena such as wake rotation and blade tip losses.

The cubic relationship between wind speed and power (V^3) means that spatial variations across the rotor have a strong, non-linear impact on total energy yield. As modern offshore turbines feature large rotor diameters, the upper portion of the swept area frequently experiences markedly higher wind velocities than the lower portion due to vertical wind shear, as observed earlier in Figure 2.1. Because the wind speed variable is cubed in the power equation, the additional energy captured at the top of the highly sheared rotor does not simply cancel out the

energy lost at the bottom. This non-linear effect is the primary reason why turbines can experience notable over- or under-performance relative to standard neutral baseline curves when operating in complex stability regimes, and this will be investigated further in the following chapters.

2.3.2. Environmental Dependencies

Beyond the baseline wind velocity, the turbine's power output is directly modulated by the thermodynamic properties of the incoming air mass. Air density (ρ) acts as a linearly proportional variable in the power equation, meaning a denser air mass carries more extractable kinetic energy. This physical density is governed by the Ideal Gas Law ($P = \rho RT$), which demonstrates that air density fluctuates continuously with dynamic changes in ambient temperature and atmospheric pressure.

In offshore environments, atmospheric humidity introduces a critical secondary effect on air density through the concept of virtual temperature. Because water vapor molecules are physically lighter than the dry air molecules they displace (predominantly nitrogen and oxygen), a humid air mass is less dense than a dry air mass occupying the same volume at an identical temperature and pressure. Accurately accounting for this moisture-driven density reduction is particularly vital for offshore wind resource assessments, where the boundary layer is continuously saturated by the underlying ocean surface.

Furthermore, the kinematic properties of the wind field, specifically wind shear and turbulence intensity (TI), interact dynamically with the rotor to alter the baseline aerodynamic performance. While shear dictates the vertical distribution of the incoming energy flux, high-frequency turbulent fluctuations (TI) cause the wind speed to constantly oscillate around its mean value. These environmental dependencies demonstrate that a turbine's performance cannot be evaluated based on wind speed alone; the thermodynamic and kinematic state of the atmosphere fundamentally dictates the true energy potential.

2.3.3. The Power Curve

The operational response of a wind turbine to varying environmental conditions is standardised and visualised through the theoretical power curve, as illustrated in Figure 2.5. This performance curve defines the expected electrical output across the full range of operational wind speeds, dividing the turbine's behaviour into distinct aerodynamic and mechanical regions.




Figure 2.5: Wind turbine power curve and operating regions.

The partial load region, often referred to as Region II, represents the operational phase where the turbine strives for maximum aerodynamic efficiency. Beginning at the cut-in wind speed, the minimum velocity required to generate usable power, the turbine's internal control system actively attempts to maintain an optimal tip-speed ratio. In this region, the cubic relationship of the power equation is in full effect, and the machine's energy yield is highly sensitive to the atmospheric stability deviations and shear profiles discussed previously.

The transition region, commonly referred to in the industry as the "knee" of the curve, connects the partial load phase to the turbine's maximum operational limits. As the wind speed approaches the machine's rated capacity, the aerodynamics become non-linear and complex. This region requires careful analytical attention, as the control system begins to shift its primary strategy from maximising energy capture to regulating structural and mechanical loads.

Finally, in the rated power region (Region III), the turbine actively limits energy capture to protect its mechanical and electrical components from excessive stress. Using smart pitch control, the turbine pitches its blades out of the wind to shed excess aerodynamic torque, maintaining a constant "rated" power output regardless of further increases in wind speed. Because the turbine's aerodynamics are being deliberately spoiled in this phase, the environmental effects of wind shear and thermodynamic stability are effectively neutralised. This state persists until the wind reaches the cut-out speed, at which point the turbine shuts down entirely for structural safety.

2.4. Remote Sensing and Measurement Technology

Since this research campaign relies exclusively on nacelle-mounted Lidar data for characterising the inflow conditions, ensuring the fidelity of these measurements is important. To properly interpret the resulting dataset and distinguish between physical atmospheric trends and potential measurement artefacts, it is necessary to first establish the fundamental working principles of the device and acknowledge the specific methodological challenges associated with its operation.

2.4.1. Working Principles of the WindCube Nacelle Lidar

The primary measurement instrument used in this campaign is the WindCube Nacelle by Vaisala, an active remote sensor operating on the principles of pulsed Doppler Light Detection and Ranging (Lidar). The system functions by transmitting laser pulses successively along four distinct lines of sight (LOS) into the atmosphere (Figure 2.6), where the light interacts with suspended particulates such as dust, water droplets, and aerosols. As outlined in foundational Lidar theory (Weitkamp 2005), a portion of this light is backscattered to the receiver, converted into an electrical signal, and processed to extract the frequency shift caused by the motion of the particles relative to the beam.




Figure 2.6: WindCube Nacelle Lidar system illustrating the configuration of the four distinct Lines of Sight (LOS0, LOS1, LOS2, and LOS3) used for wind field reconstruction, Vaisala (2023).

This frequency shift, known as the Doppler shift, allows for the direct calculation of the radial wind speed along each specific line of sight, while the time-of-flight of the pulse determines the measurement range. The system's internal processing algorithms reconstruct the three-dimensional wind vector from these radial components, outputting not only horizontal wind speed and direction but also derived atmospheric variables such as the shear exponent (α) and Turbulence Intensity (TI). However, as emphasised by Letizia et al. (2023), these derived metrics must be approached with scepticism, as the geometric assumptions inherent in the mathematical reconstruction can introduce uncertainties when compared to traditional point measurements.

2.4.2. Nacelle-Mounted Lidar Challenges

A prominent challenge in standard offshore measurement campaigns is the reliance on nacelle-mounted remote sensing without the independent validation typically provided by a meteorological mast. While the absence of

an independent, offshore met mast removes the "ground truth" reference for isolating instrumental errors, this limitation is partially mitigated by using the turbine's standard SCADA nacelle anemometry. Cross-referencing the Lidar data with the SCADA anemometer provides a secondary verification mechanism, useful in particular for confirming temporal synchronisation and basic directional alignment.

The physical nature of Lidar measurements also introduces the phenomenon of volume averaging, which distinguishes remote sensing from in-situ anemometry. Unlike a cup or sonic anemometer that measures wind at a discrete point, a pulsed Lidar averages the wind speed over the length of the probe volume along the beam path. This spatial averaging acts as a low-pass filter, smoothing out high-frequency turbulent fluctuations. If not corrected by specific post-processing algorithms, this volumetric filtering can lead to a systematic underestimation of the true Turbulence Intensity.

Perhaps the most important limitation of multi-beam Lidar systems is their reliance on the assumption of horizontal and vertical flow homogeneity. To reconstruct a single wind vector, the Lidar's algorithm assumes that the wind field is uniform across the entire geometric plane spanning the four divergent beams. As the measurement distance increases (e.g., beyond several hundred metres), however, the spatial separation between the individual lines of sight becomes large. In strongly turbulent atmospheres, such as those characterised by unstable stratification and localised convective gusts, the distinct beams may measure entirely different wind structures at the same time. This discrepancy violates the homogeneity assumption, introducing "beam decorrelation errors" that can artificially skew the reconstructed wind speed profile at long distances.

Finally, the accuracy of the reconstructed wind profile is highly sensitive to the physical alignment of the device relative to the rotor axis. Any static offsets in the tilt or pitch of the installation can propagate as structural biases in the calculated wind speed and shear profiles, potentially distorting the perceived stability regime. To mitigate the risks associated with such alignment uncertainties, researchers often rely on analytical reconstruction frameworks. For instance, methodologies proposed by Bahamonde and Litrán (2019) draw on surface boundary conditions, such as atmospheric stability (Ri_b) and sea surface roughness (z_0), to establish a theoretical baseline for evaluating wind profiles compromised by alignment offsets. While dynamic turbine motion (pitching and rolling) can also introduce transient measurement errors offshore, the high structural rigidity of the target turbine's foundation in this campaign renders such dynamic effects negligible.

2.5. State of the Art: Stability Impacts on Power Performance

This literature review closes with an examination of how the previously defined stability-driven flow features directly influence the electrical output of offshore wind turbines. While the thermodynamic definitions of the Atmospheric Boundary Layer (ABL) are well-established, definitively linking these atmospheric regimes to mechanical turbine performance remains a subject of ongoing research and debate. Because standard, neutrally-assumed power curves fail to account for stability-dependent variables, critical inaccuracies frequently emerge in yield assessments. This section draws together the current academic consensus, and the prevailing conflicts, regarding these aerodynamic impacts, thereby identifying the specific research gap this thesis addresses.

2.5.1. Impact on Energy Yield: The Stability Conflict

A nuanced conflict exists within current academic literature regarding the direction and magnitude of turbine performance deviations under stable stratification. Field observations diverge into two primary and opposing perspectives.

The first perspective argues for **over-performance in stable air**. Researchers such as St. Martin et al. (2016) and Berg (2008) report that stable conditions yield significantly higher power output relative to neutral baselines, particularly at wind speeds approaching rated power. This observation is rooted in the non-linear aerodynamics of the rotor. Because stable atmospheres induce high vertical wind shear, the upper portion of the large rotor swept area experiences markedly higher wind velocities. Due to the cubic relationship between wind speed and power (V^3), the exponential energy gain at the top of the rotor outweighs the energy lost at the bottom, resulting in a net surplus in harvested energy.

A second school of thought, by contrast, observes distinct **under-performance in stable air**. Studies by Dörenkämper et al. (2014) have shown that, even in non-wake scenarios, stable stratification can lead to energy deficits of up to 15% compared to unstable or neutral regimes. This opposing view points out that stable boundary layers are not merely defined by velocity shear; they also suppress vertical turbulent mixing and frequently introduce

pronounced wind veer. The lack of turbulence starves the rotor of kinetic energy entrainment, while high directional veer introduces severe aerodynamic misalignment across the blade span, effectively spoiling the angle of attack and negating the theoretical benefits of the high-shear profile.

Resolving this contradiction is not straightforward, as the influence of atmospheric stability on power generation is a multivariate problem. Site-specific factors, including geographical location, local wind climatology, shallow-water thermodynamics, and the distinct control logic of the turbines employed, interact dynamically to dictate which of these aerodynamic effects ultimately prevails. However, understanding this conflict requires a closer examination of the methodologies and wind speed definitions employed by these researchers.

Interestingly, both Dörenkämper et al. (2014) and St. Martin et al. (2016) derived their opposing conclusions by analysing measured operational field data rather than relying on numerical weather prediction models, like the Weather Research and Forecasting (WRF) WRF. An important distinction, however, lies in how they approached the reference wind speed. While St. Martin et al. (2016) initially evaluated a mathematically integrated wind speed across the rotor disk to account for vertical shear, the authors determined that for the specific measurement campaign, this integrated metric did not deviate enough from the standard hub-height velocity to warrant its implementation. Both studies therefore largely base their observed power deviations on a single, hub-height wind speed measurement.

This shared reliance on a single-point measurement to explain highly stratified, stability-driven flow features exposes a fundamental limitation in standard performance evaluations, pointing clearly towards the industry's proposed analytical solution.

2.5.2. Rotor Equivalent Wind Speed (REWS)

The dependency of power performance on these shear profiles shows the inadequacy of using a single hub-height wind speed measurement for modern, large-rotor offshore turbines. Relying on a single point to represent a vertically distributed and highly stratified energy flux introduces substantial scatter into the empirical power curve, masking the true aerodynamic efficiency of the machine across varying stability regimes.

To address this theoretical limitation, the industry standard introduces the concept of Rotor Equivalent Wind Speed (REWS), which mathematically integrates the wind speed at multiple discrete heights across the swept area. By calculating a weighted velocity average, REWS provides a high-fidelity representation of the actual kinetic energy available to the turbine.

While REWS represents the theoretical gold standard for shear normalisation, however, its practical implementation in offshore environments presents serious instrumental challenges. The strict IEC guidelines require simultaneous wind speed measurements at a minimum of three distinct heights across the vertical rotor plane to accurately calculate the equivalent velocity. While this is achievable with heavily instrumented meteorological masts or advanced 3D scanning Lidars, standard forward-looking nacelle Lidars frequently lack the vertical geometric spread required to capture these heights simultaneously, particularly at closer distance gates. As a result, researchers must often work through the difficult trade-off between theoretical aerodynamic ideals and the physical constraints of remote sensing technology when designing stability-driven correction methodologies.

3

Methodology & Computational Framework

This chapter details the data processing and analytical framework developed to investigate the aerodynamic impact of atmospheric stability on offshore wind turbine performance. Bridging raw meteorological observations with power assessment techniques, the methodology establishes a computational pipeline that transforms high-frequency offshore data into quantifiable energy metrics. As illustrated in the flowchart (Figure 3.1), this pipeline is structured into five core phases. The following sections outline: the physical site characterisation and measurement campaign; the temporal integration of independent data streams; the signal processing and quality control protocols; the formulation of kinematic and thermodynamic stability classifications; and finally, the standardised power performance assessment and energy yield extrapolation.




Figure 3.1: Scheme of the proposed methodology, detailing the data pipeline from raw ingestion to final power deviation quantification.

3.1. Site Characterisation and Measurement Campaign

Accurate assessment of wind turbine performance requires a fundamental understanding of the local offshore environment and the specific instrumentation used. This section details the physical layout of the target wind farm and its surrounding clusters, followed by an overview of the nacelle-mounted Lidar and supporting meteorological data streams employed to capture the inflow conditions.

3.1.1. Offshore Site and Layout

The measurement campaign was conducted at a commercial offshore wind farm situated in the North Sea. The facility is located tens of kilometres offshore and comprises multiple multi-megawatt wind turbines. The farm of interest also sits within a larger operational offshore wind cluster comprising four distinct wind farms. The spatial arrangement of these neighbouring wind farms within the cluster relative to the primary measurement turbine, hereafter referred to as the target turbine, introduces aerodynamic interactions that must be accounted for during the data filtering process.

Due to commercial non-disclosure agreements regarding asset performance, the exact coordinates and specific turbine models are omitted from this methodology. However, the localised wind climatology, Lidar orientations, and relative array spacing remain preserved to ensure the scientific validity of the aerodynamic wake exclusion calculations.




Figure 3.2: Offshore wind farm cluster characteristics.

The local wind climatology at the target turbine exhibits distinct directional patterns. As illustrated in the local wind rose (Figure 3.2b), the most frequent incoming wind directions recorded during the campaign originate primarily from the south, south-west, and south-east sectors. While this distribution generally reflects the dominant regional meteorological trends of the North Sea, the inflow measured from certain sectors may also be subjected to macroscopic blockage effects or wake interactions generated by the adjacent upwind clusters. Directional filtering is therefore required so that the analysed data represents only the undisturbed, free-stream atmospheric boundary layer prior to aerodynamic performance assessment.

3.1.2. Instrumentation and Data Streams

Acquiring high-fidelity environmental and operational data is central to this research. This section details the primary data streams used in the measurement campaign: remote sensing from a nacelle-mounted Lidar, thermodynamic and meteorological measurements from dedicated environmental sensors, and operational turbine data from the SCADA system.

Nacelle-Mounted Lidar

The primary instrument for characterising the inflow wind field is a nacelle-mounted Lidar system (Figure 3.3a). Specifically, the campaign employs the WindCube Nacelle manufactured by Vaisala (2020). This device is a forward-looking, pulsed Doppler Lidar. As previously detailed in Sections 2.4.1 and 2.4.2, understanding its specific working principles and inherent measurement challenges is important for proper data interpretation.


To capture a complete view of the incoming wind, the Lidar is configured with four distinct Line of Sight (LOS) measurement beams. This multi-beam geometry enables the vectorial reconstruction of variables across the rotor swept area, including horizontal wind speed, wind shear, wind veer, and Turbulence Intensity (TI). This long-range version of the instrument also offers modifiable distance gates. This flexibility allows for measurements ranging from 50 m up to 700 m upstream of the turbine. To align with industry standards, the resulting spatial dataset provides these readings as 10-minute interval averages for each configured distance gate.

Temperature Sensors

To evaluate the thermodynamic stability of the atmospheric boundary layer, dedicated high-precision temperature datasets are integrated into the analysis. The ambient air temperature was recorded using a Vaisala HMP155 sensor, an industry-standard meteorological probe known for its reliability in harsh marine environments. Additionally, the Sea Surface Temperature (SST) was captured using a CT15 infrared radiation pyrometer. This pyrometer is actively compensated for downward sky radiation and cloud coverage reflections, ensuring an accurate reading of the exact sea surface skin temperature. Both of these sensors were securely installed on the turbine's transition piece, as depicted in Figure 3.3b.



(a) The WindCube Nacelle Lidar during installation.



(b) Meteorological sensors (HMP155, CT15, and PTB330) mounted on the transition piece.

Figure 3.3: Primary environmental monitoring instrumentation deployed on the target turbine during the measurement campaign.

These vertical temperature gradients are crucial to calculate the thermodynamic stability metrics, in particular the Bulk Richardson Number (Ri_b). Because the raw SST dataset is initially recorded at a high frequency of 1 Hz, a temporal aggregation procedure is required. These high-frequency readings are then time-averaged to match the standard 10-minute intervals of the Lidar data, ensuring temporal consistency for all downstream data processing.

Pressure and Humidity Sensors

In addition to temperature gradients, accurately quantifying the virtual potential temperature of the atmospheric boundary layer requires precise barometric and moisture data. Atmospheric pressure was continuously measured using a Vaisala PTB330 digital barometer, which provides high accuracy and long-term stability across varying meteorological conditions. Relative humidity was simultaneously captured by the integrated Vaisala HMP155 probe. Similar to the temperature instruments, these environmental sensors were mounted directly on the transition piece of the target turbine. The raw data streams from these instruments were also time-averaged into 10-minute intervals to integrate with the Lidar and SCADA datasets.

SCADA System

Finally, operational data from the wind turbine itself is collected via the onboard Supervisory Control and Data Acquisition (SCADA) system. This system uses the standard instrumentation pre-installed on the turbine to provide the mechanical and electrical performance metrics, most notably the actual power output and the nacelle yaw position. In addition to performance tracking, the SCADA system continuously records the standard wind speed and wind direction measured directly at the nacelle. These secondary wind measurements serve a verification

role; they are cross-referenced with the primary Lidar dataset to validate the remote sensing observations, correlate wind direction shifts, and to confirm that all data streams are accurately time-aligned before the performance assessment begins.

3.2. Data Integration and Synchronisation

Before isolating the aerodynamic impacts of atmospheric stability, the disparate meteorological and operational data streams must be unified into a cohesive, analysis-ready dataset. To handle the large volume of remote sensing and SCADA measurements, a systematic data processing pipeline was established. As previously illustrated in Figure 3.1, raw data from the Lidar, SCADA, and meteorological sensors are systematically ingested, merged, filtered, and categorised by atmospheric stability. This step-by-step processing isolates measurement anomalies early in the computational pipeline, supporting reproducibility and maintaining analytical flexibility when making any adjustments.

3.2.1. Temporal Alignment

The first phase of this integration is verifying the temporal synchronisation between the independent measurement instruments. This alignment acts as a primary quality control measure to identify any potential drift or malfunctions in the internal clocking mechanisms of the sensors. If the raw data streams are temporally misaligned, any subsequent correlations between localised atmospheric stability and transient turbine performance will be flawed.

To mathematically verify this synchronisation, a cross-correlation analysis was conducted between the strongly coupled kinematic variables of the independent sensors. Specifically, the horizontal wind speed measured at hub height by the Lidar was cross-correlated against the standard nacelle anemometry data recorded by the SCADA system. The analysis applied discrete time shifts at 10-minute intervals, matching the native temporal resolution of the measurement campaign, across various Lidar distance gates to identify the point of maximum correlation.

The results of this analysis confirm that the internal clocks of the measurement devices are correctly synchronised. As depicted in Figure 3.4, the maximum correlation coefficient across all evaluated distance gates consistently occurs at a 0-minute time lag. This distinct zero-lag peak provides high confidence that the datasets are time-aligned and ready for multi-variable integration.




Figure 3.4: Cross-correlation plot of Lidar vs. SCADA wind speeds across selected distance gates, indicating peak correlation at zero lag.

The height of this correlation is not uniform across the measurement volume: it is lowest at the nearest gate and varies mildly with distance (note that for visual legibility, only a representative subset of distance gates is plotted in Figure 3.4). This does not affect the temporal synchronisation, which depends on the lag of the peak rather than its magnitude. The spatial pattern is of interest in its own right and is examined in Section 4.4.2 (Figure 4.18).

3.3. Data Filtering and Quality Control

The integrity of the wind turbine performance assessment relies heavily on the quality and physical validity of the ingested datasets. Therefore, raw measurements must undergo careful screening to eliminate environmental distortions, instrumental errors, and non-representative operational states. The following subsections detail the exclusion of aerodynamic wakes, the quality assurance of the remote sensing and meteorological instruments, and the filtering of turbine underperformance and curtailment events.

3.3.1. Environmental and Directional Filtering

To accurately assess baseline turbine performance, it is necessary to isolate free-stream atmospheric flow from conditions skewed by the aerodynamic wakes of neighbouring turbines. When a Lidar measures a wind field partially obstructed by an upstream rotor, the resulting spatial heterogeneity introduces severe artefacts into the derived wind speed and directional data. Therefore, the dataset must be restricted to a specific range of incoming wind directions where the target turbine experiences undisturbed, homogeneous flow.

The exclusion of these wake-disturbed sectors is calculated mathematically using the geometric parameters of the wind farm, in accordance with the IEC 61400-50-3 standards for nacelle-mounted Lidars IEC (2022). The first variable in this standard is the adjusted measurement distance, R_b , which represents the Lidar's range corrected for the geometric orientation of the beams. This is defined as:

$$R_b = \frac{R_{\text{conf}}}{\cos(\beta_{\text{max}}/2)} \quad (3.1)$$

where R_{conf} represents the selected Lidar measurement range (distance gate), and β_{max} is the largest horizontal angle between any two beam positions. For the Vaisala WindCube Nacelle used in this campaign, β_{max} is a fixed geometric constraint of 30° . Because R_{conf} is an adjustable parameter, the resulting wake exclusion zones are dependent on the chosen distance gate.

With the beam radius established, the angular width of the wake region generated by an upstream turbine can be calculated. Under the standard, the excluded wind sector, θ_{wake} , is given by the following empirical relationship:

$$\theta_{\text{wake}} = 1.3 \tan^{-1} \left(2.5 \frac{D_n}{L_n - R_b} + 0.15 \right) + 10^\circ \quad (3.2)$$

where D_n is the rotor diameter of the neighbouring, wake-producing turbine, and L_n is the absolute distance from the Lidar to the centre of that neighbouring tower.

The downstream impact of these wakes is evaluated up to a specific relative distance threshold to ensure all relevant aerodynamic disturbances are captured. Following industry best practices, the wake exclusion calculation is applied to all surrounding turbines that fall within a relative distance of 20 rotor diameters, since wake recovery beyond that threshold in offshore marine environments is generally considered complete. This dimensionless parameter is determined by the equation:

$$\text{Relative distance} = \frac{L_n - R_b}{D_n} \quad (3.3)$$

Applying these standard calculations to the layout of the offshore site reveals that undisturbed inflow is limited to the westerly side of the target turbine, as illustrated in Figure 3.5. Due to the dense clustering of neighbouring turbines to the east, north, and south, the mathematically derived free-flow sector is restricted to the range of 218° to 304.9° (where 360° represents True North). It must be noted that these boundaries are sensitive to the chosen R_{conf} value, and any shift in the prioritised distance gate during the analysis requires a corresponding recalculation of this clear-sector window.




Figure 3.5: Sectors excluded due to wakes of neighbouring and operating WTGs.

Complementary Validation Based on Relative Wind Direction

A complementary validation of this free-flow sector is achieved by analysing the relative wind direction as a function of the nacelle yaw position. Figure 3.6 visualises the impact of neighbouring turbines on the Lidar's directional readings, with the dashed red lines indicating the physical azimuths of nearby wind turbines. When the Lidar operates near these boundaries, partial wake interactions occur, for instance, when one or two of the four Lidar beams intersect a wake while the others remain in free air. This inhomogeneous flow causes the Lidar's reconstruction algorithm to calculate a large, incorrect relative wind direction, even if the nacelle is perfectly aligned with the true ambient wind. However, within the calculated 218° to 304.9° sector, these spurious directional spikes are absent, visually confirming that the selected window is effectively undisturbed.




Figure 3.6: Scatter plot showing relative wind direction vs. yaw position with neighboring turbines marked by dashed red lines.

Based on these aerodynamic distortions, the filtering methodology prioritises the nacelle yaw position over the raw wind direction signal to isolate the free-flow data. Filtering by the yaw parameter provides a consistent mechanism to ensure the retained dataset represents true, undisturbed inflow. This filtering and validation technique aligns with methods established in recent remote sensing literature Rozenn Wagner and Davoust (2013).

A cross-examination between the SCADA-recorded wind direction and the Lidar-derived wind direction is provided in Appendix A. This supplementary analysis visually and statistically demonstrates the efficacy of the filtering process in removing spatial discrepancies and aligning the two measurement systems.

Complementary Validation Based on Turbulence Intensity Differences

Further validation is conducted by examining the variance of Turbulence Intensity (TI) across the four distinct lines of sight. Because Lidar beams measuring within a homogeneous wind field theoretically capture similar turbulence structures, the TI values along the individual lines of sight should be comparable. By contrast, in an inhomogeneous wind field, such as a partial wake intersection, the resulting TI measurements diverge significantly. As demonstrated in Figure 3.7, the variance in TI is minimised for yaw values between 218° and 304.9° . The prominent peak observed at approximately 220° in the raw dataset is successfully filtered out in Figure 3.7b. This homogeneity in turbulence across the beams provides a final confirmation that the selected sector effectively isolates the true free-stream flow.




Figure 3.7: Radial speed turbulence intensity along the four line-of-sight vs. turbine yaw position for (a) the raw dataset and (b) the filtered dataset between 218° and 304.9° following artefact removal and quality control. Yaw position is averaged into 10° bins.

3.3.2. Lidar and Meteorological Data Quality Control

Before deriving physical atmospheric metrics, the raw signals acquired from the remote sensing and temperature instruments must be carefully filtered. This instrument-level quality control isolates the high-fidelity data required for stability classification by removing instrumental noise and mathematically implausible readings.

The first quality control measure applied to the Lidar dataset is the beam availability filter. To ensure a robust vectorial reconstruction of the wind field, this filter requires a minimum 50% accepted data value for each of the four individual measurement beams, aligning with conservative operational guidelines Vaisala (2023). However, recent academic literature Rozenn Wagner and Davoust (2013) suggests that Lidar accuracy is effectively not compromised by lower availability rates, provided the availability does not drop below critical measurement thresholds (e.g., 25%).

Following the availability check, a physical plausibility filter is implemented to act as a sanity check against erroneous turbulence readings. Because remote sensing artefacts can occasionally produce severe, non-physical

data spikes, this filter removes physically suspicious TI values. Specifically, TI measurements are retained if they fall below the mathematically defined boundary of:

$$TI < 5 \cdot RWS^{-1.25} \quad (3.4)$$

where RWS represents the measured Radial Wind Speed. This ensures that the reported turbulence remains aerodynamically logical for the corresponding velocity.

The final remote sensing filter isolates high-confidence measurements based on the Carrier-to-Noise Ratio (CNR). Acting as the primary indicator of signal clarity, the acceptable data range is bounded between -26 dB and 0 dB. These thresholds mirror those imposed by Schneemann et al. (2021), who used a similar nacelle-mounted Lidar configuration to assess offshore wind fields.

3.3.3. Filtering Turbine Underperformance and Curtailment

To accurately quantify power deviations driven exclusively by atmospheric stability, it is necessary to decouple environmental effects from mechanical or grid-driven power reductions. If the wind turbine is artificially derated or experiencing downtime, the resulting power curve will not reflect the true aerodynamic potential of the incoming wind field. Consequently, two operational filters are applied to the SCADA dataset.

The first operational filter removes grid-driven curtailment events by cross-referencing SCADA status codes and active power setpoints. During any analysed 10-minute interval, the wind turbine must be attempting to maximise its energy yield without being subjected to active power capping or externally commanded derating from the grid operator. Because the core objective of this research is to evaluate the true aerodynamic capability of the turbine under varying atmospheric conditions, it is important that the machine is operating optimally. Any external power restrictions would artificially mask the natural aerodynamic response to the wind field that this study aims to measure.

A second, temporal threshold is also enforced. The SCADA logs are filtered to require that the turbine is in full operational mode for the entirety of the 600 seconds comprising each 10-minute data point. This requirement ensures that any observed variance in power output is a true physical representation of the atmospheric boundary layer at that moment, rather than an artificial consequence of transient mechanical downtime, braking interventions, or startup sequences.

3.4. Atmospheric Stability Classification

Following the filtering of unwanted aerodynamic and instrumental artefacts, the finalised dataset is processed to categorise the prevailing atmospheric stability regimes. This categorisation is a challenging procedure within offshore wind energy research, as the true thermodynamic state of the boundary layer is difficult to measure directly using standard industry equipment.

Accurately quantifying these stratification regimes relies on Monin-Obukhov Similarity Theory (MOST). However, the direct application of MOST requires extensive, specialised equipment, such as high-frequency sonic anemometers, capable of capturing three-dimensional turbulent velocity fluctuations and kinematic surface heat fluxes. Because this level of high-fidelity data is rarely available in standard commercial offshore campaigns, this framework employs a multi-tiered classification strategy. Initially, the methodology uses widely available kinematic variables, specifically wind shear and Turbulence Intensity (TI), as practical classification proxies. It then incorporates thermodynamic metrics, namely the Bulk Richardson Number (Ri_b) and the derived Monin-Obukhov length (L), to establish a physics-based categorisation.

3.4.1. Wind Shear and Turbulence Intensity Proxies

The initial phase of the stability classification relies on the empirical behaviour of the vertical wind profile and the associated mechanical turbulence. While isolated variables like the wind shear exponent (α) and TI bypass the underlying thermodynamic derivations of the boundary layer, they are widely used as practical proxies in academic literature for primary stability categorisation.

To implement this proxy-based classification, discrete numerical limits must be established to separate the dataset into unstable, neutral, and stable regimes. The specific threshold values applied in this framework, detailed in Table 3.1, are derived from established offshore boundary layer literature Dörenkämper et al. (2014) and St. Martin et al. (2016).

Table 3.1: Classification thresholds for atmospheric stability based on kinematic proxies.

Stability Class	Wind Shear (α)	Turbulence Intensity (TI)
Unstable conditions	$0 \leq \alpha < 0.11$	$TI > 0.045$
Neutral conditions	$0.11 \leq \alpha \leq 0.17$	$0.03 \leq TI \leq 0.045$
Stable conditions	$\alpha < 0$ or $\alpha > 0.17$	$TI < 0.03$

In this framework, profiles exhibiting negative shear ($\alpha < 0$) are assigned to the stable class. A negative exponent indicates that the measurement heights lie at or above the nose of a low-level jet, a localised wind-speed maximum within the lower boundary layer. Offshore, such jets develop predominantly under stable stratification, where the suppressed vertical mixing and layer decoupling discussed in Section 2.2 permit their formation; their negative-shear signature is therefore consistent with the stable regime.

3.4.2. Thermodynamic Metrics

To achieve a more complete physical characterisation of the atmospheric boundary layer, this methodology incorporates thermodynamic stability metrics. While the kinematic proxies discussed previously provide a useful preliminary baseline, they do not directly account for the thermal stratification that drives or suppresses atmospheric turbulence. This phase of the analysis therefore draws on direct temperature, pressure, and humidity measurements to construct a physics-driven classification.

Bulk Richardson Number

The Bulk Richardson Number (Ri_b) serves as the primary site-specific thermodynamic metric, evaluating atmospheric stability by directly comparing buoyancy forces against mechanical wind shear. While some practical offshore applications use an approximation based on absolute temperature gradients alone St. Martin et al. (2016), achieving true thermodynamic accuracy requires accounting for atmospheric moisture and pressure variations. This study therefore adopts the theoretical formulation using virtual potential temperature (θ_v), as detailed by Stull (1988) and applied by Zhang, Walker, and Markfort (2025).

For the 10-minute averaged periods measured during the campaign, the Bulk Richardson Number is calculated as:

$$Ri_b = \frac{g \Delta \bar{\theta}_v \Delta z}{\bar{\theta}_v [(\Delta \bar{U})^2 + (\Delta \bar{V})^2]} \quad (3.5)$$

where:

- g is the gravitational acceleration (9.81 m/s²),
- $\bar{\theta}_v$ is the mean virtual potential temperature of the evaluated atmospheric layer,
- $\Delta \bar{\theta}_v$ is the difference in virtual potential temperature across the vertical height interval Δz ,
- $\Delta \bar{U}$ is the corresponding difference in the longitudinal component of the mean horizontal wind speed,
- $\Delta \bar{V}$ is the corresponding difference in the lateral component of the mean horizontal wind speed.

While the full formulation of Ri_b includes both the streamwise (U) and transverse (V) horizontal velocity components, remote sensing Lidar instruments operate under the assumption of flow homogeneity across their measurement volume. Transverse wind speed differences are therefore typically negligible in the mean directional profile, meaning the mechanical shear acting on the turbine is effectively driven by the primary horizontal wind speed gradient ($\Delta \bar{U}$).

A spatial adjustment was also made to the temperature datasets. Because the ambient air temperature sensor was located on the turbine's transition piece, the direct temperature reading did not represent the true hub-height boundary conditions required for the full rotor-plane gradient. To resolve this, the ambient temperature was extrapolated to hub height by applying a vertical temperature lapse rate. This lapse rate was calculated directly from the difference in ambient temperature between the transition piece height and hub height using ERA5 regional reanalysis data, giving a representation of the thermal profile across Δz .

To compute the virtual potential temperature (θ_v) required for Equation 3.5, readings of this extrapolated absolute temperature T , barometric pressure P , and relative humidity RH are used. First, the saturation vapour pressure

P_{sat} is approximated using Tetens' Equation, Murray (1967):

$$P_{sat}(T) \approx 0.6113 \exp \left(17.2694 \frac{T - 273.15}{T - 35.86} \right) \quad (3.6)$$

The actual vapour pressure P_{vap} is then extracted using the measured relative humidity:

$$P_{vap} = RH \cdot P_{sat}(T) \quad (3.7)$$

Then, the standard potential temperature θ is calculated, which standardizes the measured temperature to a reference pressure of 100 kPa:

$$\theta = T \left(\frac{100}{P} \right)^{R/C_p} \quad (3.8)$$

where $R/C_p \approx 0.286$ represents the gas constant of air divided by its specific heat capacity at constant pressure.

To account for the effect of moisture on air density, the water-air mixing ratio w is determined from the vapour pressure:

$$w = 0.622 \left(\frac{P_{vap}}{P - P_{vap}} \right) \quad (3.9)$$

Lastly, the mixing ratio and standard potential temperature are combined to derive the virtual potential temperature (θ_v):

$$\theta_v = \theta \left(\frac{1 + w/0.622}{1 + w} \right) \quad (3.10)$$

Finally, regarding the vertical height interval (Δz), the potential influence of tidal variations on the absolute elevation of the transition piece and corresponding measurement heights was evaluated. While an open-source tidal simulation was considered to track dynamic sea-level fluctuations, the resulting variance was found to be aerodynamically negligible for this classification framework. Because the reference height variable (Δz) is located in the numerator of the Bulk Richardson equation, small tidal fluctuations exert only a minor, linear scaling effect on the final Ri_b value, which is insufficient to alter the broader discrete stability classifications.

Monin-Obukhov Length

To complement the discrete parameter analysis, a site-specific estimation of the Monin-Obukhov length (L) is derived directly from the previously calculated Bulk Richardson Number. Because direct sonic anemometry was unavailable, this estimation uses the established empirical relationship formulated by Dyer (1974), as defined in Equation 3.11:

$$\frac{z}{L} = \begin{cases} Ri_b, & Ri_b \leq 0 \\ \frac{Ri_b}{1 - 5Ri_b}, & 0 < Ri_b < 0.2 \end{cases} \quad (3.11)$$

where z represents the designated measurement reference height and L is the resulting local Obukhov length.

To categorise the wind farm data based on these thermodynamic principles, numerical boundaries must be applied. Table 3.2 outlines the limits adopted for both the calculated Bulk Richardson Number and the corresponding Monin-Obukhov Length. It is important to note that because the true stability parameter is dimensionless (z/L), the absolute thresholds for L have been scaled according to the turbine's specific hub height to maintain thermodynamic equivalence with the Ri_b limits.

Table 3.2: Classification thresholds for atmospheric stability based on thermodynamic metrics.

Stability Class	Bulk Richardson Number (Ri_b)	Monin-Obukhov Length (L)
Unstable conditions	$Ri_b < -0.02$	$-5000 < L < 0$
Neutral conditions	$-0.02 \leq Ri_b \leq 0.02$	$ L \geq 5000$
Stable conditions	$Ri_b > 0.02$	$0 < L < 5000$

3.5. Power Performance Assessment

The final stage of the analytical pipeline translates the filtered meteorological and operational data into quantifiable turbine performance metrics. Because atmospheric stability alters the energy flux moving through the rotor swept area, it is important to accurately construct power curves that isolate these specific environmental effects. The following subsections detail the environmental data corrections, the empirical construction of the power curves via the Method of Bins, the estimation of Annual Energy Production (AEP), and the mathematical framework used to quantify true performance deviations across the identified stability regimes.

3.5.1. Environmental Corrections and Normalisation

Before an accurate power curve can be plotted, the raw meteorological and operational data must be standardised to isolate the effects of atmospheric stability. Because environmental parameters such as wind shear, air density, and turbulence constantly fluctuate during a measurement campaign, they introduce variations in power output that can obscure the true aerodynamic performance. To capture the energy flux accurately, these corrections are applied in a fixed sequential order: spatial shear correction, thermodynamic density normalisation, and finally turbulence normalisation.

Wind Shear Correction (REWS)

The first correction in the sequence accounts for the vertical variation of wind speed across the turbine's large rotor swept area. Standard power performance assessments historically relied on a single hub-height wind speed measurement; however, this simplified approach is inadequate for the large rotor swept areas of modern multi-megawatt offshore turbines, which operate across highly variable shear profiles. To resolve this, the Rotor Equivalent Wind Speed (REWS) is calculated to integrate the velocity profile across the entire rotor disk. As defined by the IEC (2017), the swept area is divided into discrete horizontal segments, as shown in Figure 3.8, and the equivalent wind speed is derived using the following equation:

$$V_{eq} = \left(\sum_{i=1}^n V_i^3 \frac{A_i}{A} \right)^{1/3} \quad (3.12)$$

where:

- V_{eq} is the rotor equivalent wind speed;
- V_i is the measured wind speed at the height of segment i ;
- A_i is the area of the i -th horizontal segment of the rotor swept area;
- A is the total rotor swept area;
- n is the total number of segments.




Figure 3.8: Diagram showing the rotor swept area divided into horizontal segments (A_i) to illustrate the REWS calculation.

Due to the operational constraints of the available dataset, a continuous vertical wind profile was not available. Specifically, the Lidar provided only two direct measurement heights, with the hub height profile being reconstructed from these two data points. While a discretized approach using these three heights was implemented and tested, it was ultimately excluded from the final results due to the limited number of measurement points. For a detailed breakdown of this tested methodology, the geometric derivation of the Lidar measurement heights, and an analysis of the wind shear correction's impact, see Annex B.

Air Density Normalisation

Following the spatial shear correction, the newly derived equivalent wind speed must undergo a temporal normalisation to account for fluctuating air density. This step must occur after the REWS calculation, as the physical air mass acts upon the total representative velocity field, not just the hub-height velocity. So that observed power deviations are not an artefact of varying air mass, the local 10-minute air density ($\rho_{10 \text{ min}}$) is determined from the measured absolute air temperature, air pressure, and relative humidity:

$$\rho_{10 \text{ min}} = \frac{1}{T_{10 \text{ min}}} \left(\frac{B_{10 \text{ min}}}{R_0} - \Phi P_w \left(\frac{1}{R_0} - \frac{1}{R_w} \right) \right) \quad (3.13)$$

where:

$\rho_{10 \text{ min}}$	is the derived 10 min averaged air density;
$T_{10 \text{ min}}$	is the measured absolute air temperature averaged over 10 min [K];
$B_{10 \text{ min}}$	is the air pressure corrected to hub height averaged over 10 min [Pa];
R_0	is the gas constant of dry air;
Φ	is the relative humidity;
R_w	is the gas constant of water vapour;
P_w	is the vapour pressure [Pa].

Once the localised air mass is established, the wind speed is normalised to a consistent reference air density (ρ_0). Because the evaluated offshore turbine features modern active pitch and power control systems, this normalisation must be applied directly to the wind speed measurements rather than the power outputs:

$$V_n = V_{eq} \left(\frac{\rho_{10 \text{ min}}}{\rho_0} \right)^{1/3} \quad (3.14)$$

where V_n is the density-normalised wind speed. Adjusting the wind speed to an "equivalent" velocity at standard density represents the shifting kinetic energy flux without violating the turbine's internal pitch and speed control logic.

Turbulence Intensity Normalisation

With the wind speed standardised for both spatial shear and thermodynamic density, the final environmental correction isolates the impact of turbulence by normalising the power output to a standard reference Turbulence Intensity (TI). This must be the final step in the sequence because the TI normalisation algorithm relies on simulating the turbine's theoretical power curve; simulating this curve using uncorrected wind speeds would yield flawed baselines. Utilising the procedures outlined in IEC 61400-12-1 Annex M, this step removes the artificial power enhancements or deficits caused purely by high-frequency wind fluctuations. The correction adjusts the measured 10-minute active power ($P_{10 \text{ min}}$) to reflect a consistent baseline TI (e.g., 6.5%) using the following relationship:

$$P_{I, \text{norm}} = P_{10 \text{ min}} - P_{\text{sim}}(V_n, I_{\text{measured}}) + P_{\text{sim}}(V_n, I_{\text{ref}}) \quad (3.15)$$

where:

$P_{I, \text{norm}}$	is the TI-normalised power output;
$P_{10 \text{ min}}$	is the raw measured active power;
$P_{\text{sim}}(V_n, I_{\text{measured}})$	is the simulated power calculated using the density-normalised wind speed and the actual measured TI;
$P_{\text{sim}}(V_n, I_{\text{ref}})$	is the simulated power calculated using the density-normalised wind speed and the chosen reference TI.

For the detailed analysis of the iterative mathematical fitting, parameter convergence criteria, and the step-by-step mathematical implementation of this procedure, refer to Annex C.

3.5.2. Power Curve Construction

With the dataset standardised for shear, density, and turbulence, the empirical power curves are constructed using the industry-standard Method of Bins. This methodology, mandated by the IEC 61400-12-1 guidelines, reduces the inherent scatter of high-frequency meteorological data by grouping time-series observations into discrete wind speed intervals. The density-normalised wind speeds and their corresponding corrected active power outputs are sorted into contiguous bins with a fixed width of 0.5 m/s. This binning process is applied across the entire operational envelope of the turbine, from the cut-in wind speed up to the cut-out limit, giving a full representation of the machine's aerodynamic performance.

Once the data points are sorted into their respective 0.5 m/s intervals, the representative wind speed for each bin is determined by calculating the arithmetic mean of all normalised wind speed readings falling within that interval. According to the IEC guidelines, the averaged wind speed for bin i is defined by Equation 3.16:

$$V_i = \frac{1}{N_i} \sum_{j=1}^{N_i} V_{n,i,j} \quad (3.16)$$

where:

- V_i is the normalised and averaged wind speed in bin i ;
- $V_{n,i,j}$ is the normalised wind speed of the j -th 10-minute dataset contained within bin i ;
- N_i represents the total number of 10-minute average datasets collected in bin i .

A parallel averaging procedure is applied at the same time to the power output data within the same bins. The representative power output for each interval is calculated by averaging all environmentally corrected power readings contained within it. Using the TI-normalised active power derived in the previous section, the averaged power output for bin i is calculated using Equation 3.17:

$$P_i = \frac{1}{N_i} \sum_{j=1}^{N_i} P_{n,i,j} \quad (3.17)$$

where:

- P_i is the normalised and averaged power output in bin i ;
- $P_{n,i,j}$ is the corrected power output of the j -th 10-minute dataset contained within bin i .

To ensure the statistical validity of these binned averages, a minimum data-count threshold is enforced. While the IEC 61400-12-1 standard permits a minimum of three 10-minute records per bin, this framework adopts a stricter threshold of six records (one hour), applied uniformly to both the baseline all-data reference curve and the regime-specific (stable, neutral, and unstable) power curves. Bins that fail to meet this threshold are excluded, so that the resulting curves reflect sustained aerodynamic trends rather than transient meteorological anomalies.

3.5.3. Annual Energy Production (AEP) Estimation

The estimated total energy production in one standard year is among the most important parameters derived during a wind resource assessment campaign. To calculate the Annual Energy Production (AEP), the empirical power curves generated via the Method of Bins are mathematically integrated with a probability model of the available wind speed resource. This metric directly determines the expected energy yield and economic viability of a wind project, and serves as the foundational figure for developers and investors. The following paragraphs detail the generic AEP estimation methodology adopted in this study, which enables a generalised comparison across diverse operational contexts.

To keep the research findings broadly applicable, a generic AEP calculation methodology is employed rather than a site-specific assessment. This idealised approach means the resulting analysis and performance comparisons are not tied to a single location, which extends the applicability of the results to other potential offshore deployment environments.

To model the wind speed probabilities across these generalised environments, the calculation incorporates the standard Rayleigh cumulative distribution function. The Rayleigh distribution is a simplified form of the Weibull distribution with a fixed shape factor (typically $k = 2.0$, Figure 3.9). Within industry standards and the relevant academic literature, this distribution is commonly accepted as an accurate approximation of the offshore climate

and resulting wind power density found across the Northwest Europe region (Osamu, Rikiji, and Gelder 2002). In accordance with IEC 61400-12-1, the probability $F(V)$ of the wind speed not exceeding a given velocity V is calculated using Equation 3.18:

$$F(V) = 1 - \exp\left(-\frac{\pi}{4} \left(\frac{V}{V_{\text{ave}}}\right)^2\right) \quad (3.18)$$

where V is the evaluated wind speed and V_{ave} is the annual average wind speed at hub height.




Figure 3.9: Rayleigh distribution, probability density function (a) & cumulative distribution function (b), Ren (2025).

Finally, the generic AEP is quantified by combining the completed stability-classified power curves with this Rayleigh resource model. The estimation procedure calculates the expected energy yield by multiplying the probability of the wind occurring within each discrete interval by the average power produced in that bin, and summing these yields across the entire operational range. Following the IEC 61400-12-1 standard, the expected Annual Energy Production is calculated using Equation 3.19:

$$AEP = N_h \sum_{i=1}^N [F(V_i) - F(V_{i-1})] \left(\frac{P_{i-1} + P_i}{2} \right) \quad (3.19)$$

where:

- AEP is the Annual Energy Production;
- N_h is the total number of hours in one year (8760);
- N is the total number of bins;
- V_i is the normalised and averaged wind speed in bin i ;
- P_i is the normalised and averaged power output in bin i .

This standardised methodology evaluates the expected power output across generic reference average wind speeds (V_{ave}) ranging from 4 m/s to 11 m/s. It also employs specific mathematical patching and stitching techniques to support comparability, even when the underlying regime-specific datasets have incomplete low- or high-wind speed bins.

4

Results

The purpose of this chapter is to present and critically analyse the findings from applying the established methodology to the operational offshore measurement campaign. The primary objective is to quantify the discrete influence of atmospheric stratification on wind turbine power performance and to evaluate the efficacy of various stability classification frameworks. To this end, the chapter is structured into four phases. First, the characteristics and baseline climatology of the filtered, free-stream dataset are established to confirm the underlying data quality. Second, the meteorological conditions are categorised into distinct stability regimes, providing a comparison between kinematic proxies and thermodynamic metrics. Third, the direct aerodynamic impact of these isolated stability regimes is quantified through power anomalies and extrapolated Annual Energy Production (AEP) deviations. Finally, the analysis expands to further performance dynamics, examining the seasonal variability of energy yield in shallow offshore environments and attempting to decouple kinematic proxies from thermal stratification to isolate their distinct aerodynamic effects.

4.1. Characteristics of the Finalized Dataset

The foundation of any offshore power performance assessment relies on the quality and representativeness of the underlying data. Following the filtering protocols detailed in Chapter 3, the finalised dataset comprises 8,597 valid 10-minute observations across 260 meteorological and operational variables. These measurements, recorded at an offshore site in the North Sea, span a continuous campaign from the beginning of 2025 through early 2026. Because raw operational data is obscured by non-atmospheric artefacts, such as turbine wakes, sensor inaccuracies, and grid curtailment, evaluating this baseline dataset is a necessary prerequisite before attempting to isolate the specific aerodynamic impacts of atmospheric stability.

The following subsections outline the characteristics of this retained free-stream dataset. First, the overall data yield is quantified to show the progressive attrition of measurements through the quality control filters. Next, the baseline climatology is presented to confirm the capture of a representative spectrum of operational wind conditions. Finally, the scatter distributions of the concurrent Lidar and SCADA power metrics are examined to demonstrate the effectiveness of the filtering process, while highlighting the residual aerodynamic variance that motivates the following stability analysis.

4.1.1. Data Retention and Attrition


Table 4.1 summarises the progressive data reduction throughout the quality control framework. The largest attrition occurs during the initial isolation of free-stream conditions. Specifically, the selection of the wake-free directional sector removes a large majority of the raw data. This is immediately followed by the target range-gate selection, which isolates the single Lidar measurement distance at 2.5 rotor diameters (2.5D) upstream, discarding 19 of the 20 available distance measurements per timestamp to obtain an undisturbed flow assessment.

Table 4.1: Data Filtering Summary

Filtering Stage	Remaining Observations	Data Removed [%]
Total Raw Dataset (Initial)	1,318,900	—
Wake-Free Sector Selection	440,340	66.6 %
Target Range-Gate Selection	22,017	95.0 %
Lidar Beam Availability	18,016	18.2 %
Kinematic Physicality Screening	17,997	0.1 %
Carrier-to-Noise ratio (CNR)	17,971	0.1 %
Curtailed Operation Exclusion	11,069	38.4 %
Fully operational (10 min)	8,597	22.3 %
Final Filtered Dataset	8,597	Total: 99.3 %

Figure 4.1 contrasts the retained clean dataset against the excluded raw measurements. The heavy scatter characterising the rejected data is primarily driven by neighbouring turbine wakes, aerodynamic blockage at non-optimal Lidar range gates, and standard SCADA instrumentation errors or grid curtailments (observable as a continuous horizontal line of zero-power output). Filtering these non-representative artefacts is essential to prevent artificial biases in the aerodynamic assessment.

Furthermore, the evaluated turbine features two distinct operational power configurations: a standard mode and a "boost" mode. Because the boost mode events made up a much smaller portion of the overall dataset, these data points were excluded to remove operational noise and keep a consistent, single aerodynamic baseline. The clean scatter data and the manufacturer reference curve used throughout this analysis therefore represent only the standard, non-boost operational profile.

**Figure 4.1:** Scatter power curve illustrating the finalized retained data (Clean) against the cumulative excluded data points (Noise).

4.1.2. Baseline Climatology

Isolating free-stream conditions is the necessary prerequisite before introducing atmospheric stability variables. Because the target turbine occupies a perimeter position within the offshore cluster (Figure 3.5), restricting the analysis to the discrete directional sector defined in Section 3.3.1 strips away artificial aerodynamic noise. This geographical advantage means the remaining dataset captures purely undisturbed incoming flow, free from the mechanical turbulence and velocity deficits induced by upstream rotors.

The physical characteristics and distribution of this undisturbed flow are visualised in Figure 4.2. Figure 4.2a presents a spatial reconstruction of the retained wind field upstream of the turbine, interpolating the measurements

from all 20 Lidar distance gates via the Ordinary Kriging method. This geostatistical approach was selected to provide unbiased, optimal estimations of the wind field by accounting for spatial autocorrelation (Isaaks and Srivastava 1989), implemented using the PyKriging package (Müller 2026). This interpolation presents the mean wind speed distribution across the selected sector. At the same time, Figure 4.2b displays the probability density of the filtered hub-height wind speeds. The histogram is overlaid with a theoretical Rayleigh distribution fit, showing that the surviving dataset captures a broad, continuous spectrum of operational wind conditions. This Rayleigh fit is also used for the following Annual Energy Production (AEP) extrapolations.




Figure 4.2: Baseline environmental climatology of the finalized measurement dataset.

4.1.3. Lidar and Power Curve Scatter Validation

To validate the efficacy of the applied filtering protocols, the retained 10-minute dataset is plotted to construct unclassified scatter power curves. Figure 4.3 presents these observations, plotting the power output against the concurrent wind speed measurements after all wake exclusion, instrumental quality control, and operational curtailment filters have been applied. To verify the remote sensing measurements, this visualisation compares the distributions derived from the nacelle-mounted Lidar (Figure 4.3a) alongside the conventional nacelle anemometer (Figure 4.3b).




Figure 4.3: Unclassified 10-minute scatter power curves of the filtered dataset, comparing remote sensing (a) and conventional anemometry (b).

A comparison of the resulting distributions confirms that the filtering process removed severe mechanical and operational outliers, yielding a relatively tight variance around the OEM reference curve for both instruments. However, a distinct behavioural difference between the sensors can be observed. The Lidar measurements align

more closely with the theoretical production curve, particularly in the lower wind speed ranges. By contrast, the nacelle anemometer shows an underestimation bias at these lower velocities.

4.2. Atmospheric Stability Characterization

Building on the filtered baseline climatology established in the previous section, the next phase of the analysis categorises these meteorological conditions into distinct atmospheric stability regimes. Because direct measurement of thermal stratification in offshore environments is often restricted by limited data availability, industry practitioners frequently rely on kinematic proxies or reanalysis data (ERA5) as practical substitutes for direct thermodynamic metrics. Therefore, a core objective of this research is to evaluate the reliability of these different classification methodologies and determine their efficacy in isolating aerodynamic performance variations. The following subsections first detail the distribution of stability regimes derived from kinematic proxies, then present the classifications based on thermodynamic metrics, and finally compare the two to determine which methods perform most reliably and whether kinematic variables can serve as trustworthy alternatives when thermodynamic data is unavailable.

The stability characterisation in this section is performed on the full meteorologically valid dataset. Because atmospheric stability is an independent property of the boundary layer, its classification does not require the turbine to be operational; the curtailment and full-availability filters of Chapter 3 are therefore not applied here, retaining 16,475 valid records. The power-performance analysis in Section 4.3 instead uses the operationally filtered subset.

4.2.1. Classification via Kinematic Proxies

In accordance with the analytical framework established in Section 3.4.1, this phase of the research uses wind shear and Turbulence Intensity (TI) as the kinematic proxies for stability classification. These variables are selected as they represent the most prevalent and accepted stability indicators within both wind energy industry standards and academic research. Analysing these parameters as a function of wind speed provides the first insight into the behaviour of the captured offshore boundary layer prior to full regime disaggregation.

The operational behaviour of these kinematic variables is first examined across the wind speed spectrum. Figure 4.4 presents the scatter distributions and mean trends for the wind shear exponent (α) and Turbulence Intensity, with the error bars denoting ± 1 standard deviation. At lower wind speeds, both variables exhibit high variance. Despite this wide scatter, the mean wind shear remains relatively stable across the entire operational envelope. By contrast, TI demonstrates a characteristic trend: high initial values drop steeply as wind speed increases. This reduction occurs because stronger, mechanically mixed wind fields homogenise the flow, lowering the relative turbulent fluctuations. However, at higher wind speeds, the mean TI begins to marginally pick up again, caused by the increased aerodynamic roughness of the sea surface due to wind-driven wave generation.




Figure 4.4: Behaviour of kinematic stability proxies across the operational wind speed spectrum. Error bars indicate ± 1 standard deviation from the mean.

Applying the discrete quantitative thresholds established in Table 3.1, the dataset is then partitioned into the three

primary stability regimes. Table 4.2 details the resulting percentage distributions derived independently from the wind shear and TI classifications. While both proxies identify unstable conditions as the most frequently occurring atmospheric state, their overall distributions diverge significantly. The shear-based classification yields a relatively balanced partition across the three regimes. By contrast, the TI-based classification skews heavily towards unstable conditions, categorising a much lower proportion of the dataset as stable. This divergence provides an initial indication that these kinematic variables, while related, do not respond identically to the underlying boundary layer characteristics.

Table 4.2: Percentage distribution of 10-minute data among atmospheric stability regimes derived from kinematic proxies.

Atmospheric Stability	Wind Shear Classification (α)	Turbulence Intensity Classification (TI)
Unstable Conditions	48.2%	64.7%
Neutral Conditions	18.0%	20.5%
Stable Conditions	33.8%	14.8%

4.2.2. Classification via Thermodynamic Metrics

While the previously analysed kinematic proxies offer a convenient estimation of boundary layer conditions, a thorough evaluation requires the direct quantification of thermal stratification. As detailed in the theoretical framework of Section 3.4.2, thermodynamic metrics directly capture the buoyancy forces that either suppress or enhance turbulent mixing. This phase of the analysis therefore classifies the dataset using two thermodynamic approaches: the Bulk Richardson Number (Ri_b) and the Monin-Obukhov Length (L).

Bulk Richardson Number

The Bulk Richardson Number provides a localised, site-specific quantification of atmospheric stability by directly comparing buoyancy forces against mechanical shear. To calculate this dimensionless parameter in accordance with Equation 3.5, a temperature gradient was required. The ambient air temperature was recorded using a Vaisala HMP155 sensor, an industry-standard meteorological probe, and the Sea Surface Temperature (SST) was measured using a CT15 infrared radiation pyrometer, which was actively corrected for downward sky reflections.

By integrating this temperature differential with the wind speed data from the Lidar, the behaviour of the boundary layer's thermal stratification becomes apparent. Figure 4.5a illustrates the scatter distribution of the calculated Bulk Richardson Number across the recorded wind speed spectrum. A physical trend emerges from this distribution: as the wind speed increases, the Ri_b values rapidly converge towards zero, which represents a neutral atmospheric state. This convergence occurs because, at higher wind velocities, the strong mechanical shear forces mix the boundary layer, overpowering the thermal convective forces and neutralising the atmospheric stratification regardless of the existing temperature gradient.

Monin-Obukhov Length

To complement the discrete temperature gradient analysis, the stability classification is further evaluated using the Monin-Obukhov Length (L). Because direct computation of L via high-frequency sonic anemometry was unavailable at the offshore site, this research derives the length scale using the empirical methodology outlined in Section 3.4.2. An established empirical expression (Equation 3.11) is applied to convert the Bulk Richardson Number directly into an estimated inverse Obukhov length ($1/L$), as shown in Figure 4.5b.

Following the calculation and refinement of these thermodynamic metrics, the dataset is segregated into the three defining stability regimes. Table 4.3 provides a quantitative summary of the resulting data distributions based on both the Bulk Richardson Number and the Monin-Obukhov Length classification thresholds.




Figure 4.5: Comparison of thermodynamic metrics against hub-height wind speed. Includes filtered data for normal turbine operations

The clustering results show that the two classification methods match closely, both with a bias towards unstable conditions at the site. The visual comparison in Figure 4.5 confirms that both metrics show similar physical behaviour across the operational envelope. In the low wind speed regions, there is a much higher spread of data, indicating that strong thermal gradients and buoyancy forces are driving the boundary layer. However, as the wind speed increases, mechanical shear forces begin to dominate, overriding the thermal effects and forcing the data points in both metrics to converge to a neutral state.

It is worth noting that when using the Obukhov length for stability classification, the absolute numerical thresholds defining these regimes must be adjusted according to the measurement reference height (z). Because boundary layer physics are governed by the dimensionless stability parameter (z/L) rather than L alone, standard literature thresholds derived from traditional 10-metre meteorological masts cannot be applied directly to modern multi-megawatt turbine hub heights. To maintain thermodynamic equivalence and prevent misclassification, the absolute limits for L must scale proportionally with the measurement height to accurately capture the corresponding atmospheric state at the rotor.

Table 4.3: Percentage distribution of 10-minute data among atmospheric stability regimes derived from thermodynamic metrics.

Atmospheric Stability Class	Bulk Richardson Number Classification (R_{ib})	Monin-Obukhov Length Classification (L)
Unstable Conditions	42.1%	42.1%
Neutral Conditions	32.3%	33.7%
Stable Conditions	25.6%	24.2%

4.2.3. Comparative Evaluation

The comparative evaluation of the classification methodologies requires establishing a baseline to serve as the analytical ground truth. Because the Bulk Richardson Number (R_{ib}) incorporates the complete set of direct instrument readings to derive virtual potential temperature, it is treated as the most accurate physical representation of the site's thermal stratification. When evaluated against this baseline, the Obukhov length (L), derived via the empirical expressions of Dyer (1974), yields closely comparable results. This confirms that, provided stability thresholds are properly scaled to the measurement reference height, empirical extrapolations remain reliable for modern offshore applications, extending beyond the traditional surface-layer limits.

A primary objective of this evaluation is determining the viability of kinematic proxies when thermodynamic instrumentation is unavailable. To assess their diagnostic value, the recorded kinematic variables are cross-

examined against the established Ri_b baseline regimes in Figure 4.6. The resulting boxplots confirm a fundamental physical correlation between the atmosphere's thermodynamic state and its kinematic flow characteristics. As anticipated by boundary layer theory, the strong vertical mixing of unstable conditions results in higher Turbulence Intensity (TI) and lower wind shear (α). By contrast, the thermally suppressed vertical momentum exchange of stable conditions yields the opposite behaviour: lower TI and higher shear.




Figure 4.6: Distribution of Turbulence Intensity (a) and Wind Shear Exponent (b) across the three atmospheric stability regimes defined by the Bulk Richardson Number. Boxes indicate the interquartile range (IQR), horizontal lines represent the median, and whiskers extend to the extremes of the distribution ($1.5 \times \text{IQR}$).

Despite this macroscopic correlation, substantial overlapping variance exists between the stability regimes. The interquartile ranges and whisker spans overlap considerably, showing that localised mechanical turbulence can obscure underlying thermal buoyancy signals, causing a single kinematic value to frequently correspond to multiple thermodynamic states. However, when assessing absolute classification performance via the percentage distributions (Tables 4.2 and 4.3), a hierarchy emerges. Although both kinematic proxies suffer from overlapping variance, the regime distribution produced by the wind shear classification aligns more closely to the Ri_b baseline than the TI classification. Therefore, while imperfect, wind shear serves as the preferred fallback metric for offshore environments lacking thermodynamic instrumentation.

To quantify the diagnostic capabilities of these kinematic proxies, their discrete stability classifications were evaluated against the Ri_b (BRN) baseline. The statistical agreement was measured using overall classification accuracy and Cohen's Kappa coefficient (κ), a metric that accounts for the possibility of agreement occurring by chance. The resulting confusion matrices, normalised by the baseline regimes, are presented in Figure 4.7.

The assessment reinforces the visual observations derived from the macroscopic distributions. The wind shear exponent (α) outperforms Turbulence Intensity (TI) as a stand-alone classification metric. The shear-based classification achieves an overall accuracy of 51.3% and a Cohen's Kappa of 0.27 (indicating "fair" statistical agreement). As shown in Figure 4.7a, it identifies unstable conditions in 71.8% of cases and captures 56.8% of true stable conditions. Its primary limitation lies in differentiating neutral conditions from unstable ones, frequently misclassifying true neutral states as unstable (49.5%).

In contrast, the TI-based classification demonstrates predictive limitations, yielding an overall accuracy of 42.3% and a Cohen's Kappa of 0.12 (indicating "slight" or poor agreement). Figure 4.7b shows a substantial predictive bias towards the unstable regime. While it correctly flags 97.3% of true unstable periods, it classifies 86.0% of true neutral periods and 39.5% of true stable periods as unstable. This high false-positive rate indicates that TI is sensitive to mechanically generated wake and surface-roughness effects, stripping it of the nuance needed to discern the underlying thermal stratification boundaries.

For this reason, while neither kinematic proxy fully replicates the thermodynamic baseline, wind shear retains markedly more diagnostic fidelity. It is therefore recommended as the alternative method for characterising atmospheric stability in offshore environments where temperature-gradient data is absent.




Figure 4.7: Confusion matrices evaluating the predictive performance of kinematic proxies against the R_{th} (BRN) baseline. Values represent the row-normalised recall percentages.

4.3. Impact of Stability on Power Performance

4.3.1. Environmental Normalisation and Correction

To ensure that the performance deviations quantified in this study are driven by atmospheric stability rather than ordinary environmental fluctuations, the raw SCADA power data must be normalised before the empirical power curves are constructed. This process isolates the aerodynamic response of the rotor by removing the effects of varying air mass and kinematic turbulence at the same time.

Air Density Normalisation

The absolute kinetic energy available in the free-stream flow is directly proportional to air density. Because the measurement campaign spans a full annual cycle, the boundary layer experiences considerable thermodynamic variance, ranging from cold, dense winter air masses to warm, less dense summer conditions. If left uncorrected, these seasonal density shifts would skew the binned power output, masking the true aerodynamic efficiency. Therefore, prior to analysis, all measured wind speeds are standardised to a constant reference air density utilising the site-specific pressure, temperature, and humidity data, so that the kinetic energy flux remains mathematically consistent across all evaluated timestamps.




Figure 4.8: Analysis of site air density variance and its seasonal impact on wind speed normalisation.

As illustrated in Figure 4.8, the thermodynamic behaviour of the boundary layer exhibits a pronounced seasonal cycle that drives the required wind speed corrections. The probability density distribution in panel (a) shows that the standard reference density used for the normalisation ($\rho_0 = 1.232 \text{ kg/m}^3$) is representative of the site's actual climatology, aligning with the dataset's annual mean ($\rho_{mean} = 1.233 \text{ kg/m}^3$). The practical impact of this variance is evident in the seasonal shifts shown in panel (b). During the colder winter months the air is cold and dense, carrying more absolute kinetic energy for a given measured wind speed. On the other hand, during the warmer summer months, the air is inherently less dense. To account for this reduced kinetic energy flux, the normalisation procedure scales the measured wind speeds downwards, producing negative corrections ($\Delta V < 0$). This seasonal oscillation shows the need for the normalisation process; without it, winter turbine performance would appear artificially inflated, while summer performance would be disproportionately penalised.

Turbulence Intensity (TI) Normalisation

While air density normalisation addresses the absolute available energy, Turbulence Intensity (TI) normalisation addresses how that energy is averaged by the turbine's control system. Applying the IEC 61400-12-1 Annex M methodology is important because a wind turbine's power curve is highly non-linear.

Because power extraction is capped at the turbine's rated limit, turbulent wind fluctuations behave asymmetrically in the transition region (the "knee" of the power curve). If the mean 10-minute wind speed sits just below the rated capacity, turbulent gusts will frequently push the instantaneous wind speed into the rated zone, where the turbine pitches its blades and cannot capture the excess energy. By contrast, turbulent lulls drop the wind speed into the partial-load region, resulting in an immediate loss of power. Hence, highly turbulent winds yield less average power at the transition region than low-turbulence winds of the exact same mean velocity.

To remove this kinematic artefact, the algorithm simulates a theoretical zero-turbulence power curve. It then calculates the expected power output at the measured TI, compares it against the expected power at a fixed reference TI (defined here as the mean TI across the dataset), and applies the difference as a discrete power correction (ΔP), as explained in Annex C. The impact of this normalisation is detailed in Figure 4.9.




Figure 4.9: Validation of the Turbulence Intensity normalisation process, illustrating the algorithm's effect on the transition region and the magnitude of power adjustments based on measured TI.

As visualized in Figure 4.9b, the required power correction depends strongly on both the normalised wind speed and the ambient turbulence. A prominent downward "drop", representing a negative power correction, is observed at the transition region (normalised wind speeds between 0.85 and 1.0) for low-TI data points (blue).

This negative correction supports the physical validity of the algorithm. Because low-turbulence wind fields do not suffer from the asymmetric energy clipping described above, their raw empirical power curves exhibit a sharp, efficient knee. To normalise these calm data points to the higher, average dataset reference TI, the algorithm penalises them by subtracting power, effectively smoothing their transition curve to match the baseline. In contrast, high-TI data points (red) receive a positive power correction in this same region, restoring the energy mathematically lost to high-frequency gusts.

Beyond this sensitive transition zone, the magnitude of the applied correction is mild. Across the lower partial-load region and the fully rated power region, the algorithm mostly serves to smooth isolated statistical outliers. By removing these kinematic peaks across the operational envelope, the dataset is stripped of its mechanical artefacts, so that any remaining performance stratification is driven by thermodynamic boundary layer stability.

4.3.2. Baseline and Stability-Classified Power Curves

Before analysing the aerodynamic impacts of atmospheric stratification, it is necessary to contextualise the volume and statistical reliability of the retained dataset. The filtering protocols established in Chapter 3, specifically the wake sector exclusions and Lidar data quality controls, were needed to isolate free-stream conditions. While these filters reduced the overall data volume, they provided the high fidelity required for comparative aerodynamic research. Rather than adhering rigidly to the IEC (2017) completeness criteria, which are in principle designed for commercial warranty validation, this assessment evaluates the data distribution to prioritise statistical robustness. Table 4.4 summarises the general characteristics of the finalised dataset used for the following power curve construction. Table 4.4 summarises the characteristics of the operationally filtered subset used for power curve construction.

Table 4.4: General statistics of the filtered, free-stream dataset used for power curve construction.

Dataset Characteristic	Value
Total Sampled Data (Hours)	717.17 h
Total Valid 10-Minute Observations	4,303
Maximum Normalised Wind Speed Coverage	1.27
Minimum Data Points Required per Bin (10-min intervals)	6 (1 hour)

While the dataset exhibits a deficit in valid high-wind-speed observations (falling short of capturing 1.5 times the wind speed at 85% rated power), it covers the partial-load and transition regions where stability impacts are most prominent. Figure 4.10 presents the baseline operational profile derived from all filtered concurrent measurements, plotting binned active power and the calculated power coefficient (C_p) against the standardised hub-height wind speed. To preserve commercial sensitivity, both active power and incoming wind speed are normalised, and the measured operational profile is overlaid directly against the normalised original equipment manufacturer (OEM) reference power curve for validation.




Figure 4.10: Baseline measured power curve and power coefficient (C_p) curve derived from the filtered dataset.

To support statistical reliability and a representation of the underlying distributions, the construction of each regime-specific power curve follows fixed binning requirements. The framework retains only wind speed bins

containing a minimum of six data points. While the baseline IEC standard permits a minimum of three points, increasing this threshold to six is a deliberate adjustment.

Building on this validated baseline, the finalised dataset is partitioned into the stable, neutral, and unstable atmospheric regimes using the Bulk Richardson Number (R_{ib}) classification methodology. For statistical reliability and consistency, the construction of each regime-specific power curve follows the same binning requirements outlined previously, such as retaining only wind speed bins containing a minimum of six data points. Figure 4.11 provides a direct comparison of these regime-specific power curves, with shaded regions representing ± 1 standard deviation (σ) to show the observed variance within each atmospheric state.

A preliminary analysis of these operational curves shows distinct patterns across the velocity spectrum. At very low wind speeds, and again after reaching the rated wind speed (marked by the vertical dashed line), the aerodynamic differences between the stability regimes are minimal. Deviations do occur, however, within the partial-load region. This performance stratification is most pronounced approaching the "knee" of the power curve, pointing to an aerodynamic sensitivity to thermal stratification that is quantified and investigated in the following sections.




Figure 4.11: Comparison of wind turbine power curves disaggregated by atmospheric stability regime. Shaded areas represent ± 1 standard deviation.

4.3.3. Quantification of Performance Deviations

Power Deviation

To quantify the impact of atmospheric stratification on turbine efficiency, the power output is evaluated as the deviation, or "anomaly", from the theoretical manufacturer power curve. While the Bulk Richardson Number (R_B) is often treated as the standard thermodynamic metric, it is important to compare how different atmospheric classification methodologies partition the data and capture these aerodynamic performance shifts. By isolating the discrete power differential (ΔP) across distinct stability regimes, the absolute influence of the boundary layer conditions can be analysed.

When evaluating these performance deviations, it is important to verify that the observed differences are physically meaningful rather than artefacts of statistical noise. Figure 4.12 visualizes the power anomaly across the normalised wind speed spectrum for four classification methods: wind shear, turbulence intensity, the Bulk Richardson Number, and the Obukhov length. In these subplots, the solid horizontal baseline at 0 kW represents exact agreement with the manufacturer's reference power curve. The horizontal axis is normalised against the rated wind speed, with the vertical dashed line at 1.0 marking the transition into the rated power region. To support statistical reliability, the data processing framework filters out wind speed bins with low data volumes. The resulting bin-averaged deviations are plotted alongside shaded envelopes representing the 95% Confidence Interval (CI), derived from the Standard Error of the Mean (SEM), providing a threshold for statistical significance.




Figure 4.12: Comparison of absolute power anomalies (ΔP) from the manufacturer baseline across four atmospheric classification methods: (a) Wind Shear, (b) Turbulence Intensity, (c) Bulk Richardson Number, and (d) Obukhov length. Shaded envelopes denote the 95% Confidence Interval. The vertical dashed line indicates the normalised rated wind speed.

Analysis of the anomaly profiles shows several distinct patterns. Most prominently, across all classification methodologies, the power anomalies consistently converge towards zero as the normalised wind speed approaches and exceeds 1.0. This shows that in the rated power region, the turbine's active pitch control governs energy extraction, largely mitigating the influence of atmospheric stability.

Within the partial-load region (normalised wind speeds < 1.0), by contrast, a clear performance stratification emerges. Unstable atmospheric conditions consistently drive aerodynamic overperformance relative to the manufacturer baseline, while stable conditions induce noticeable underperformance. This regime separation is clearer when using the thermodynamic metrics (Figure 4.12c and 4.12d). While the kinematic proxies are noisier, a difference in fidelity is observed: the shear-based classification (Figure 4.12a) approximates the thermodynamic behaviour more closely than the turbulence intensity proxy (Figure 4.12b), particularly in the positioning of the neutral regime relative to the zero-anomaly baseline. This suggests that vertical wind gradients offer a more reliable kinematic reflection of thermal stratification than local turbulence levels, although direct thermodynamic classification remains necessary for a precise quantification of rotor aerodynamics.

AEP Deviation

While quantifying the discrete power deviation (ΔP) provides insight into the turbine's instantaneous aerodynamic behaviour, translating these variances into long-term energy yield is essential for practical wind farm assessment. Using the methodology detailed in Section 3.5.3, the impact of atmospheric stability is projected onto the Annual Energy Production (AEP). This extrapolation is achieved by applying the regime-specific power curves to a Rayleigh wind speed distribution, calculating the Generic AEP to normalise the data for comparative purposes.

It must be noted that standard IEC guidelines state that the measured AEP should account for at least 95% of the extrapolated AEP to be considered fully complete, as mentioned in Table 4.4. Therefore, it must be stated that the AEP estimates derived in this analysis do not constitute a formal power performance test equivalent to those used for commercial warranty validation. A formal warranted test requires strict, certifiable adherence to all IEC

61400-12-1 measurement criteria, which are not fully satisfied by the available operational dataset. However, as established in Section 4.3.2, while these criteria are mandated for commercial guarantees, the current dataset is adequate for the specific research purposes of this thesis. Here, the Generic AEP serves only as a standardised comparative metric to quantify the relative sensitivity of long-term energy yield to atmospheric stability.

To assess the site's energy potential, the extrapolated AEP was calculated across the full spectrum of annual average wind speeds recommended by the IEC (2017) standard. To establish a standardised comparative baseline, all calculated yields were normalised against the theoretical AEP derived from the manufacturer's reference curve at $V_{ave} = 8$ m/s. Figure 4.13 presents these normalised AEP distributions, showing the projected energy yield assuming the turbine operated continuously under a single, isolated stability regime for an entire year.




Figure 4.13: Extrapolated Annual Energy Production (AEP) by stability regime across standard annual average wind speeds. All values are normalised against the manufacturer's theoretical AEP at $V_{ave} = 8$ m/s.

The trends in the projected long-term energy yield mirror the power deviations identified in Section 4.3.3. Across all annual average wind speeds, unstable atmospheric conditions consistently drive an overperformance in energy yield relative to the manufacturer's baseline. Stable conditions, by contrast, sharply curtail energy extraction, resulting in an underestimation of the theoretical yield. Neutral conditions also show a slight underperformance but remain the closest approximation to the OEM warranty curve. This indicates that traditional, neutrally-assumed power curves overestimate production in stable offshore environments.

While this stratification persists across the operational envelope, the $V_{ave} = 8$ m/s scenario warrants specific evaluation as it closely aligns with the long-term climatology of the offshore site. Table 4.5 quantifies the AEP deviations relative to the OEM baseline at this reference wind speed, showing the sensitivity of the energy yield to different kinematic correction methodologies.

Table 4.5: Percentage deviation of extrapolated AEP by stability regime relative to the OEM reference at $V_{ave} = 8$ m/s, illustrating the sensitivity of yield results to different normalisation and kinematic correction methodologies.

Stability Regime	Air Density Norm. Only	Wind Shear Correction	TI Correction	TI & Shear Correction
All Data	+0.18%	-0.53%	+0.01%	-0.70%
Stable	-1.59%	-2.86%	-1.81%	-3.08%
Neutral	-0.68%	-1.40%	-0.90%	-1.59%
Unstable	+1.82%	+1.59%	+1.80%	+1.57%

As detailed in Table 4.5, the pattern remains consistent regardless of the applied data normalisation method. However, wind shear corrections are not taken into account for the final assessment. This is due to limitations in the deployed Lidar technology, specifically its inability to simultaneously measure wind velocities at a minimum of three distinct vertical heights within the rotor swept area. This constraint prevents an accurate calculation of the Rotor Equivalent Wind Speed (REWS), making shear-based normalisations unreliable for this dataset. The analysis therefore uses the air density and Turbulence Intensity (TI) corrections as the final results. As shown in the table, the TI correction provides the best match between the aggregated dataset ("All Data") and the manufacturer's warranty curve, with a negligible +0.01% deviation. This method also captures a symmetrical aerodynamic response to thermal stratification, with an energy deficit of -1.81% under stable conditions and a proportional surplus of $+1.80\%$ under unstable conditions.

4.3.4. Indicative Socio-Economic Implications

The preceding AEP deviations answer the third research question posed in Section 1.3 in aerodynamic terms, but their practical weight is clearer once expressed in monetary form. This subsection translates the regime-specific yield deviations of Section 4.3.3 into indicative economic and social terms. It is not a financial appraisal: the generic AEP is itself a normalised comparative metric rather than a site yield, so the figures below are order-of-magnitude indicators of why the stability-driven spread matters, not a revenue forecast for any specific asset.

To keep the translation independent of project size, the deviations are expressed as a revenue sensitivity per GWh of baseline annual energy yield. The price used is a representative North-West European wholesale day-ahead value of approximately €82/MWh, obtained by converting the 2025 EU average wholesale price of USD 95/MWh (IEA 2026) at the mid-2026 exchange rate of about €0.86 per US dollar. Each entry is the percentage deviation applied to this price, so multiplying by a project's own baseline yield recovers an absolute figure. Table 4.6 lists these sensitivities for the four correction methods, with the air-density and TI correction adopted here shown in bold. Under that correction, an all-stable year would forgo about €1,480 per GWh of baseline yield and an all-unstable year would add a similar amount, a spread of roughly €2,960 per GWh between the two regimes. The figures scale linearly with price, so a different price assumption shifts them in proportion without changing the relative picture.

Table 4.6: Indicative revenue sensitivity by stability regime (€/yr per GWh of baseline annual yield) at $V_{ave} = 8$ m/s, assuming a wholesale price of €82/MWh. The adopted air-density and TI correction is shown in bold.

Stability Regime	Air Density Norm. Only	Wind Shear Correction	TI Correction	TI & Shear Correction
Stable	-1,304	-2,345	-1,484	-2,526
Neutral	-558	-1,148	-738	-1,304
Unstable	+1,492	+1,304	+1,476	+1,287

Two points follow for system integration. First, across the measured mix of conditions the aggregate deviation is small (+0.01% for the adopted correction), so the result is not that the turbine loses energy on balance; rather, a power curve assumed to be neutral misallocates expected yield between conditions and biases the regime- and season-dependent estimates that feed P50/P90 yield estimates, financing cases, and the forecasting used for grid integration. Second, this bias does not average out where stability is seasonally persistent. As shown later in Section 4.4.1, the site favours stable conditions in winter and unstable conditions in summer, so a neutral assumption distorts the seasonal yield profile that operators and offtakers depend on rather than cancelling over the year.

The same sensitivities give a sense of scale beyond a single turbine. Applied across 1 TWh of annual offshore generation, the stable-to-unstable spread is of the order of €3.0 million per year, and the corresponding energy difference of about 36 GWh is enough to cover the annual electricity use of roughly 10,000 average European households (Odyssey-MURE 2024). These scaled figures are illustrative only: the assessed turbine sits in wake-free flow on the array perimeter, and interior turbines experience a different inflow, so the percentage cannot be applied uniformly across a farm. The wider value is therefore methodological: stability-aware power curves give more accurate and more bankable offshore yield assessments, and a clearer basis for integrating variable offshore wind into the energy system.

4.4. Advanced Performance Dynamics

While the preceding sections established the impact of atmospheric stratification on absolute power production and long-term energy yield, standard generic assessments often obscure the dynamic, time-dependent nature of the offshore boundary layer. To fully understand the operational efficiency of this site, it is necessary to go beyond standard metrics and examine the performance dynamics driven by local atmospheric phenomena. Specifically, this section examines how seasonal thermodynamic cycles influence long-term energy yield, investigates the spatial evolution of the upstream wind profile to identify aerodynamic blockage effects, and uses a decoupling methodology to isolate the distinct aerodynamic impacts of kinematic turbulence from underlying thermal stratification.

4.4.1. Seasonal Variations in Stability and Performance

The relationship between atmospheric stability and temporal cycles is highly dependent on a site's geographical environment. In typical onshore environments, stability exhibits a strong diurnal correlation; daytime heating of the landmass predominantly drives unstable convective conditions, while nighttime cooling leads to stable, stratified layers, as shown in the field observations of St. Martin et al. (2016). However, this diurnal dependency is mostly absent in offshore environments. The high specific heat capacity of water prevents rapid temperature fluctuations, causing the Sea Surface Temperature (SST) to remain relatively constant on a daily basis. Instead, offshore stability is driven by broader advection processes and seasonal macroscopic shifts.




Figure 4.14: Seasonal stability characterization and underlying thermodynamic drivers at the offshore site.

An analysis of the distributions presented in Figure 4.14 reveals a distinct seasonal shift in both the boundary layer stability and its underlying thermodynamic drivers. Figure 4.14a shows that during the winter and spring periods, the site exhibits a strong bias towards stable conditions. In contrast, the summer and autumn months are dominated by predominantly unstable regimes. The physical mechanism driving this bimodal distribution is tracked in Figure 4.14b. The data show a distinct inversion of the temperature gradient across the year: the ambient

air is generally warmer than the sea surface (a positive $T_{\text{air}} - T_{\text{SST}}$ differential) during the spring, whereas the sea surface remains notably warmer than the incoming air masses (a negative differential) throughout the autumn.

This seasonal thermodynamic signature is characteristic of shallow coastal waters. Because the evaluated site is situated in relatively shallow water depths of 18 to 25 metres, the sea column gains and dissipates heat more rapidly than deep ocean environments, yet still at a slower rate than the adjacent landmass. Therefore, during the spring months, warm continental air or advecting weather fronts flow over the still-cooling sea. This process cools the boundary layer from below, suppressing vertical mixing and creating the strong temperature inversions that drive stable conditions (Barthelmie et al. 2015). On the other hand, in the autumn, cold continental air advects over the relatively warm, sun-heated shallow waters. This heats the lowest layers of the atmosphere from below, triggering the convective buoyancy and vertical mixing that align with the observed high frequency of unstable conditions.

These seasonal stability shifts have a substantial impact on turbine performance. Naturally, the energy yield is generally higher during the winter season simply due to the higher baseline wind speeds. However, while the overall ordering of aerodynamic efficiency remains constant across the year, unstable conditions perform best, followed by neutral, and finally stable, the proportional variation within those specific regimes yields useful insights.




Figure 4.15: Seasonal comparison of Annual Energy Production (AEP) across stability classes. Values are normalised against the theoretical yield derived from the OEM reference power curve.

As shown in Figure 4.15, the performance output of the various stability regimes reacts very differently to seasonal changes. Stable and neutral conditions show a significantly higher deviation in AEP between the summer and winter seasons than unstable conditions, where the seasonal performance difference effectively disappears.

This behaviour can be explained by the fundamental fluid dynamics of the atmospheric boundary layer. Under unstable conditions, convective heat transfer drives strong vertical mixing. This mixing homogenises the wind field across the entire rotor swept area, breaking down wind shear and eliminating directional veer. Because the flow is thoroughly mixed, the rotor continuously experiences a blunt, uniform wind profile. Once this convective homogenisation is achieved, the background seasonal temperatures become aerodynamically irrelevant; the turbine extracts power from a highly uniform flow field, resulting in consistent seasonal performance.

Stable conditions, by contrast, suppress vertical mixing, decoupling the atmospheric layers. Under these stratified conditions, the turbine rotor experiences high wind shear, variable inflow angles, and wind veer. As established in Section 4.3.3, this offshore site exhibits reduced aerodynamic efficiency during stable conditions. This occurs because the limited vertical momentum exchange and directional veer alter the aerodynamic angle of attack across the rotor.

To quantify these boundary layer characteristics, Table 4.7 details the mean wind shear and occurrence frequencies across the winter and summer months (with occurrence frequencies taken directly from the distributions presented in Figure 4.14a).

Table 4.7: Seasonal distribution of atmospheric stability regimes and their corresponding mean wind shear (α).

Stability Regime	Winter		Summer	
	Frequency [%]	Mean Shear (α)	Frequency [%]	Mean Shear (α)
Stable	50.6	0.1544	15.3	0.2502
Neutral	19.2	0.1038	19.7	0.1931
Unstable	30.2	0.0477	65.0	0.0712

Site data indicate that the mean wind shear (α) during summer stable events (0.2502) is higher than during winter stable events (0.1544). This difference suggests that summer stable conditions feature a more stratified boundary layer, likely driven by temperature inversions as warm continental air settles over cooler shallow waters. Because stratification inherently reduces aerodynamic efficiency at this site, the stronger stable conditions in summer, characterised by higher shear, increased directional veer, and limited kinetic energy mixing, lead to greater energy losses. As a result, the summer-stable power curve underperforms compared to the winter-stable curve. This seasonal variation in boundary layer stability explains why the extrapolated Annual Energy Production (AEP) decreases during the summer months, even when normalised against a constant baseline wind speed distribution.

4.4.2. Upstream Flow and Blockage Effects

To further investigate the aerodynamic interactions preceding the turbine, the evolution of the incoming flow was analysed using the full measurement volume of the long-range, nacelle-mounted Lidar. By extracting concurrent wind speeds from all 20 available distance gates, spanning from 57 m to 657 m upstream, a continuous wind speed profile was reconstructed. To isolate the localised aerodynamic effects, these profiles are normalised against the measurement gate situated at 407 m (approximately 2.5 rotor diameters upstream). This distance is selected as the free-stream reference: it is far enough to remain largely uninfluenced by the rotor's immediate induction zone, yet close enough to limit the measurement uncertainties associated with extended Lidar ranges.




Figure 4.16: Normalised upstream wind speed profiles across varying distance gates, categorized by stability regime. The panels represent: (a) partial load region, (b) transition region, (c) above-rated wind speed region, and (d) all available data. Profiles are normalised against the 407 m reference gate.

Figure 4.16 presents these normalised wind profiles, disaggregated by atmospheric stability regime and operating region, with shaded envelopes indicating the 95% Confidence Interval.

An initial assessment of the aggregated data (Figure 4.16d) reveals that the neutral and stable profiles behave in accordance with established aerodynamic theory. Neutral conditions maintain a stable, undisturbed signal across the upstream domain, overcoming the turbine's blockage effect until immediately preceding the rotor plane. By contrast, the stable regime shows a sharp, pronounced deceleration as it approaches the turbine. Because vertical momentum exchange is thermally suppressed under stable stratification, the incoming flow cannot easily expand over the rotor and is instead forced to decelerate and divert laterally, resulting in an amplified aerodynamic blockage effect.

The unstable flow profile, however, exhibits an anomalous, nearly linear variance extending from the immediate rotor proximity to the furthest upstream distances, failing to flatten into the expected free-stream baseline. Sub-regional analysis reveals that while this structural behaviour is only apparent in the partial-load region (Figure 4.16a), the magnitude of the deviation here is smaller than in the aggregated dataset (Figure 4.16d). This indicates that the pronounced variance observed in the overall dataset is not solely an artefact of active power extraction, but is also influenced and artificially skewed by non-operational data points that amplify these remote sensing errors.

This deviation is unlikely to represent a true physical acceleration of the mean flow towards the turbine; rather, it points to remote sensing artefacts happening in low wind speeds. As the measurement distance increases, the spatial separation between the Lidar's diverging beams widens considerably, as illustrated in Figure 2.6. Under unstable atmospheric conditions, which are characterised by intense vertical mixing, elevated Turbulence Intensity (TI), and large, localised convective gusts, the Lidar assumption of horizontal and vertical flow homogeneity across the scanning volume breaks down. The kinematic differences encountered simultaneously by the left and right beams induce substantial beam decorrelation errors, artificially skewing the reconstructed wind speed profile at extreme distances. Fully isolating and quantifying these sensor-induced artefacts requires further dedicated investigation.

To ensure that the performance variations observed across the stability regimes are driven by true aerodynamic differences rather than measurement distortion at the reference gate (2.5D), the impact of vertical wind shear on the Lidar measurements was investigated. Under varying atmospheric stability, distinct shear profiles alter the kinetic energy distribution across the rotor swept area. A highly stratified (stable) profile could theoretically cause the Lidar's internally processed hub-height wind speed (HWS) to diverge from the true Rotor Equivalent Wind Speed (REWS), introducing a measurement bias.

To quantify this potential bias, theoretical vertical shear profiles were constructed using the power law equation 2.2, using the mean shear exponent (α) and mean hub-height wind speed for each regime. The Lidar-reported hub-height wind speed was then compared directly against the mathematically integrated REWS across the rotor swept area.




Figure 4.17: Theoretical vertical shear profiles by stability regime, overlaid with Lidar beam heights and rotor span.

Table 4.8: Lidar measurement bias analysis at the 2.5D reference gate.

Regime	α	HWS [m/s]	REWS [m/s]	Bias [%]
Neutral	0.129	13.22	13.09	1.01
Stable	0.182	10.15	10.01	1.33
Unstable	0.055	9.37	9.32	0.47

The results of this analysis are detailed in Table 4.8 and shown in Figure 4.17. The maximum calculated bias between the Lidar-reported HWS and the REWS is minimal, peaking at approximately 1.3% under stable conditions. Because this measurement discrepancy remains consistently low across the stable, neutral, and unstable regimes, vertical shear variations can be taken not to distort the reference wind speed at 2.5D in any meaningful way. The performance differences and flow decelerations observed near the rotor are therefore primarily attributable to true aerodynamic phenomena rather than remote sensing artefacts at the reference gate. It should be noted that this calculation relies on a theoretical power-law extrapolation based on only two measurement heights; while absolute REWS values carry some inherent uncertainty, the relative inter-regime bias is low enough to rule out systematic distortion.




Figure 4.18: Correlation coefficient between the Lidar and nacelle SCADA hub-height wind speeds at zero time lag, plotted against Lidar distance gate for all 20 gates (57–657 m).

A further observation comes from the temporal-alignment step in Section 3.2.1. Besides fixing the time lag, the cross-correlation also records how closely the two instruments agree at each measurement distance. Figure 4.18 plots this zero-lag correlation coefficient against distance gate for all 20 gates, from 57 m to 657 m. Across most

of the range the agreement is high and fairly flat, holding to roughly $r \approx 0.93 - 0.97$. The exception is the nearest gate at 57 m, which sits clearly below the rest at about $r \approx 0.83$; no other gate shows a comparable drop. After a maximum in the near range (around 130 - 200 m), the correlation eases off slowly with distance and levels out near $r \approx 0.94 - 0.95$, with small gate-to-gate scatter rather than a steady decline.

The low value at the nearest gate is isolated: it appears only at 57 m and does not build up over the near-rotor region, so it does not follow the gradual pattern that a flow-development effect across the induction zone would produce. The slow easing at longer ranges is consistent in direction with the larger probe volume and wider beam separation at distance noted in Section 2.4.2, though the present data do not isolate a single cause.

4.4.3. Isolating Kinematic Turbulence from Thermodynamic Drivers

To further isolate the specific drivers of turbine performance, it is necessary to decouple the kinematic effects of turbulence from the underlying thermodynamic stability. Inspired by the collaborative methodologies established by the Power Curve Working Group, Lee et al. (n.d.), this decoupling can be shown through a power deviation matrix. This matrix maps the operational performance space by evaluating the difference between the actual measured energy yield and the theoretical yield derived from the OEM reference power curve, binned by both Normalised Wind Speed and Turbulence Intensity (TI).

By analysing this performance space both before and after applying the IEC Annex M TI-normalisation, the distinct aerodynamic impact of the turbulent kinetic energy can be isolated from the underlying atmospheric stratification. Figure 4.19 presents these paired matrices, where deviations are expressed as fractional deficits or surpluses relative to the ideal OEM baseline.




Figure 4.19: Power deviation as a fraction of the theoretical OEM yield, mapped across Normalised Wind Speed and Turbulence Intensity. Cool colours (blue) indicate overperformance; warm colours (red) indicate underperformance.

An initial assessment of the raw performance space (Figure 4.19a) shows several distinct aerodynamic regimes that align with established boundary layer theory. First, at wind speeds above the rated capacity ($V_{\text{norm}} \geq 1.0$), the power deviations reliably converge towards zero regardless of the prevailing TI. In this region, the turbine actively regulates its aerodynamic efficiency via blade pitch control to maintain constant rated power, neutralising any environmental or stability-driven aerodynamic advantages.

However, within the highly sensitive partial-load region ($V_{\text{norm}} < 1.0$), a vertical gradient emerges. At very low TI values, which act as a primary kinematic proxy for stable atmospheric stratification, the turbine shows sustained underperformance. In the upper TI bounds associated with unstable, highly convective atmospheres, by contrast, the matrix reveals distinct overperformance relative to the OEM baseline.

A further insight emerges when comparing this raw data to the TI-normalised matrix (Figure 4.19b). Following the mathematical correction, the absolute magnitude of the deviations is dampened, represented visually by a shift towards paler colours. This reduction is consistent, as the correction algorithm neutralises the artificial energy peaks and troughs caused by high-frequency turbulent fluctuations.

Crucially, while the absolute intensity decreases, the underlying aerodynamic pattern remains intact. The low-TI bins continue to show structural underperformance, and the high-TI bins maintain an energy surplus. The persistence of this gradient after the explicit removal of turbulence effects is strong empirical evidence for this study's primary hypothesis: the observed power deviations are not merely an artefact of varying turbulence. Rather, they are driven by the underlying thermodynamic stratification that persists in stable and unstable boundary layers regardless of the isolated turbulent kinetic energy.

5

Discussion

This research aimed to reliably decouple the underlying thermodynamic stability of the offshore boundary layer from interconnected kinematic variables to quantify its isolated impact on wind turbine power performance. The discussion is structured into two sections. First, Section 5.1 addresses the inherent methodological constraints and uncertainties associated with the measurement campaign and the analytical framework. Then, Section 5.2 contextualises these findings within the broader offshore wind sector. It demonstrates how the quantified aerodynamic deviations can be used to reduce uncertainty in pre-construction yield assessments, improve grid forecasting capability, and optimise the design of future remote sensing campaigns.

5.1. Methodological Constraints and Uncertainties

While the data processing pipeline and normalisation frameworks used in this study are robust and aligned with contemporary industry standards, the broad applicability of the findings is bounded by specific methodological constraints. These limitations emerge primarily from instrumental configurations, the physical mechanics of remote sensing, and the intrinsic site-specificity of atmospheric research.

Equipment Limitations

Although the measurement campaign used advanced remote sensing technology and accurate thermodynamic sensors, the experimental setup lacked high-frequency 3D sonic anemometers. This absence prevented the direct measurement of turbulent kinematic surface heat fluxes. Therefore, the most rigorous, direct application of Monin-Obukhov Similarity Theory (MOST) via eddy covariance was not possible. This specific equipment constraint necessitated the adoption of alternative, yet established, stability classification methods to measure the thermodynamic state of the boundary layer without high-frequency flux data.

Furthermore, the campaign lacked an independent meteorological mast, as the original site mast had been disassembled prior to this study. This removed the traditional reference for cross-validating the remote sensing data, forcing the temporal and directional alignment procedures to rely exclusively on the turbine's internal SCADA nacelle anemometry. While the presence of a met mast and sonic anemometry would have allowed a more accurate boundary layer analysis, their absence mirrors the practical reality of modern commercial offshore wind farms. Overcoming these specific equipment deficits ultimately strengthens the applied methodology, demonstrating that high-fidelity stability filtering and yield correction can be carried out using the constrained datasets commonly faced by the industry.

Nacelle Lidar Technical Limitations

The reliance on a forward-looking, nacelle-mounted Lidar introduced distinct aerodynamic and geometric measurement challenges. Foremost was the inability to reliably implement the Rotor Equivalent Wind Speed (REWS) correction. While REWS is the industry standard for addressing extreme wind shear across large rotor swept areas, the geometric spread of the deployed Lidar could not simultaneously capture the required minimum of three vertical measurement heights across the rotor disk. This geometric constraint forced the methodology to abandon

shear normalisation and rely on air density and Turbulence Intensity (TI) corrections to isolate the stability effects.

Additionally, the physical operation of the multi-beam Lidar introduced spatial uncertainties at extended measurement ranges. As observed during the upstream flow analysis, the spatial separation between the diverging Lidar beams becomes substantial at far distances. Under unstable atmospheric conditions, which are characterised by localised convective gusts, this wide horizontal beam spread violates the instrument's internal assumption of flow homogeneity (Held and Mann 2019). The resulting beam decorrelation introduces artificial skewing into the reconstructed wind speed profile. Compounding this, the finite probe volume over which each beam integrates the radial velocity acts as a spatial low-pass filter, attenuating the smaller turbulent scales and biasing the Lidar-derived turbulence intensity on which the stability correction depends (Sathe et al. 2015; Newman and Clifton 2017). This geometric divergence also produces vertical discrepancies; as detailed in Annex B and illustrated in Figure B.1, the cone angle of the forward-looking Lidar dictates that at extended ranges the lowest measuring beams intersect elevations progressively closer to the ocean surface. This vertical spread forces the instrument to sample distinct aerodynamic layers, some influenced by surface wave friction, which increases the uncertainty of the long-range measurements and can reduce their reliability in convective offshore conditions (Robey and Lundquist 2022).

A brief comparison with continuous-wave (CW) Lidar clarifies that these range-dependent effects follow in part from the pulsed measurement principle itself. A pulsed Lidar gates the backscattered signal by time of flight, so its probe length is fixed by the pulse duration and stays roughly constant with distance, allowing the instrument to sample several ranges far upstream of the rotor. A CW system instead focuses the beam at a chosen distance: this yields finer spatial resolution close to the instrument, but the probe length grows with the square of the focus distance, so resolution degrades quickly at longer ranges and CW systems are generally limited to short-range, hub-height profiling (*Remote Sensing for Wind Energy* 2011). The two principles therefore carry different uncertainties: the pulsed system is subject to the beam-separation and probe-volume averaging effects described above, while a CW system trades these for stronger range-dependent volume averaging and a sensitivity of its Doppler spectrum to focus distance (Peña, Yankova, and Mallini n.d.). For the present application, profiling the inflow several rotor diameters upstream, the multi-range reach of the pulsed system was the more suitable choice; a CW instrument would have lowered volume-averaging uncertainty only in the near field, at the cost of the upstream coverage the induction-zone analysis requires (Antoine Borraccino et al. 2017).

Site-Specific Climatology and Turbine Characteristics

Finally, it must be acknowledged that the exact magnitudes of the quantified power deviations (e.g., the specific energy surplus in unstable air and the proportional deficit in stable air) are inherently tied to the specific climatology of the North Sea site and the aerodynamic control logic of the target multi-megawatt turbine. This site-specificity is an intrinsic, unavoidable limitation within all power performance research.

While this study attempted to generalise the findings by stripping away site-specific wind distributions and calculating a Generic Annual Energy Production (AEP) using an idealised Rayleigh distribution (as detailed in Section 3.5.3), the thermodynamic behaviour of the boundary layer remains location-dependent. Offshore environments with different shallow-water thermodynamics, varying coastal proximity, or older stall-regulated wind turbines will naturally exhibit different magnitudes of aerodynamic deviation. Therefore, while the general patterns of stability-driven performance are transferable, the discrete percentage deviations must be calibrated to local environmental and mechanical conditions.

5.2. Industry Applications and Yield Optimisation

The quantification of stability-driven power deviations holds direct, actionable value for the offshore wind sector. Beyond advancing fundamental boundary layer meteorology, the integration of thermodynamic stability into standard aerodynamic assessments offers solutions to refine project financing, improve grid reliability, and optimise the design of future measurement campaigns.

Refining Pre-Construction Energy Yield Assessments

The most direct application of this research lies in the improvement of pre-construction Annual Energy Production (AEP) estimates. Standard yield assessments frequently rely on manufacturer-provided power curves that inherently assume neutral atmospheric stratification. However, this study demonstrates that such an assumption can

lead to systematic yield miscalculations in environments dominated by complex thermodynamics. For instance, if a specific offshore site experiences a high annual frequency of stable stratification, relying solely on a neutral baseline power curve will result in an overestimation of the aerodynamic energy capture.

By acknowledging that stability alters the kinetic energy flux across the rotor, project developers can implement stability-weighted probability distributions. Rather than using a single generic Weibull or Rayleigh distribution, developers can partition the local wind resource by stability regime and apply the corresponding empirical power curve (Stable, Neutral, or Unstable). This granular approach reduces the uncertainty in P50 and P90 energy yield probabilities. Ultimately, lowering this aerodynamic uncertainty directly reduces the financial risk profile of the wind farm, which can secure more favourable financing terms and in turn lower the Levelized Cost of Energy (LCOE).

Identifying Seasonal Stability Trends for Grid Planning

As offshore wind farms grow into major contributors to national energy grids, the predictability of their power output becomes an important factor in maintaining grid stability. This research shows that the aerodynamic efficiency of a turbine is not solely a function of wind speed; the transition between air masses, such as a cold continental flow advecting over warmer coastal waters (unstable) or warm air moving over a cooler sea (stable), produces distinct shifts in power generation. At the offshore site studied here, these transitions are not random but follow a seasonal cycle, as set out in Section 4.4.1.

The value of this finding lies less in short-term operational forecasting than in characterising these seasonal trends. Intra-day or ten-minute forecasting from real-time remote sensing struggles to outperform simple persistence models (Giebel et al. 2011), and the present analysis does not attempt to show day-ahead predictive skill. What it does establish is a repeatable seasonal pattern: the site is biased towards stable conditions in winter and spring and towards unstable conditions in summer and autumn, driven by the shallow-water temperature gradient between the air and the sea surface. This pattern has a systematic and quantifiable effect on yield, with the seasonal performance gap between stability regimes shown in Figure 4.15. Recognising that a site carries a recurring seasonal stability signature allows operators and planners to set more realistic seasonal yield expectations and to anticipate periods of over- or under-performance relative to a neutral baseline, which supports longer-horizon planning rather than real-time balancing.

Optimising Future Offshore Measurement Campaigns

Finally, the methodological observations recorded during this study provide valuable feedback for the design of future remote sensing campaigns. The identification of Lidar beam decorrelation under highly unstable, convective conditions points to an important boundary condition for the use of multi-beam remote sensing systems.

Because unstable atmospheres are characterised by intense, localised thermal mixing, the core Lidar assumption of horizontal flow homogeneity breaks down at extended measurement ranges where the spatial separation between beams is large. To maintain high-fidelity spatial data, this research indicates that industry practitioners should exercise caution when evaluating long-range remote sensing data (e.g., beyond several rotor diameters) in highly convective environments. In such regimes, data analysis should be restricted to closer distance gates (such as $2.5D$), or campaigns should consider supplementing standard pulsed systems with continuous-wave or high-resolution 3D scanning Lidars capable of resolving complex, heterogeneous turbulence structures.

Conclusions & Recommendations

6.1. Conclusions

The primary objective of this thesis was to isolate and quantify the specific impacts of atmospheric stability on offshore wind turbine power performance, independent of interconnected flow characteristics. By decoupling the underlying thermodynamic stability of the atmospheric boundary layer from localised kinematic variables, this research aimed to provide a clearer representation of how the atmosphere dictates energy yield. To address this, the main research question guiding the study was: *How does atmospheric stability influence the aerodynamic power performance of a modern offshore wind turbine, and to what extent can these thermodynamic effects be decoupled from interconnected kinematic characteristics?*

To arrive at the final answer to this primary question, the research was structured around three foundational sub-questions, the conclusions of which are detailed sequentially below.

6.1.1. What are the critical data processing requirements to develop a framework that extracts atmospheric and operational variables from nacelle-mounted Lidar, meteorological instruments and SCADA datasets?

Answering this question required the engineering of a systematic, multi-stage computational pipeline capable of handling disparate, high-frequency offshore data streams. The framework was divided into three distinct operational phases:

Data Foundation and Ingestion

The initial phase required the deployment of algorithms to extract the necessary raw signals from the nacelle-mounted Lidar, the transition-piece thermodynamic sensors, and the turbine's Supervisory Control and Data Acquisition (SCADA) system. For temporal consistency and to align with wind energy industry standards, high-frequency raw measurements were mathematically aggregated and matched into synchronised 10-minute intervals. Using a fully modular logic, these scripts were executed in series to merge the disparate signals into a unified dataset.

Data Processing and Stability Framework

Once ingested, the dataset was subjected to a careful series of Quality Control and filtering protocols. This began with temporal alignment and cross-validation between the Lidar and SCADA anemometry to confirm instrument synchronisation. The data were then refined through three distinct filtering layers: operational filtering (to remove turbine curtailment and downtime), instrumental filtering (enforcing Carrier-to-Noise Ratio (CNR) bounds and data availability limits on the Lidar), and environmental filtering (removing directional wakes from neighbouring turbines). The thoroughly filtered dataset was then passed into a modular stability framework, which categorised the discrete 10-minute periods into atmospheric regimes based on the specific classification methodologies applied.

Outputs and Analysis

In the final phase, the categorised data were used to conduct advanced aerodynamic and performance analyses. This included the construction of empirical power curves in accordance with IEC 61400-12-1 standards, applying sequential data normalisation techniques for varying air density and Turbulence Intensity (TI). This output process allowed the comparison of power and energy yield deviations across regimes and the study of complex upstream aerodynamic phenomena, such as seasonal variations and localised blockage effects.

6.1.2. How can offshore atmospheric stability regimes be accurately classified utilising a combination of kinematic proxies and bulk thermodynamic metrics?

To address this question, the study evaluated four distinct classification methodologies: two relying on kinematic proxies (vertical wind shear and Turbulence Intensity) and two based on thermodynamic metrics (the Bulk Richardson Number, Ri_b , and the Obukhov Length, L). The objective was to determine how well the simpler, widely available kinematic proxies performed in the absence of high-fidelity thermodynamic data.

The application of the Bulk Richardson Number (Ri_b), using virtual potential temperature measurements, proved to be the most consistent classification method. This metric provided a clear separation in energy deviations between the stability classes and exhibited relatively low variance, successfully stratifying the thermodynamic state of the offshore site.

Interestingly, the vertical wind shear exponent, despite being a much simpler empirical proxy, proved effective. Although its simplicity inherently caused some misclassification, it identified the overall thermodynamic patterns with relatively high accuracy. By contrast, while classification via Turbulence Intensity (TI) was generally capable of identifying the broad trends between stable and unstable regimes, the volatility of turbulent fluctuations introduced considerable noise and a high spread of output data, making it the least reliable stand-alone classifier for performance assessments.

6.1.3. To what extent do the overall energy yield and power performance deviate from theoretical baselines across stable, neutral, and unstable atmospheric conditions?

Based on the Ri_b thermodynamic classification, the turbine's power performance exhibited distinct, quantifiable deviations from the standard baseline, as illustrated in Table 4.5. As detailed in Chapter 4, the Rotor Equivalent Wind Speed (REWS) spatial shear correction was discarded due to Lidar geometric constraints. The isolation of these deviations therefore relied on air density normalisation and IEC Annex M Turbulence Intensity (TI) corrections.

Applying these normalisations yielded high-confidence results. At an average reference wind speed of 8 m/s, the aerodynamic efficiency of the turbine diverged symmetrically based on stability: stable atmospheric conditions exhibited a sustained underperformance of -1.81% , while unstable, convective conditions showed an overperformance of $+1.80\%$.

When calculating the generic Annual Energy Production (AEP) across the entire aggregated dataset regardless of stability, the total deviation from the OEM warranty curve was a negligible $+0.01\%$. This near-perfect alignment, falling well within the margin of statistical error, is an encouraging indicator for future industry collaboration, as it validates the accuracy of the baseline warranty curve when integrated over a full climatological cycle. However, when isolating the neutral conditions, the data showed an unexpected underperformance bias of -0.9% . This finding challenges the standard OEM assumption that idealised neutral boundary layers inherently provide optimal baseline energy capture, suggesting that neutral atmospheres may carry subtle aerodynamic deficits compared to true long-term averages.

6.2. Main Conclusion

With the underlying methodologies and discrete aerodynamic deviations established, the primary research question can be answered.

The aerodynamic power performance of modern offshore wind turbines is dynamically modulated by the thermodynamic stability of the Atmospheric Boundary Layer. While early academic literature presented contradictory

theories regarding whether stable or unstable air yields higher production, this research establishes that performance is intrinsically tied to a combination of localised climatology and turbine scale. Modern offshore wind turbines feature very large rotor swept areas that operate across a wide vertical domain, exposing them to highly stratified shear and density gradients.

Addressing the extent to which these thermodynamic effects can be decoupled from interconnected kinematic boundary layer variables points to a substantial analytical challenge. There is an inherently strong, physical correlation between kinematic characteristics (such as wind shear and turbulence) and atmospheric stability regimes. In fact, classical boundary layer meteorology often defines these kinematic behaviours as intrinsic components of the physical characterisation of atmospheric stability itself. Because these variables are practically interconnected in the natural environment, as illustrated in Figure 4.6, decoupling them is an advanced aerodynamic concept that requires a careful methodological approach. This research addressed this entanglement by using discrete environmental filtering and applying Turbulence Intensity (TI) normalisation to minimise high-frequency kinematic confounding effects and better isolate the underlying thermodynamic drivers.

When these thermodynamic effects are decoupled from the raw turbulent kinetic energy, a consensus emerges: within the partial-load operational region, unstable atmospheric conditions consistently outperform stable conditions, exhibiting significant statistical separation. The convective mixing inherent in unstable air creates a blunt, uniform kinetic energy flux that maximises energy extraction across the rotor disk, while the highly sheared and directionally misaligned flow of stable layers starves the lower rotor and reduces overall efficiency.

This performance pattern, where unstable conditions yield the highest aerodynamic efficiency, is expected to hold across the modern offshore wind industry. The sole empirical exception to this consensus would occur at sites possessing a highly concentrated wind speed distribution tightly centred around the turbine's transition region. In such rare, poorly assessed, scenarios, the pronounced shear of stable or neutral conditions might push the upper rotor into rated power sooner, potentially causing those regimes to outperform unstable air. However, for standard offshore climates, accounting for the distinct power surplus of unstable regimes and the symmetrical deficit of stable regimes is important for producing accurate, risk-averse wind resource assessments.

6.3. Recommendations for Future Research

While this research successfully quantified the aerodynamic impacts of atmospheric stability and established a decoupling methodology, it also points to several avenues for further scientific inquiry. The following recommendations focus on overcoming the identified instrumental limitations, advancing the classification algorithms, and scaling these findings to improve broader industry practices in wind resource assessment.

6.3.1. Enhanced Instrumentation and REWS Implementation

The primary methodological constraint of this study was the inability to reliably execute the Rotor Equivalent Wind Speed (REWS) correction due to the geometric limitations of the forward-looking nacelle Lidar. Future measurement campaigns should prioritise the deployment of vertically profiling Lidars or advanced 3D scanning Lidars capable of simultaneously mapping the wind field at multiple discrete heights across the entire rotor disk. Fulfilling this geometric requirement would allow researchers to integrate REWS alongside air density and Turbulence Intensity (TI) normalisations, providing a complete decoupling of spatial shear from thermodynamic stability. Complementing this remote sensing setup with high-frequency 3D sonic anemometers would enable the calculation of turbulent kinematic heat fluxes via eddy covariance, allowing the direct application of Monin-Obukhov Similarity Theory (MOST) to cross-validate the empirical Bulk Richardson Number (Ri_b) methodology.

6.3.2. Advanced Multivariate Classification via Machine Learning

This study relied on discrete, rigid numerical thresholds to separate the dataset into stable, neutral, and unstable regimes. However, the atmospheric boundary layer exists on a continuous thermodynamic spectrum. Future research should examine advanced unsupervised machine learning techniques, such as Gaussian Mixture Models (GMMs) or K-means clustering, to classify stability regimes. Rather than relying on a single rigid threshold, GMMs evaluate a multivariate feature space, simultaneously clustering data points based on combinations of Ri_b , TI, wind shear (α), and temperature gradients. Because GMMs assign probabilistic boundaries rather than hard limits, they can capture the nuanced transitions between regimes (e.g., "weakly stable" or "near-neutral"), potentially eliminating the classification noise inherent in strict threshold methodologies.

6.3.3. Comprehensive Uncertainty Quantification

While the calculated power deviations (-1.81% to +1.80%) show a high level of statistical confidence, a formal sensitivity and uncertainty analysis is recommended to further validate these metrics. Future studies should use uncertainty propagation frameworks, such as Monte Carlo simulations or the strict analytical methods outlined in IEC 61400-12-1. This analysis should quantify how individual measurement uncertainties, such as Lidar volume averaging artefacts, standard errors in the transition-piece temperature sensors, and anemometry calibration drift, propagate through the Ri_b calculation and ultimately affect the final Annual Energy Production (AEP) integration. Establishing strict uncertainty bands around the stability-driven deviations will be important for convincing financial stakeholders to adopt these corrections.

6.3.4. Pre-Construction Application via Floating Remote Sensing

A particularly valuable progression of this research would be translating these post-construction operational findings into pre-construction yield prediction algorithms. The industry currently struggles to accurately predict stability-driven deviations before a wind farm is built. Future research should investigate whether the Ri_b -based methodology developed in this thesis can be reliably replicated using stand-alone Floating Lidar Systems (FL-iDAR) equipped with integrated met-ocean buoys (measuring sea surface and air temperatures). If a floating platform can gather the required thermodynamic data, project developers could apply stability-weighted probability distributions to their preliminary wind resource data, resulting in optimised, risk-averse pre-construction Energy Yield Assessments (EYA).

6.3.5. Geographic and Technological Transferability

Finally, the findings of this thesis reflect the specific climatology of a shallow-water North Sea site and the aerodynamic characteristics of a specific modern multi-megawatt turbine. To build industry confidence in these performance deviations, geographic and technological diversification of this framework is needed. Future studies should replicate this exact modular methodology across different climates, such as deep-water floating offshore sites (where sea surface temperature inertia behaves differently) or coastal environments dominated by severe diurnal land-sea breezes. Testing this framework on next-generation 15+ MW turbines, which feature considerably larger rotor swept areas and different active pitch control logic, will determine whether the magnitude of these stability-driven aerodynamic deviations scales linearly with rotor diameter.

Acknowledgments


For the past seven months, I have had the opportunity to conduct this research thesis in collaboration with RWE and TU Delft. Investigating offshore wind turbine performance has been a challenging yet incredibly rewarding journey, and I truly hope that you enjoy reading this work. More importantly, I hope it encourages future engineers to take on the harder problems in offshore wind, where the atmosphere rarely behaves the way our models assume, and to help carry the energy transition forward.

I would first like to express my deep gratitude to my university supervisor and chairman of the thesis committee, Prof. Dr. Simon Watson, for his continuous support, constructive feedback, and expert insight throughout this project. I am also very thankful to Prof. Dr. Sara Porchetta for agreeing to be a member of my thesis committee.

At RWE, I owe my sincere appreciation to my company supervisor, Justin Burstein. Despite his busy schedule, he always went out of his way to help me, sometimes quite literally, climbing the turbine himself to retrieve the data I needed. I would also like to thank the Wind Resource Assessment team, whose help was essential to seeing this project through. My appreciation also goes to Fraunhofer IWES for kindly allowing me to use their Lidar data, and to Jörg Schneemann and Janna Seifert from Oldenburg University for their crucial help with the SST and installation processes.

Finally, to my family, thank you for giving me the opportunity to pursue an MSc degree in the Netherlands. I am also grateful to my friends, both here in Delft and back home in Greece, for their support, and to my partner, Aga, for being there for me.

*Dimitrios Gkoufas
June 2026
Delft, Netherlands*



Data Validation

To verify the reliability of the remote sensing data and the effectiveness of the filtering protocols, the upstream LiDAR measurements were also cross-checked against the turbine's standard SCADA anemometry. Figure A.1 illustrates this comparison at the 407 m reference distance for both wind direction and hub-height wind speed.

The raw dataset exhibits significant measurement scatter, which is largely driven by wake interactions from neighbouring turbines and localized flow distortions. Applying the directional and operational filters successfully removes this aerodynamic noise. As shown by the linear regression fits, the retained dataset demonstrates strong agreement between the two measurement systems, with coefficients of determination (R^2) exceeding 0.97 for both metrics. This tight correlation confirms that the filtering process effectively isolates the undisturbed free-stream flow, establishing a reliable baseline for the subsequent power performance assessment.




Figure A.1: Validation and cross-check analysis between nacelle anemometer and LiDAR measurements at the 407 m distance gate. Background data points represent the raw 10-minute records, while the highlighted subset shows the remaining data after applying the filtering criteria.




Figure A.2: Wind speed validation between SCADA and LiDAR measurements, segmented by atmospheric stability regime. Atmospheric stability was classified using the Bulk Richardson Number (BRN). Each panel displays the filtered operational dataset alongside its corresponding linear regression fit and coefficient of determination (R^2).

In Figure A.2, the linear regression slopes between the two instruments vary by stability regime. Unstable conditions show near-perfect agreement (slope of 0.99), while the correlation drops slightly under stable (0.97) and neutral (0.93) conditions.

This variation likely points to aerodynamic limitations in the nacelle anemometry rather than a bias in the forward-looking LiDAR. Because nacelle anemometers sit directly behind the rotor, they measure wind that has already been distorted by the blades and the nacelle structure. The turbine applies a Nacelle Transfer Function (NTF) to estimate the free-stream wind speed from these measurements. However, varying atmospheric stability fundamentally alters the incoming wind shear and turbulence, which in turn changes the shape and intensity of the rotor wake. NTF algorithms often struggle to fully correct for these complex, stability-driven flow distortions.

B

REWS Correction

While the theoretical framework for the Rotor Equivalent Wind Speed (REWS) assumes a dense vertical measurement profile across the rotor swept area, the practical implementation for this assessment was constrained by the availability of only three discrete vertical measurement points. Consequently, the standard continuous integral was approximated by discretizing the rotor swept area into $n = 3$ segments.

The available dataset provided two physical Lidar measurements from beams offset by 10° relative to each other, equating to a vertical offset of approximately 5° above and below the hub height. As illustrated in Figure B.1, at a standard measurement distance of $2.5D$, these beam angles correspond to specific measurement heights within the rotor swept area.




Figure B.1: Geometric visualization of the nacelle-mounted Lidar beams. At a horizontal measurement distance of $L = 2.5D$, the $\pm 5^\circ$ beam elevation angles result in a vertical offset (Δz) from the hub height (H). Both physical measurement heights (z_{low} and z_{high}) fall within the turbine's rotor swept area.

The dynamic measurement height of the lower beam is denoted as z_{low} . For the central segment, the standard hub height wind speed (V_{hub}) was utilized. It is acknowledged as a technical limitation that this V_{hub} signal is reconstructed rather than directly measured by the Lidar; however, its inclusion was necessary to satisfy the three-measurement requirement of the methodology to construct the velocity profile.

Mathematical Formulation

The area of a circular segment integrated up to a specific height z is calculated using the geometric function $g(z)$:

$$g(z) = \Delta y \sqrt{R^2 - \Delta y^2} + R^2 \arctan \left(\frac{\Delta y}{\sqrt{R^2 - \Delta y^2}} \right) \quad (\text{B.1})$$

where:

$\Delta y = z - H$, constrained to the rotor boundaries $H - R \leq z \leq H + R$.

H is the turbine hub height (107.5 m).

R is the rotor radius (83.5 m).

Based on the three available signals, the rotor swept area was partitioned into three segments defined by the following height boundaries:

$$z_0 = H - R \quad (\text{Bottom of rotor}) \quad (\text{B.2})$$

$$z_1 = \max(z_0, \min(z_{\text{low}}, z_2)) \quad (\text{B.3})$$

$$z_2 = z_0 + 2 \left(\frac{2R}{3} \right) \quad (\text{B.4})$$

$$z_3 = H + R \quad (\text{Top of rotor}) \quad (\text{B.5})$$

The individual segment areas A_i are defined as the difference between the integrated areas at their respective upper and lower boundaries:

$$A_1 = g(z_1) - g(z_0) \quad (\text{B.6})$$

$$A_2 = g(z_2) - g(z_1) \quad (\text{B.7})$$

$$A_3 = g(z_3) - g(z_2) \quad (\text{B.8})$$

Finally, the implemented Rotor Equivalent Wind Speed (V_{eq}) is calculated by area-weighting the cubed wind speeds of the lower (V_{low}), hub (V_{hub}), and upper (V_{high}) signals:

$$V_{\text{eq}} = \left(\frac{A_1 V_{\text{low}}^3 + A_2 V_{\text{hub}}^3 + A_3 V_{\text{high}}^3}{A} \right)^{1/3} \quad (\text{B.9})$$

where $A = \pi R^2$ is the total rotor swept area.





Figure B.2: Impact analysis of the implemented Wind Shear (REWS) correction. (Left) Scatter comparison between the standard Hub Height Wind Speed and the Rotor Equivalent Wind Speed (both normalized to the 13 m/s rated wind speed). (Right) Histogram detailing the distribution of the absolute wind shear correction impact ($\Delta V = V_{eq} - V_{hub}$) in m/s, demonstrating the mean shift applied to the dataset's velocity profile.



Turbulence Intensity Normalization

This annex presents the mathematical framework for the Turbulence Intensity (TI) normalization procedure implemented in this study. This correction isolates stability-driven aerodynamic effects by neutralizing the artificial power enhancements or deficits caused by high-frequency wind speed fluctuations.

The Initial Zero-Turbulence Power Curve

The normalization routine requires the construction of an initial reference power curve representing hypothetical conditions under zero ambient turbulence ($I = 0$). This curve is derived iteratively by adjusting the parameters of the measured, bin-averaged operational profile normalized for air density.

The first-approach approximation for the zero-turbulence power curve parameters follows a strict sequence:

- a) Cut-in Wind Speed ($v_{\text{cut-in}}$): Set to the mean wind speed of the lowest velocity bin where the measured power output matches or exceeds 0.1% of the turbine's rated power (P_{rated}):

$$P_i \geq 0.001 \cdot P_{\text{rated}} \quad (\text{C.1})$$

- b) Partial Load Region (Region II): A constant power coefficient (C_p) equal to the maximum measured bin-averaged power coefficient ($C_{p,\text{max}}$) is assumed. Within this region, the power output scales strictly with the cube of the wind speed:

$$P(v) \propto v^3 \quad \text{with} \quad C_p = C_{p,\text{max}} \quad \text{for} \quad v_{\text{cut-in}} \leq v \leq v_{\text{rated}} \quad (\text{C.2})$$

The initial zero-turbulence power curve is calculated with this assumption using standard reference air density in increments no larger than 0.1 m/s.

- c) Rated Wind Speed (v_{rated}): Determined directly from the rated electrical output, rotor swept area (A), maximum power coefficient ($C_{p,\text{max}}$), and reference air density (ρ):

$$v_{\text{rated}} = \left(\frac{2 \cdot P_{\text{rated}}}{\rho \cdot C_{p,\text{max}} \cdot A} \right)^{1/3} \quad (\text{C.3})$$

where P_{rated} is set to the maximum bin-averaged power output observed across the operational dataset.

- d) Rated Power Region (Region IV): For all wind velocities exceeding v_{rated} , the power output is fixed flat at P_{rated} , up to an assumed structural threshold ($v_{\text{cut-out}} = 100$ m/s).




Figure C.1: First-approach approximation for the initial zero-turbulence power curve parameters IEC (2017).

Iterative Mathematical Fitting and Convergence

Because parameters like P_{rated} , $v_{\text{cut-in}}$, and $C_{p,\text{max}}$ are naturally dependent on the background turbulence field, the initial parameters must be mathematically adjusted through an iterative loop.

At each iteration step, the zero-turbulence power curve ($P_{I=0}$) is numerically integrated over a Gaussian wind speed distribution, $f(v)$, where the mean equals the bin wind speed and the standard deviation is the product of the binned wind speed and the binned turbulence intensity:

$$\sigma_v = v_i \cdot I_i \quad (\text{C.4})$$

The parameters are modified sequentially until the simulated power curve converges toward the measured data according to three thresholds:

- The maximum power of the simulated curve deviates from the measured curve by $\leq 0.1\%$.
- The simulated cut-in wind speed deviates from the measured curve by ≤ 0.5 m/s.
- The maximum simulated power coefficient deviates from the measured curve by ≤ 0.01 .

The Power Simulation and Normalization Function

Once the finalized zero-turbulence power curve ($P_{I=0}$) is mathematically established, the 10-minute average power output of an active pitch-controlled wind turbine can be simulated under any turbulence state using the fundamental integration function:

$$\overline{P}_{\text{sim}}(v) = \int_{v=0}^{\infty} P_{I=0}(v) \cdot f(v) dv \quad (\text{C.5})$$

where $f(v)$ is assumed to be a localized Gaussian distribution representing the high-frequency wind fluctuations occurring within that specific 10-minute period.

Equation C.5 is applied to every valid 10-minute data point in the campaign to compute two distinct, concurrent simulated power metrics:

1. $\overline{P}_{\text{sim},I}(v)$: Evaluated using the air-density normalized hub-height wind speed and the actual measured 10-minute Turbulence Intensity (I_{measured}).

2. $\overline{P_{\text{sim}, I_{\text{ref}}}}(v)$: Evaluated using the identical air-density normalized wind speed and the dataset's fixed reference Turbulence Intensity (I_{ref}).

Finally, the normalized power output ($\overline{P_{I_{\text{ref}}}}$), adjusted to match the consistent baseline turbulence profile, is calculated for each individual observation using the following difference equation:

$$\overline{P_{I_{\text{ref}}}}(v) = \overline{P}(v) - \overline{P_{\text{sim}, I}}(v) + \overline{P_{\text{sim}, I_{\text{ref}}}}(v) \quad (\text{C.6})$$

where $\overline{P}(v)$ represents the raw measured power output. These corrected power values are subsequently utilized in the core body of this thesis to construct the final stability-classified power curves and compute relative energy deviations.

References

- Anantharaman, Arjun et al. (Sept. 8, 2025). “The impact of far-reaching offshore cluster wakes on wind turbine fatigue loads”. In: *Wind Energy Science* 10.9, pp. 1849–1867. ISSN: 2366-7443. DOI: 10.5194/wes-10-1849-2025. URL: <https://wes.copernicus.org/articles/10/1849/2025/> (visited on 12/04/2025).
- Araújo Da Silva, Marcos et al. (Mar. 14, 2022). “Assessing Obukhov Length and Friction Velocity from Floating Lidar Observations: A Data Screening and Sensitivity Computation Approach”. In: *Remote Sensing* 14.6, p. 1394. ISSN: 2072-4292. DOI: 10.3390/rs14061394. URL: <https://www.mdpi.com/2072-4292/14/6/1394> (visited on 12/04/2025).
- Arya, S. Pal (1988). *Introduction to Micrometeorology*. Google-Books-ID: bLkhFx5sPEIC. Academic Press. 344 pp. ISBN: 978-0-12-064490-2.
- Bahamonde, Manuel Ignacio and Salvador P. Litrán (May 1, 2019). “Study of the energy production of a wind turbine in the open sea considering the continuous variations of the atmospheric stability and the sea surface roughness”. In: *Renewable Energy* 135, pp. 163–175. ISSN: 0960-1481. DOI: 10.1016/j.renene.2018.11.075. URL: <https://www.sciencedirect.com/science/article/pii/S0960148118313946> (visited on 12/22/2025).
- Bardal, Lars Morten et al. (June 20, 2018). “Evaluation of methods for estimating atmospheric stability at two coastal sites”. In: *Wind Engineering* 42, p. 0309524X1878037. DOI: 10.1177/0309524X18780378.
- Barthelmie, R J et al. (Mar. 2015). “The role of atmospheric stability/turbulence on wakes at the Egmond aan Zee offshore wind farm”. In: *Journal of Physics: Conference Series* 625.1, p. 012002. ISSN: 1742-6596. DOI: 10.1088/1742-6596/625/1/012002. URL: <https://doi.org/10.1088/1742-6596/625/1/012002> (visited on 05/28/2026).
- Berg, G. P. van den (Mar. 2008). “Wind turbine power and sound in relation to atmospheric stability”. In: *Wind Energy* 11.2, pp. 151–169. ISSN: 1095-4244, 1099-1824. DOI: 10.1002/we.240. URL: <https://onlinelibrary.wiley.com/doi/10.1002/we.240> (visited on 12/04/2025).
- Betz, Albert (1920). *Das Maximum der theoretisch möglichen Ausnützung des Windes durch Windmotoren. Zeitschrift für das gesamte Turbinenwesen*. URL: <https://www.scrip.org/reference/referencespapers?referenceid=1642440> (visited on 05/28/2026).
- Borraccino, A and R Wagner (n.d.). “Power curve measurement using ∞ estimates from nacelle lidars”. In: ().
- Borraccino, Antoine et al. (Mar. 6, 2017). “Wind Field Reconstruction from Nacelle-Mounted Lidars ShortRange Measurements”. In: DOI: 10.5194/wes-2017-10. URL: <https://wes.copernicus.org/articles/2/269/2017/wes-2-269-2017-discussion.html> (visited on 12/19/2025).
- Businger, J. A. et al. (Mar. 1, 1971). “Flux-Profile Relationships in the Atmospheric Surface Layer”. In: *Journal of the Atmospheric Sciences* 28.2, pp. 181–189. ISSN: 0022-4928, 1520-0469. DOI: 10.1175/1520-0469(1971)028<0181:FPRITA>2.0.CO;2. URL: https://journals.ametsoc.org/view/journals/atsc/28/2/1520-0469_1971_028_0181_fprita_2_0_co_2.xml (visited on 01/28/2026).
- Cao, Jiufa et al. (Dec. 15, 2025). “Impact of atmospheric stability on wake interactions and power performance of offshore wind farm clusters”. In: *Energy Conversion and Management* 346, p. 120421. ISSN: 0196-8904. DOI: 10.1016/j.enconman.2025.120421. URL: <https://www.sciencedirect.com/science/article/pii/S0196890425009458> (visited on 12/19/2025).
- Delfos (2025). *How Much Does Wind Power Cost in Europe? ROI, Financing & Payback*. URL: <https://www.delfos.energy/blog-posts/cost-of-wind-energy-in-europe-investment-and-financial-returns> (visited on 05/22/2026).
- Dörenkämper, Martin et al. (Dec. 16, 2014). “Atmospheric Impacts on Power Curves of Multi-Megawatt Offshore Wind Turbines”. In: *J. Phys.: Conf. Ser.* 555, p. 0122029. DOI: 10.1088/1742-6596/555/1/012209.
- Dyer, A. J. (Nov. 1974). “A review of flux-profile relationships”. In: *Boundary-Layer Meteorology* 7.3, pp. 363–372. ISSN: 0006-8314, 1573-1472. DOI: 10.1007/BF00240838. URL: <http://link.springer.com/10.1007/BF00240838> (visited on 01/28/2026).
- Giebel, Gregor et al. (Jan. 28, 2011). *The State of the Art in Short-Term Prediction of Wind Power A Literature Overview, 2nd Edition*. DOI: 10.13140/RG.2.1.2581.4485.

- Held, Dominique P. and Jakob Mann (July 11, 2019). “Detection of wakes in the inflow of turbines using nacelle lidars”. In: *Wind Energy Science* 4.3, pp. 407–420. ISSN: 2366-7443. DOI: 10.5194/wes-4-407-2019. URL: <https://wes.copernicus.org/articles/4/407/2019/> (visited on 12/27/2025).
- Hellmann, Gustav (1916). *Über die Bewegung der Luft in den untersten Schichten der Atmosphäre*. *Meteorologische Zeitschrift*. 416–421.
- IEA (2026). *Prices – Electricity 2026 – Analysis*. IEA. URL: <https://www.iea.org/reports/electricity-2026/prices> (visited on 06/02/2026).
- IEC (2017). *Wind energy generation systems: Part 12-1: Power performance measurements of electricity producing wind turbines*. BSI British Standards. DOI: 10.3403/30434175. URL: <https://linkresolver.bsigroup.com/junction/resolve/000000000030434175?restype=standard> (visited on 01/12/2026).
- (2022). *Wind energy generation systems. Part 50-3 Use of nacelle-mounted lidars for wind measurements*. OCLC: 1324578676. London: British Standards Institution.
- Isaaks, Edward H. and R. Mohan Srivastava (1989). *Applied Geostatistics*. Google-Books-ID: vC2dcXFLI3YC. Oxford University Press. 592 pp. ISBN: 978-0-19-505013-4.
- Lee, Joseph C Y et al. (n.d.). “The Power Curve Working Group’s Assessment of Wind Turbine Power Performance Prediction Methods”. In: ().
- Letizia, Stefano et al. (Nov. 2, 2023). “Characterization of wind turbine flow through nacelle-mounted lidars: a review”. In: *Frontiers in Mechanical Engineering* 9. ISSN: 2297-3079. DOI: 10.3389/fmech.2023.1261017. URL: <https://www.frontiersin.org/journals/mechanical-engineering/articles/10.3389/fmech.2023.1261017/full> (visited on 05/28/2026).
- Martynowicz, Paweł, Piotr Ślimak, and Desta Kalbessa Kumsa (Jan. 2026). “Advancing Offshore Wind Capacity Through Turbine Size Scaling”. In: *Energies* 19.7, p. 1625. ISSN: 1996-1073. DOI: 10.3390/en19071625. URL: <https://www.mdpi.com/1996-1073/19/7/1625> (visited on 05/22/2026).
- Monin, A S and A M Obukhov (1954). “Basic laws of turbulent mixing in the surface layer of the atmosphere”. In: ().
- Müller, Sebastian (May 24, 2026). *GeoStat-Framework/PyKrige*. original-date: 2014-10-01T20:19:24Z. URL: <https://github.com/GeoStat-Framework/PyKrige> (visited on 05/25/2026).
- Murray, F. W. (Feb. 1, 1967). “On the Computation of Saturation Vapor Pressure”. In: *Journal of Applied Meteorology and Climatology* 6.1, pp. 203–204. ISSN: 1520-0450. DOI: 10.1175/1520-0450(1967)006<0203:OTCOSV>2.0.CO;2. URL: https://journals.ametsoc.org/view/journals/apme/6/1/1520-0450_1967_006_0203_otcosv_2_0_co_2.xml (visited on 05/08/2026).
- Newman, Jennifer F. and Andrew Clifton (Feb. 10, 2017). “An error reduction algorithm to improve lidar turbulence estimates for wind energy”. In: *Wind Energy Science* 2.1, pp. 77–95. ISSN: 2366-7451. DOI: 10.5194/wes-2-77-2017. URL: <https://wes.copernicus.org/articles/2/77/2017/> (visited on 12/04/2025).
- Odysee-MURE (2024). *Electricity consumption per dwelling | Electricity dwelling | ODYSSEE-MURE*. URL: <https://www.odyssee-mure.eu/publications/efficiency-by-sector/households/electricity-consumption-dwelling.html> (visited on 06/02/2026).
- Osamu, Kiyomiya, Tatsuomi Rikiji, and P.H.A.J.M. Gelder (Jan. 1, 2002). “Dynamic Response Analysis of On-shore Wind Energy Power Units during Earthquakes and Wind”. In: ().
- Peña, Alfredo, Ginka G Yankova, and Vasiliki Mallini (n.d.). “On the lidar-turbulence paradox and possible countermeasures”. In: ().
- Pérez, Christy et al. (Jan. 2023). “Influence of Atmospheric Stability on Wind Turbine Energy Production: A Case Study of the Coastal Region of Yucatan”. In: *Energies* 16.10, p. 4134. ISSN: 1996-1073. DOI: 10.3390/en16104134. URL: <https://www.mdpi.com/1996-1073/16/10/4134> (visited on 12/04/2025).
- Pérez Albornoz, C. et al. (July 1, 2022). “Review of atmospheric stability estimations for wind power applications”. In: *Renewable and Sustainable Energy Reviews* 163, p. 112505. ISSN: 1364-0321. DOI: 10.1016/j.rser.2022.112505. URL: <https://www.sciencedirect.com/science/article/pii/S1364032122004099> (visited on 02/09/2026).
- Remote Sensing for Wind Energy* (2011). Report. Publication Title: Remote Sensing for Wind Energy. Risø National Laboratory for Sustainable Energy, Technical University of Denmark.
- Ren, Huilong (Dec. 5, 2025). “Thermodynamically consistent damage models derived from cumulative distribution functions, with FEM implementations”. In: *Computational Mechanics*, pp. 1–29. DOI: 10.1007/s00466-025-02726-x.
- Robey, Rachel and Julie K. Lundquist (Aug. 15, 2022). “Behavior and mechanisms of Doppler wind lidar error in varying stability regimes”. In: *Atmospheric Measurement Techniques* 15.15, pp. 4585–4622. ISSN: 1867-1381.

- doi: 10.5194/amt-15-4585-2022. URL: <https://amt.copernicus.org/articles/15/4585/2022/> (visited on 12/04/2025).
- Ruiz-Leo, María, Isabel Cristina Gil García, and Ana Fernández-Guillamón (Nov. 18, 2025). “From Resource Assessment to AEP Correction: Methodological Framework for Comparing HAWT and VAWT Offshore Systems”. In: *Journal of Marine Science and Engineering* 13, p. 2183. doi: 10.3390/jmse13112183.
- Sathe, A. et al. (Feb. 12, 2015). “A six-beam method to measure turbulence statistics using ground-based wind lidars”. In: *Atmospheric Measurement Techniques* 8.2, pp. 729–740. issn: 1867-8548. doi: 10.5194/amt-8-729-2015. URL: <https://amt.copernicus.org/articles/8/729/2015/> (visited on 12/04/2025).
- Schneemann, Jörg et al. (Apr. 13, 2021). “Offshore wind farm global blockage measured with scanning lidar”. In: *Wind Energy Science* 6.2, pp. 521–538. issn: 2366-7443. doi: 10.5194/wes-6-521-2021. URL: <https://wes.copernicus.org/articles/6/521/2021/> (visited on 12/04/2025).
- Shimada, Susumu et al. (June 6, 2022). “Validation of near-shore wind measurements using a dual scanning light detection and ranging system”. In: *Wind Energy* 25. doi: 10.1002/we.2757.
- Souaiby, Marwa and Fernando Porté-Agel (Sept. 11, 2025). “Atmospheric stability effect on wind farm flow and performance”. In: *Physics of Fluids* 37.9, p. 095152. issn: 1070-6631. doi: 10.1063/5.0280027. URL: <https://doi.org/10.1063/5.0280027> (visited on 12/19/2025).
- St. Martin, Clara M. et al. (Nov. 1, 2016). “Wind turbine power production and annual energy production depend on atmospheric stability and turbulence”. In: *Wind Energy Science* 1.2, pp. 221–236. issn: 2366-7443. doi: 10.5194/wes-1-221-2016. URL: <https://wes.copernicus.org/articles/1/221/2016/wes-1-221-2016.html> (visited on 12/04/2025).
- Stull, Roland B. (1988). *An Introduction to Boundary Layer Meteorology*. (Visited on 05/28/2026).
- Tempel, J van der et al. (2010). “Design of support structures for offshore wind turbines”. In: *Wind Power Generation and Wind Turbine Design*. Ed. by null Wei Tong. Southampton: WIT Press, pp. 559–594. isbn: 978-1-84564-205-1. URL: <https://www.scopus.com/pages/publications/77957877667> (visited on 04/19/2026).
- UNFCCC (n.d.). *Offshore wind power plants with fixed foundations* | UNFCCC. URL: <https://unfccc.int/technology/offshore-wind-power-plants-with-fixed-foundations>.
- Vaisala (Aug. 29, 2020). *Vaisala | Wind Energy WindCube Nacelle*. URL: <https://www.vaisala.com/en/products/wind-energy-windcube-nacelle> (visited on 03/10/2026).
- (2023). *Hardware version 2.2 - Software version 3.1 WindCube Nacelle PPT Configuration User manual M212516EN-D*. Vaisala.
- Wagner, Rozenn and Samuel Davoust (2013). *Nacelle lidar for power curve measurement - Avedøre campaign*. Report. Publication Title: Nacelle lidar for power curve measurement - Avedøre campaign. DTU Wind Energy.
- Wallace, John M. and Peter V. Hobbs (Mar. 24, 2006). *Atmospheric Science: An Introductory Survey*. Google-Books-ID: HZ2wNtDOU0oC. Elsevier. 507 pp. isbn: 978-0-08-049953-6.
- Weitkamp, Claus, ed. (2005). *Lidar- Range-Resolved Optical Remote Sensing of the Atmosphere*. Vol. 102. Springer Series in Optical Sciences. New York: Springer-Verlag. isbn: 978-0-387-40075-4. doi: 10.1007/b106786. URL: <http://link.springer.com/10.1007/b106786> (visited on 05/28/2026).
- Zhang, Wei, Elliott Walker, and Corey D. Markfort (Jan. 2025). “Influence of Surface Complexity and Atmospheric Stability on Wind Shear and Turbulence in a Peri-Urban Wind Energy Site”. In: *Energies* 18.19, p. 5211. issn: 1996-1073. doi: 10.3390/en18195211. URL: <https://www.mdpi.com/1996-1073/18/19/5211> (visited on 12/04/2025).