A Learning-Based Algorithm for Modeling Blade Degradation in Wind Turbines

by Decoupling Lift and Drag Coefficient for a Physics-Informed Approach

Master Thesis Project

Mats Heine

Delft University of Technology

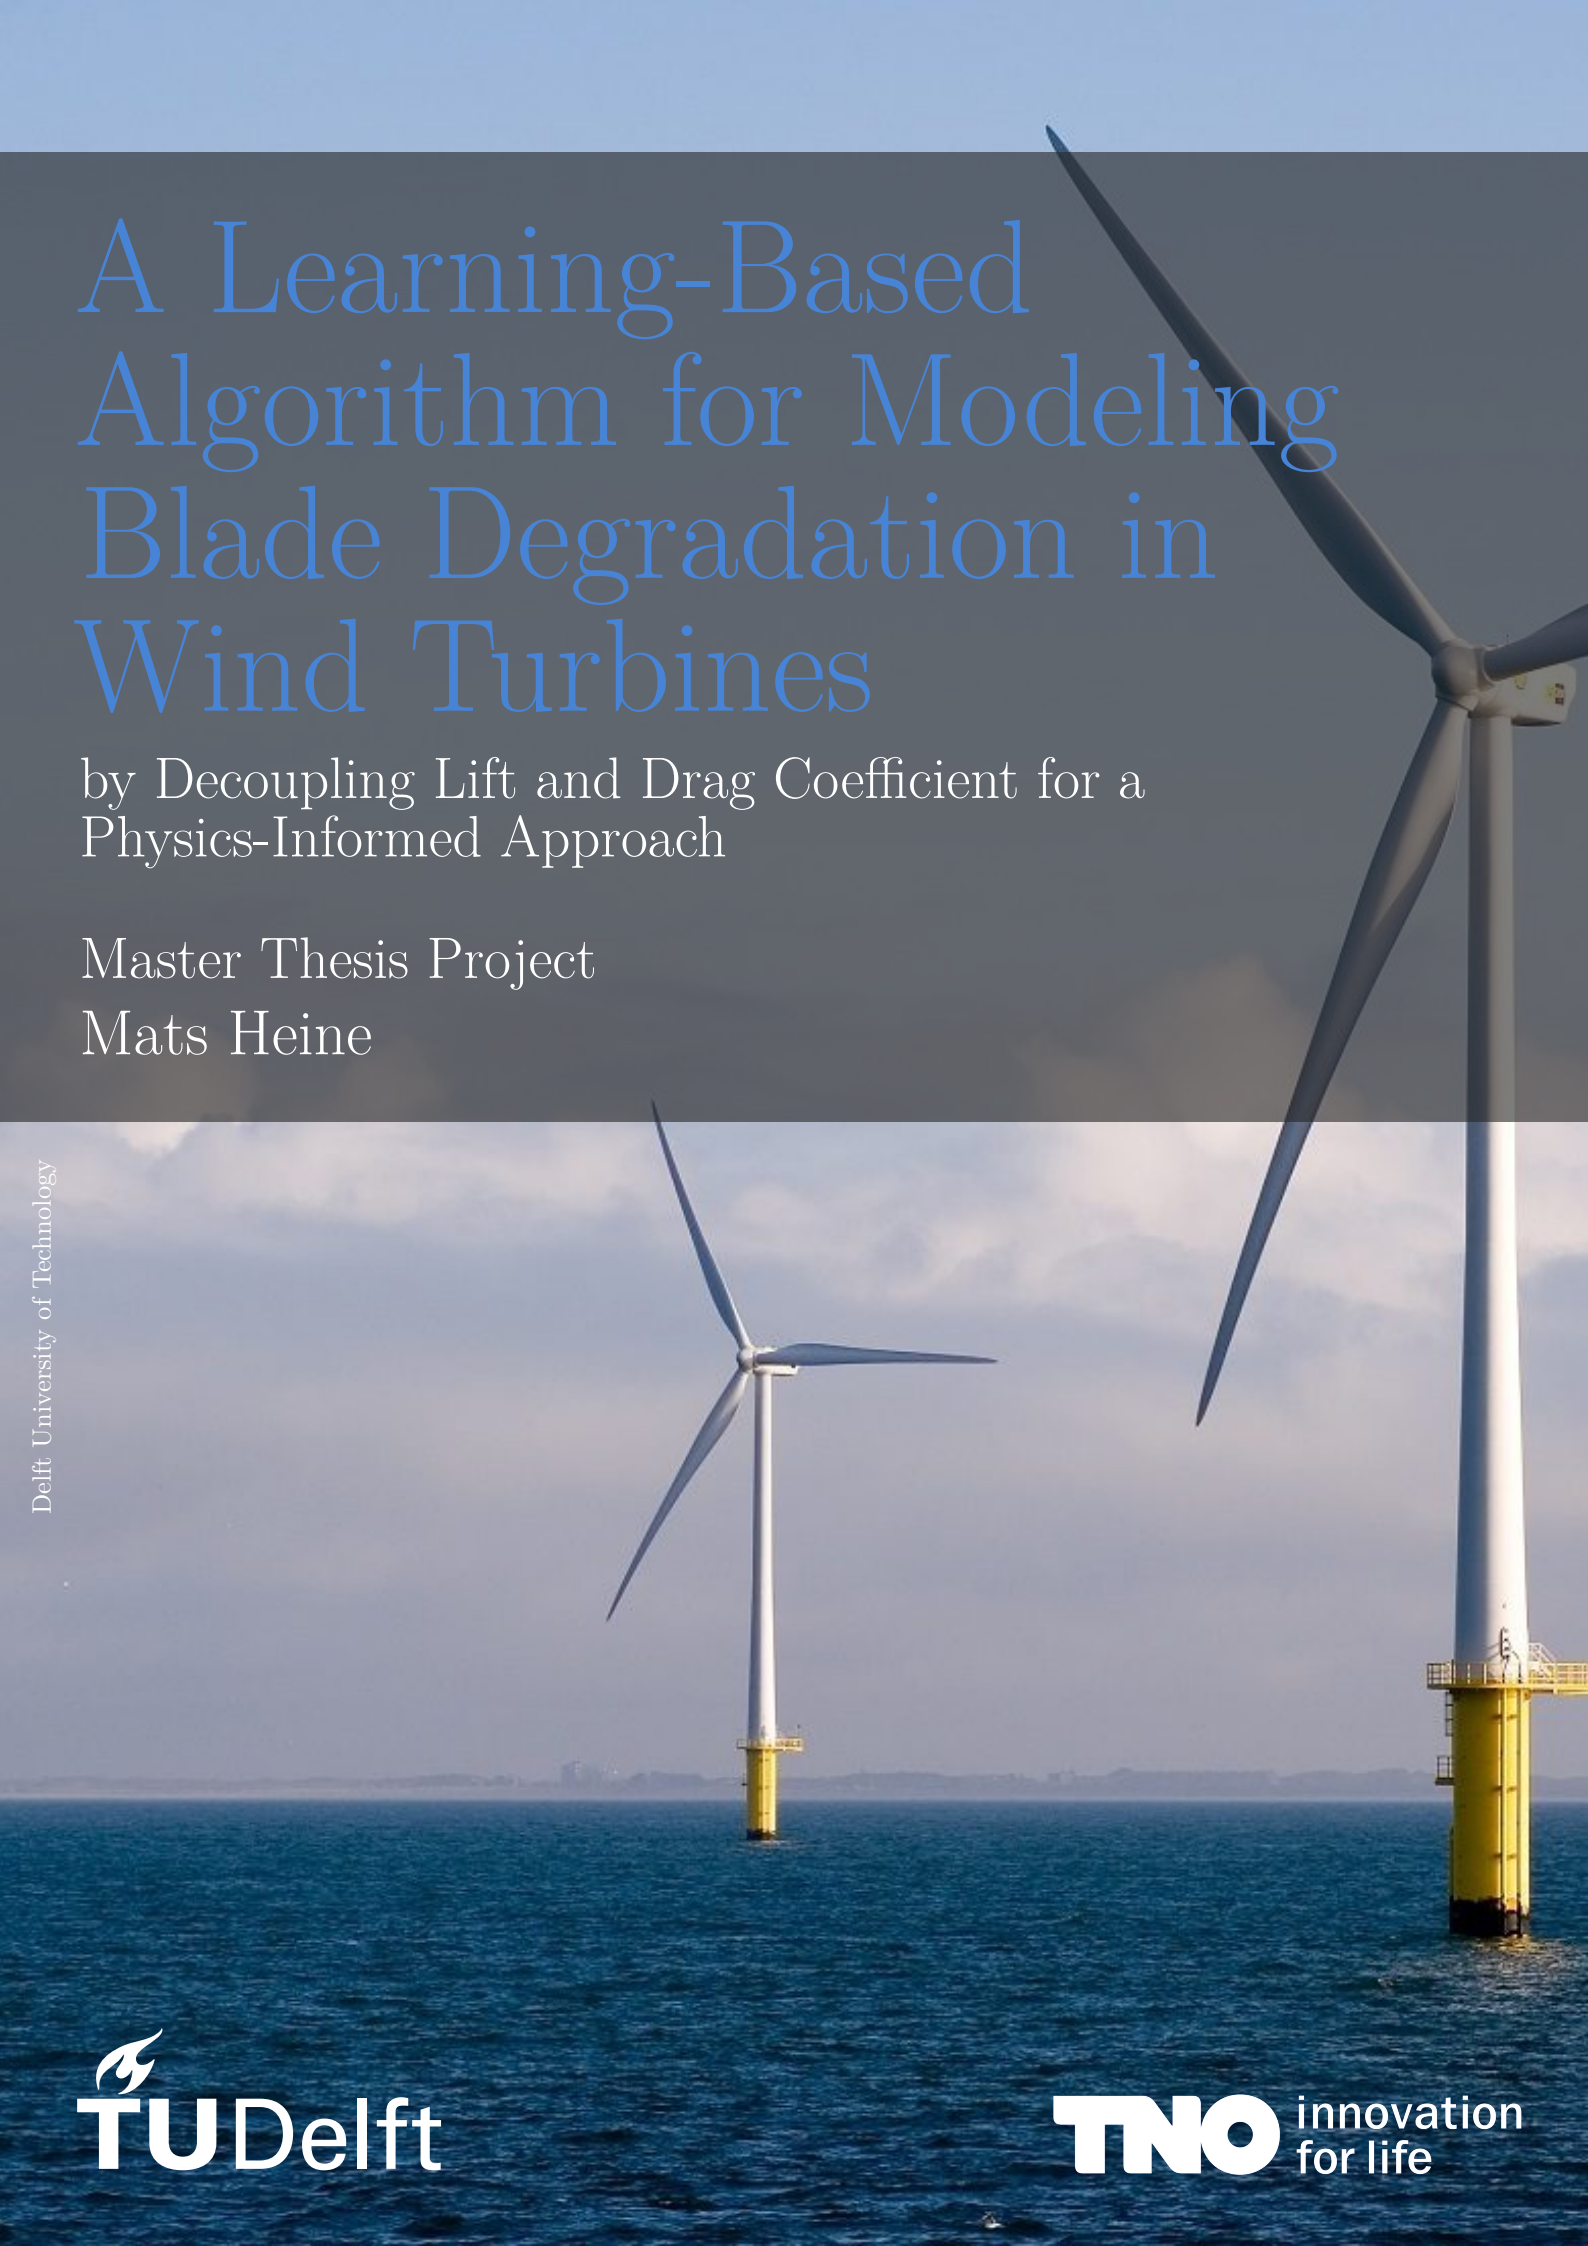

TU Delft

TNO innovation for life

A Learning-Based Algorithm for Modeling Blade Degradation in Wind Turbines

by Decoupling Lift and Drag Coefficient for a Physics-Informed Approach

by

Mats Heine

Student number

5137373

Responsible TNO Supervisor:	Dr. Y. Liu
Responsible TU Supervisor:	Prof. Dr. H.M. Schuttelaars
Additional Supervisor:	Dr. Ir. S.P. Mulders
Additional Supervisor:	Dr. D. Zappalá
Project Duration:	November, 2024 - July, 2025
Faculty:	Faculty of Applied Mathematics, Delft

Cover: Offshore wind farm from an article on offshore wind energy by Deltares [1].

TU Delft

Acknowledgement

This thesis marks the finalization of my master's program in Applied Mathematics, with a specialization in Computational Science and Engineering, at Delft University of Technology. In this project, I have had the opportunity to explore the interdisciplinary challenge of wind energy at the intersection of mathematics, engineering, and renewable energies. It has been a demanding but rewarding process.

I would first like to express my gratitude to my responsible academic supervisor, Prof. Dr. H.M. Schuttelaars, for your clear guidance and constructive feedback throughout the entire process. Your expertise and experience in mathematical modeling were crucial to the structure and clarity of this work. I am equally grateful to Dr. Ir. S.P. Mulders from the Faculty of Mechanical Engineering. Your knowledge of control theory and wind turbine systems helped build the foundations of my work, and your feedback helped me find the direction of the project. To Dr. D. Zappalá from the Faculty of Aerospace Engineering, thank you for your helpful discussions and critical perspective on the broader aerospace engineering context of degradation modeling.

I would like to sincerely thank my TNO supervisor, Dr. Y. Liu, for the opportunity to conduct this thesis within TNO. I appreciated the collaboration and the freedom to explore ideas while staying focused on the key goals of the project.

Finally, I want to thank my family and friends for their encouragement and support throughout my studies. Your presence has helped me find the right balance and stay motivated.

For now, I hope you enjoy reading this thesis!

Mats Heine
Delft, July 2025

Abstract

As the world shifts toward sustainable energy, offshore wind power is becoming increasingly important because of its ability to capture stronger and more consistent winds. However, offshore wind turbines operate in harsh environmental conditions, exposing wind turbine blades to loads. These conditions accelerate blade degradation through mechanisms such as leading-edge erosion, which disrupt airflow, increase drag, and decrease lift, ultimately reducing the aerodynamic efficiency and energy output of the turbine. Over time, this degradation alters the aerodynamic characteristics of the turbine, posing significant challenges to control strategies.

This work addresses this challenge by proposing a physics-informed, learning-based framework for modeling blade degradation in wind turbines. Unlike earlier studies that model degradation as a scalar reduction in performance (typically as a uniform downscaling of the power coefficient C_P), the method developed here decouples the impact on lift and drag coefficients. This approach enables the reconstruction of complex, shape-altering changes in the $C_P(\lambda)$ curve. The result is a generalized degradation framework, parameterized through sensitivity coefficients k_1 and k_2, that represents the degradation trajectory in a degradation severity space. These sensitivity coefficients are not static but evolve over time, enabling a dynamic and interpretable representation of aerodynamic degradation.

The framework is implemented into the Wind Speed Estimation - Tip-Speed Ratio (WSE-TSR) control scheme, a closed-loop scheme that is capable of learning. Within this control scheme, the degraded turbine dynamics are compared with an internal reference model to derive estimation errors, which in turn are used to iteratively update the internal degradation parameters. The result is a correcting controller that learns degradation severity and tracks the evolution of the $C_P(\lambda)$ curve over time.

Initial simulation results, based on modeled degradation paths and a realistic leading-edge erosion case, revealed limitations in the learning process caused by systematic errors in the control scheme. These errors led to discrepancies between the estimated and actual degraded aerodynamic profiles. However, after identifying and analytically correcting for this model error, the learning algorithm demonstrated significantly improved performance. It was then able to reconstruct accurate degraded $C_P(\lambda)$ curves for moderate and severe degradation scenarios. A discussion is provided on the origin of the learning inconsistencies and how they could be further investigated. The framework underscores the potential of the framework to support integration into digital twin systems for adaptive control in a degraded state.

By combining physics-based simulation, parameter learning, and a control scheme, this thesis presents a framework that enhances our ability to monitor and interpret degradation in offshore wind turbines. In particular, it provides physics-based insight into how degradation mechanisms, such as leading-edge erosion, affect aerodynamic performance through changes in lift and drag. This approach contributes to the development of more accurate and adaptive digital twin models that can account for suboptimal aerodynamic performance due to blade degradation.

Contents

Nomenclature

Symbols

Symbol	Definition	Unit
U	Flow velocity	[m/s]
U_∞	Undisturbed wind speed far upstream	[m/s]
U_d	Flow velocity at the actuator disk	[m/s]
U_W	Flow velocity far downstream	[m/s]
P	Power output of the wind turbine	[W]
P^+, P^-	Pressure in front/behind the actuator disk	[Pa]
P_∞	Ambient static pressure far upstream	[Pa]
$P_\mathrm{d}^+, P_\mathrm{d}^-$	Pressure at the actuator disk front/back	[Pa]
C_P	Power coefficient	[-]
$C_P^\diamond$	Undegraded power coefficient curve	[-]
$\widehat{C}_P$	Estimated power coefficient	[-]
C_T	Thrust coefficient	[-]
C_L	Lift coefficient	[-]
C_D	Drag coefficient	[-]
$\Gamma(\lambda)$	Degradation function over tip-speed ratio	[-]
$\gamma_{C_L}(\lambda)$	Lift-related degradation shape	[-]
$\gamma_{C_D}(\lambda)$	Drag-related degradation shape	[-]
a	Axial induction factor	[-]
λ	Tip-speed ratio (TSR)	[-]
$\lambda^\star$	Optimal tip-speed ratio for undegraded blade	[-]
λ_opt	Optimal TSR for degraded blade	[-]
α	Degradation magnitude parameter	[-]
k_1, k_2	Sensitivity parameters for lift and drag degradation	[-]
β	Blade pitch angle	[deg]
α (AOA)	Angle of attack	[deg]
ϕ	Inflow angle to the blade section	[deg]
r	Radial position along blade span	[m]
R	Rotor radius	[m]
c	Chord length of blade element	[m]
B	Number of rotor blades	[-]
A_∞	Cross-sectional area upstream	[m^2]
A_d	Cross-sectional area at the actuator disk	[m^2]
A_W	Cross-sectional area downstream	[m^2]
ρ	Air density	[$\mathrm{kg/m}^3$]
Ω	Rotor angular velocity	[rad/s]
W	Relative wind speed at blade element	[m/s]
F_n	Normal force on blade element	[N]
F_t	Tangential force on blade element	[N]
dL, dD	Differential lift/drag forces	[N]
dr	Blade element spanwise width	[m]

1

Introduction

1.1. Overview of Wind Energy Generation

Wind energy has become a core component of the transition toward more sustainable power systems, as global demand for renewable energy increases. Offshore wind energy, in particular, is experiencing rapid growth due to its ability to capture more consistent and stronger winds than onshore wind turbines [2, 3]. The latest WindEurope 2024 Report [4] highlights that Europe installed 6.4 GW of new wind power capacity in the first half of 2024, of which 5.3 GW (83%) was onshore and 1.1 GW offshore. The EU-27[1] accounted for 5.7 GW of this capacity. This is slightly below the level seen in the same period of 2023. However, year-on-year comparisons of installation volumes can be misleading, as they do not account for structural and procedural factors such as auction timings, permitting delays, grid constraints, and port capacity bottlenecks. Notably, the first half of 2024 saw a marked increase in turbine orders and awarded capacity in auctions. Europe ordered 9.5 GW of new wind turbines in H1 2024. This was 11% up on the first half of 2023. The 9.5 GW breaks down to 7.4 GW onshore and 2.1 GW offshore. This shows that the underlying pipeline of future projects is strengthening. As illustrated in Figure 1.1, most new capacity in the first half of 2024 came from Germany (1.7 GW), France (1.2 GW), and Spain (0.9 GW), reflecting the continued dominance of a few core markets.

This increased attention to wind energy is closely tied to long-term climate objectives. Wind energy plays a crucial role in achieving carbon neutrality goals, with many governments setting ambitious zero-emission targets within the next two to three decades, aligned with IPCC recommendations [5]. By 2030, the EU aims to reach 425 GW of wind capacity, though projections suggest a more realistic estimate of 350 GW due to challenges in permitting, grid expansion, and supply chain limitations [4]. Figure 1.2 illustrates the expected annual installations, emphasizing the growing role of offshore wind.

While onshore wind still represents the majority of installed capacity, policymakers and investors are increasingly prioritizing offshore development to meet long-term climate and energy goals. Larger turbines are in higher demand to optimize power production [6], and moving wind farms offshore avoids societal constraints such as land-use conflicts and public opposition [7]. Additionally, it allows for a scalable and efficient way to produce energy. However, offshore wind still faces significant challenges: high operations and maintenance (O&M) costs, limited grid integration capacity, and bottlenecks in installation logistics. O&M costs can account for up to 30% of a wind farm's lifetime costs [8]. These obstacles have led to delays in major offshore wind projects, such as Dogger Bank A in the UK, which only installed 38 MW in 2024 instead of its planned 1.2 GW [4].

In response, operators are increasingly turning to advanced digital technologies. One notable innovation is the use of Physics-Informed Digital Twins (PDTs), which integrate real-time monitoring, predictive maintenance, and fault detection to reduce downtime and extend turbine lifespan and reduce downtime [9].

[1]The EU-27 refers to the 27 member countries of the European Union, following the departure of the United Kingdom (Brexit) in 2020.

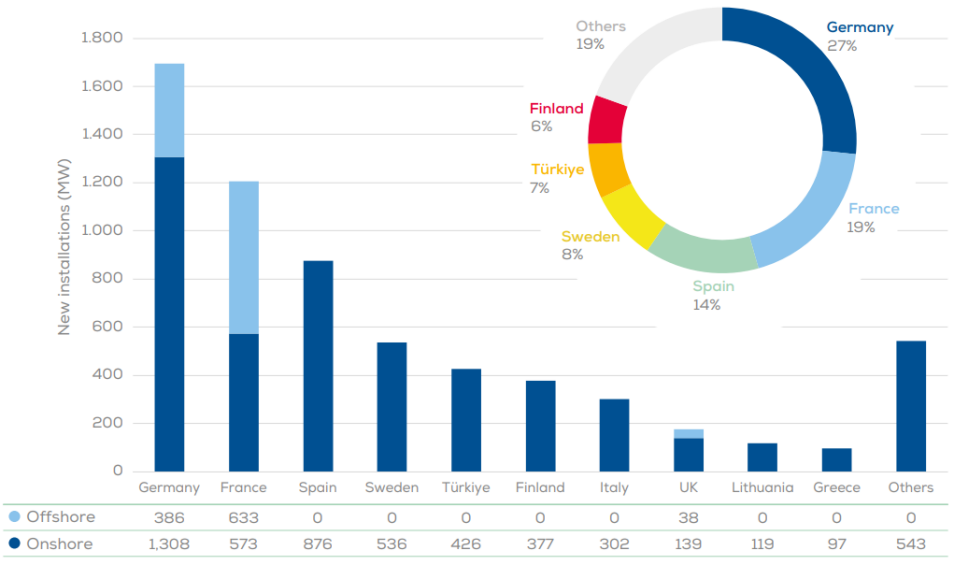

Figure 1.1: New onshore and offshore wind installations in Europe in the first half of 2024. The majority of new capacity came from Germany (1.7 GW), followed by France (1.2 GW) and Spain (0.9 GW). Offshore wind additions remain lower than anticipated, reflecting project delays. [4]

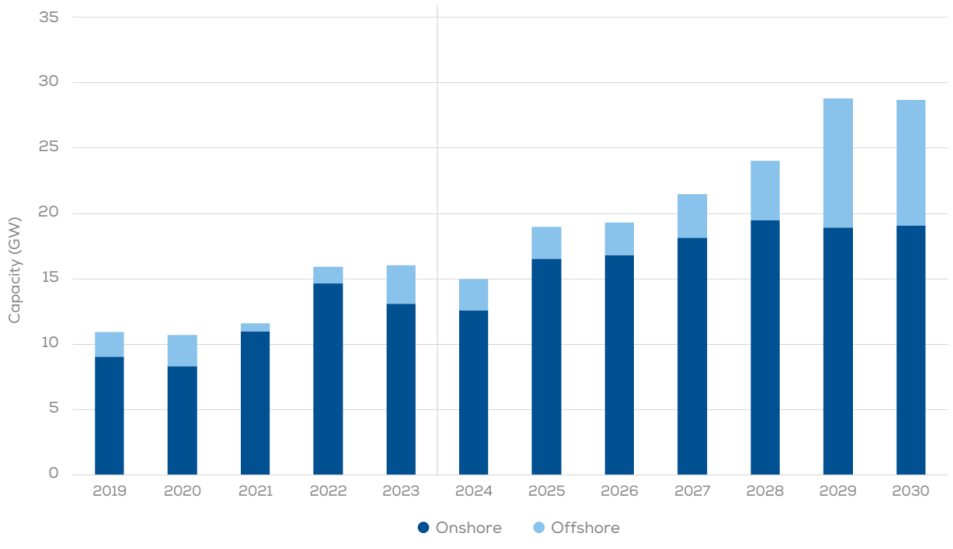

Figure 1.2: Projected annual wind power installations in the EU from 2024 to 2030. While installations are expected to approach 30 GW in 2030, the average annual rate over this period is forecast at 22 GW, which remains below the 33 GW per year required to meet the EU's 2030 climate and energy targets. [4]

1.2. Wind Turbine Generalities

To fully understand the role of wind energy in the renewable energy sector, it is essential to explore the principles of wind turbine design and operation. The following section provides an overview of wind turbine technology, key components, operational principles, including the aerodynamics and control strategies, and concluding with the degradation of wind turbines.

The evolution of modern wind turbines began in the late 19th century with Charles Brush's "Brush Mill," an early example of wind-powered electricity generation. Key advancements followed in Denmark, where Poul la Cour worked on aerodynamics and energy storage, laying the foundation for modern wind energy systems. By the mid-20th century, grid-connected prototypes had introduced pitch control and expanded wind energy's practical potential [10]. The oil crises of the 1970s and 1980s stimulated interest in wind power, which led to innovations in both horizontal and vertical-axis turbines. While larger experimental models from corporations like NASA and Boeing encountered issues with cost and reliability, smaller manufacturers succeeded with iterative improvements, scaling turbines from 50-kilowatt machines in the 1980s to one- to two-megawatt systems by the 1990s. Offshore wind development further accelerated this trend, enabling turbine capacities of five to eight megawatts by the early 21st century. Since 2016, turbine capacities have continued to increase, with prototypes reaching 14 to 16 megawatts by 2022. Notable examples include Vestas' V236-15.0 and MingYang Smart Energy's MySE 16.0-242, both with rotor diameters exceeding 230 meters. These advancements highlight the rapid pace of development and the competitive nature of the wind energy sector, particularly in offshore applications.

While technological innovations have significantly improved turbine performance, the cost of wind energy has remained a central challenge [11]. To address this, three fundamental factors have reduced the cost of wind energy: increased hub height (which allows turbines to access stronger and more consistent winds at higher altitudes), higher power rating (the maximum electrical output the turbine can generate), and larger rotor diameter (which increases the swept area and thus energy capture). These factors directly influence energy capture, which is governed by the fundamental equation for wind turbine power output:

$$P = \frac{1}{2}\rho C_{\mathrm{P}}(\lambda, \beta) A U^3, \tag{1.1}$$

where P is the instantaneous power produced, ρ is the air density, $C_{\mathrm{P}}(\lambda, \beta)$ is the power coefficient, A is the swept area of the rotor, and U is the free-stream wind velocity. The power coefficient C_{P} is a dimensionless measure of aerodynamic efficiency that indicates how effectively the turbine converts the kinetic energy of the wind into mechanical power. Importantly, C_{P} depends on two control-relevant parameters: the tip-speed ratio λ, defined as the ratio of the blade tip speed to wind speed, and the blade pitch angle β, which controls the angle of attack of the blades. Therefore, $C_{\mathrm{P}} = C_{\mathrm{P}}(\lambda, \beta)$, and optimal turbine operation involves adjusting λ and β to maximize C_{P} under varying wind conditions. A more detailed explanation of the relationship between C_{P}, λ, and β is provided in Section 1.2.2. Design improvements that increase A or allow for better control of λ and β can lead to significantly higher energy capture. More power per turbine ultimately allows for fewer installations, reduced system costs, and lower maintenance requirements [12].

While Equation (1.1) describes the instantaneous power generated by a wind turbine at a given moment, a more comprehensive performance metric is the Annual Energy Production (AEP). AEP quantifies the total amount of energy a turbine produces over an entire year, typically expressed in megawatt-hours (MWh). It accounts for varying wind speeds, turbine control strategies, availability, and efficiency losses. Because AEP integrates the turbine's power output over time and wind conditions, it is a key figure in both technical assessments and economic viability of wind energy projects.

1.2.1. Anatomy of a Wind Turbine

This subsection provides an overview of the structural and functional components of horizontal-axis wind turbines (HAWTs), which are the focus of this thesis. The description is primarily based on the Engineer's Guidebook to Wind Energy [13], which offers a comprehensive introduction to wind turbine mechanics and classifications.

The power is extracted from the wind and is converted into mechanical energy, generating electricity.

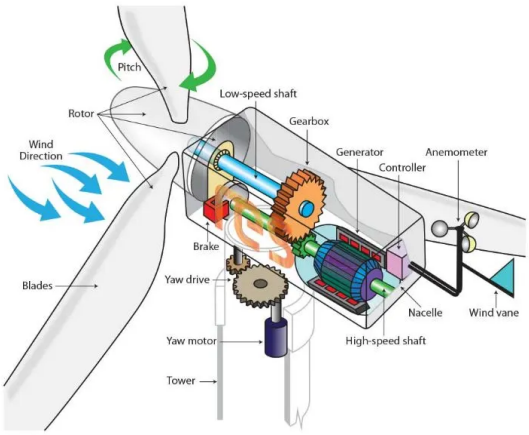

Figure 1.3: Anatomy of a Horizontal Axis Wind Turbine. This figure shows the main components of a wind turbine. Wind rotates the blades, driving the low-speed shaft connected to a gearbox, which increases rotational speed for the generator to produce electricity. The nacelle houses the gearbox, generator, and control systems. A yaw drive aligns the turbine with wind direction, while pitch control adjusts blade angles for optimal performance. Sensors like the anemometer and wind vane provide data for real-time adjustments. The tower supports the structure, elevating it to capture stronger winds. [13]

This process begins when the wind rotates the blades, which are mounted on a central component called the hub. The hub connects the blades to the rotor, forming the part of the turbine that spins as a unit. The rotor's motion is transferred through a drive train to a generator, where electricity is produced. A wind turbine typically consists of three essential components: the tower, blades, and transmission system. The tower provides structural support, the blades capture the wind's energy, and the transmission system, including the rotor, drive train, and generator, facilitates the conversion of kinetic energy into electricity.

The transmission system in a wind turbine is a sophisticated mechanical setup with several configurations. Depending on the orientation of the rotor, wind turbines are classified into two main types: vertical-axis and horizontal-axis turbines. Vertical-axis wind turbines feature rotors positioned near the ground. However, their practical use is limited due to high maintenance costs, substantial land requirements, and reduced efficiency caused by turbulent wind flows. In contrast, horizontal-axis wind turbines are more commonly used and frequently connected to electrical grids. These turbines position the rotor on a tower, where winds are stronger and less turbulent. Therefore, this thesis focuses exclusively on horizontal-axis wind turbines. Their mechanical operation is examined in more detail. As illustrated in Figure 1.3, most horizontal-axis wind turbines have a transmission system with two or three blades. The rotor, acting as the low-speed shaft, rotates at an angular velocity ω_r, generating a torque T_r. This torque is transferred to the generator through a gearbox, with the generator operating at an angular velocity ω_g, producing a torque T_g. The shafts, gearbox, and generator are housed in the nacelle. Additionally, a yaw drive system rotates the nacelle to align the turbine with the incoming wind direction. This mechanism relies on a motor and controller. Adjusting the pitch angle of the blades further optimizes performance by managing structural loads and maintaining consistent power output.

1.2.2. Operational principles

To understand how wind turbines adapt to varying wind conditions, it is useful to categorize them by their speed and pitch control strategies. Some operate with fixed speed or fixed pitch, while others use variable speed or variable pitch configurations. In many cases, these strategies are combined to optimize performance. One of the most common and effective designs is the variable speed-variable pitch (VS-VP) turbine, which adjusts both rotor speed and blade pitch angle dynamically to maximize efficiency under changing wind conditions. This control is governed by two key parameters: the tip-speed ratio (λ, TSR) and the blade pitch angle (β). The tip-speed ratio is defined as

$$\lambda = \frac{\omega_r R}{U}, \tag{1.2}$$

where ω_r is the rotor speed, R is the rotor radius and U is the free-stream wind speed, which is the undisturbed wind velocity before it interacts with the turbine rotor.

Since control strategies rely on tuning of the tip-speed ratio and the blade pitch angle, it becomes essential to quantify turbine performance using a metric that reflects this efficiency. The performance of a wind turbine is typically characterized by its power coefficient, C_P, which represents the fraction of available wind power converted into mechanical energy. The coefficient can be calculated using

$$C_\mathrm{P}(\lambda, \beta) = \frac{P}{\frac{1}{2}\rho A U^3}, \tag{1.3}$$

which is Equation (1.1) reformulated. The value of C_P is directly influenced by λ and β, meaning that different combinations of these parameters result in different levels of aerodynamic efficiency. The theoretical upper limit of C_P, known as the Betz limit, is approximately 0.593 and applies to all wind turbines. This limit signifies that no turbine can capture more than 59.3% of the kinetic energy in the wind. This limit will be proven in Section 2.1.1 of the Literature Review.

In real-world conditions, this theoretical efficiency is influenced by environmental variability, making it necessary to track how C_P responds to changes in operating conditions such as λ and β. Turbine manufacturers often provide the maximum power coefficient, $C_\mathrm{P}^{\max}$, as part of the turbine specifications. Additionally, C_P can be expressed as a function of the pitch angle (β) and the tip-speed ratio (λ). For each turbine design, there exists an optimal tip-speed ratio (λ_opt) and optimal pitch angle (β_opt) that maximize the power coefficient (C_P). An example of the dependency of C_P on the tip-speed and blade angle can be seen in Figure 1.4. These curves will be referred to as the $C_\mathrm{P} - \lambda$ curves throughout this thesis.

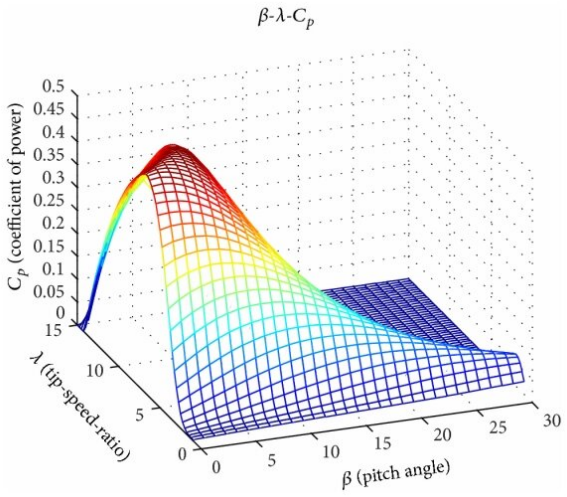

Figure 1.4: Power coefficient (C_P) as a function of tip-speed ratio (λ) and pitch angle (β). This figure illustrates how C_P varies with changes in λ and β for a typical wind turbine. The maximum C_P is achieved at an optimal combination of λ_opt and β_opt, which ensures peak aerodynamic efficiency. [14]

Although TSR-based control aims to keep the turbine at its optimal operating point, recent studies suggest that for large, flexible wind turbines, the assumption of a fixed optimal tip-speed ratio may become less effective due to blade degradation affecting aerodynamic performance [15]. As turbine sizes increase and structural dynamics become more complex, rigid control strategies that maintain a constant λ_opt and β_opt are giving way to more adaptive approaches. These emerging control methods dynamically adjust λ and β in response to real-time operating conditions [16]. Figure 1.5 illustrates

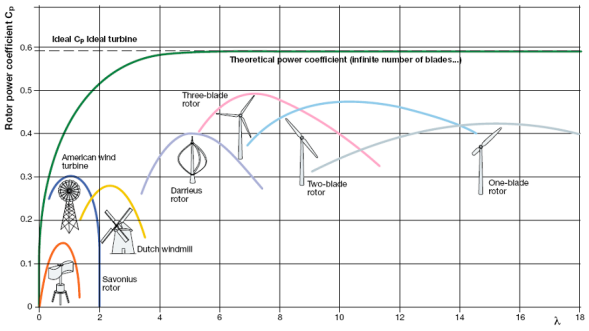

Figure 1.5: Typical power coefficient curves for different types of wind turbines. This figure illustrates the variation of the power coefficient (C_P) as a function of the tip-speed ratio (λ) for different wind turbine rotor designs. The Betz limit (0.593) represents the theoretical maximum efficiency of wind energy conversion. Generally, turbines with more blades operate efficiently at lower λ values, whereas modern horizontal-axis turbines with fewer blades achieve higher peak C_P at higher λ. Understanding these curves is crucial for optimizing wind turbine performance under different operating conditions. [17]

typical C_P curves for different types of wind turbines, with the maximum efficiency occurring at the peak of each curve. This visual highlights how optimal operational points may shift under different conditions, further motivating the need for dynamic control strategies.

To support control development, accurate aerodynamic modeling of wind turbines is essential. One such widely used model is Blade Element Momentum Theory (BEMT), which enables evaluation of the turbine's power coefficient under various conditions.

1.2.3. Blade Element Momentum Theory (BEMT)

To analyze and predict wind turbine performance under various operating conditions, aerodynamic modeling is essential. One widely used method for this is Blade Element Momentum Theory (BEMT). BEMT combines two fundamental principles: Blade Element Theory and Momentum Theory.

- **Blade Element Theory:** Divides the turbine blade into small radial segments and analyzes the local forces acting on each element based on aerodynamic coefficients.
- **Momentum Theory:** Models the rotor as an actuator disk and applies conservation of momentum to the airflow through the rotor plane.

By combining these two approaches, BEMT provides an efficient method to compute the aerodynamic loads, power output, and performance characteristics of a wind turbine based on blade geometry and flow conditions.

1.2.4. Control strategies

To ensure the highest power output, as described by Equation (1.1), the power coefficient must be maximized. Since the power coefficient depends on both the tip-speed ratio and the pitch angle, maintaining optimal values for these parameters is crucial for efficient turbine operation. However, wind conditions are constantly changing, requiring active control strategies to dynamically adjust rotor speed and blade pitch angle. The following section discusses control strategies that allow wind turbines to maximize energy capture in low to moderate wind speeds and regulate power output in high wind conditions, ensuring safe and efficient operation.

Following the categorization proposed by [18], wind turbine operation can be divided into four main regions based on wind speed, as illustrated in Figure 1.6:

1. **Region 1:** Below the cut-in wind speed, no electrical power is produced because the wind speed is too low to overcome mechanical and electrical losses. The turbine remains stationary or in idle mode, managed by supervisory control.
2. **Region 2:** Between the cut-in and rated wind speeds, the turbine operates in the partial-load regime. The goal here is to maximize power extraction. This is done by adjusting the rotor speed to maintain an optimal tip-speed ratio, ensuring that the turbine operates near its maximum

aerodynamic efficiency. During this phase, the pitch angle is held constant at its optimal value, β_{opt}.

3. **Region 3:** When wind speeds exceed the rated value, the turbine enters the full-load regime. The control objective now shifts from maximizing to limiting power. The generator torque is held constant at its rated value, and the pitch angle is actively varied to reduce aerodynamic efficiency by lowering the power coefficient C_{P}. This allows the turbine to maintain constant rated power output while protecting the structure from excessive mechanical loads.
4. **Region 4:** Above the cut-out wind speed, the turbine is shut down to prevent damage from extreme wind conditions. This region is critical for turbine safety.

Between Regions 2 and 3, there is a transitional phase where the rotor speed reaches its rated value, but the generator torque has not yet reached its limit [19]. This phase ensures a smooth transition to Region 3 by gradually increasing torque while holding the rotor speed constant.

The power coefficient $C_{\text{P}}(\lambda, \beta)$ achieves its maximum value at an optimal pitch angle (β_{opt}) and optimal TSR (λ_{opt}), specific to each turbine design. When operating in Region 2 the wind turbine rotor speed should be varied to maintain this optimal tip speed ratio as the wind changes its speed, to ensure the maximum power production. In this region the pitch angle remains constant and optimal at β_{opt}. In Region 3, the control objective changes from maximizing to limiting power, as well as limiting rotor speed and torque. The generator torque is maintained constant at its rated value and the pitch angle should be controlled to reduce the power coefficient and aerodynamic efficiency. Thereby, constant rated power is extracted from the wind.

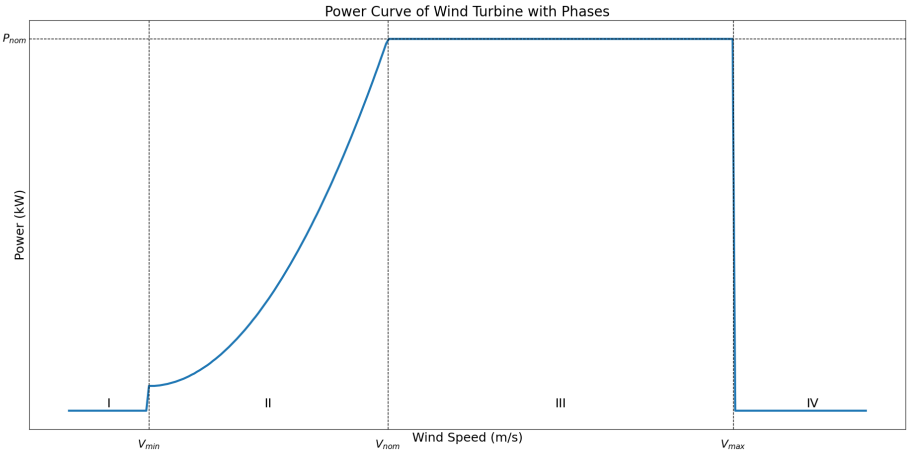

Figure 1.6: Wind turbine operational regions and corresponding power output. This figure illustrates the typical power curve of a wind turbine, divided into four operational regions based on wind speed. In Region 1, below the cut-in wind speed, the turbine remains idle as the wind is insufficient for power generation. Region 2 represents partial-load operation, where power output increases with wind speed, Region 3 marks full-load operation, where power is limited to the rated value by adjusting blade pitch to regulate aerodynamic efficiency. Finally, in Region 4, the turbine shuts down at cut-out wind speed to prevent structural damage in extreme wind conditions. This power curve highlights the control strategies required for safe and efficient turbine operation across varying wind speeds.

In this thesis, the focus will be on Region 2, the partial-load regime, where the primary control objective is to maximize energy capture by maintaining the optimal tip-speed ratio. This regime is chosen as the focus because it is the operational regime where wind turbines generate a substantial portion of their total energy, and where control objectives align directly with maximizing aerodynamic efficiency. Within this regime, the blade pitch angle is held constant at its optimal value β_{opt}, allowing the power coefficient to be modeled as a function of tip-speed ratio only.

1.2.5. Degradation of Wind Turbines

The effectiveness of control strategies is highly dependent on the physical condition of the wind turbine, particularly the state of the blades and other mechanical components. Over time, wind turbines experience degradation due to exposure to harsh environmental conditions and continuous operational

stresses. This degradation can significantly impact the aerodynamic performance of the turbine, altering the power coefficient and reducing the effectiveness of control strategies aimed at maximizing energy capture.

As wind turbines age, their failure rates, defined as the number of failures over a period divided by the operational time of the turbine, increase due to cyclic loads from environmental factors. These forces induce vibrations and stress on mechanical components, leading to progressive wear and eventual failures. The unavailability of wind turbines increases from around 3.2% in the first four years to 5–9% between 14 and 19 years [20]. Key failure mechanisms affecting wind turbines include:

- Blade failure (Wear, erosion, contact fatigue, and cracking),
- Gear and Bearing failures in the nacelle,
- Drive Train failures including Low-Speed Shaft, Gearbox, and High-Speed Shaft [21].

From these blade failure is the most critical failure mechanisms in wind turbine control due to their direct impact on aerodynamic performance, operational safety, and energy yield. The blades capture the energy from the wind, and their degradation can significantly alter the power coefficient, C_{P}, reducing turbine efficiency [22]. Unlike other mechanical components, damage to the blade or hub not only decreases energy capture but can also lead to catastrophic events such as blade cracks and detachment, which pose major safety and operational risks [23]. These failures are also the most costly, requiring specialized equipment and extended downtime for repair [24]. Given their central role in energy production and the severity of consequences with failure, blade and hub conditions remain a critical area of concern for wind turbine reliability.

Leading-edge erosion is a form of blade degradation. Blade erosion affects the entire surface of the blade, while leading-edge erosion (LEE) occurs on the leading edge, the frontmost part that directly faces the airflow. Leading edge erosion directly impacts the AEP, with moderate erosion causing losses of 3–5% and severe erosion leading to reductions of up to 25% over the turbine's lifespan [25].

Key factors contributing to blade erosion include:

- Rain: Larger, high-velocity raindrops cause significant damage due to their kinetic energy.
- Hail: The hardness and density of hailstones make them destructive.
- Abrasive particles: Sand and dust damage the surface layer of the blades.
- Insects: While less damaging, insect residue increases surface roughness.
- Environmental conditions: Humidity, temperature fluctuations, icing, and lightning accelerate material degradation.
- Manufacturing defects: Imperfections such as weak layers reduce the resistance to erosion or material loss.

The aerodynamic performance of wind turbines is directly linked to the condition of the blades. The power coefficient, C_{P}, which determines turbine efficiency, depends on the tip-speed ratio (λ) and pitch angle (β). Blade erosion, particularly LEE, disrupts airflow over the blades, increasing drag and reducing lift, which diminishes aerodynamic performance. Understanding the effect of blade degradation on the optimal power coefficient is critical for optimizing turbine performance.

1.3. Degradation Modeling

Since blade degradation has a direct impact on the power coefficient and therefore on wind turbine efficiency, degradation modeling is a critical topic in the wind energy sector. The power coefficient, C_{P}, plays a central role in determining power, any degradation that alters the $C_{\mathrm{P}} - \lambda$ curve will affect AEP. Accurately modeling such degradation effects is therefore essential for predicting and optimizing long-term energy output. The degradation model can be used in digital twins. In the case of wind energy, a digital twin can be defined as:

"A digital twin is a virtual representation of a physical object, system, or process that calculates system states and makes system information available, through integrated models and data, with the purpose of providing decision support." [26]

The wide applicability of digital twins means that implementations may vary significantly. To address this diversity, a classification system has been proposed [27]. The classification, progressing from simple to advanced stages, includes the following types: Supervisory Digital Twins, Operational Digital Twins, Simulation/Prediction Digital Twins, Intelligent/Learning Digital Twins, Autonomous/Management Digital Twins.

In the context of wind turbine operation and maintenance, digital twins, particularly Supervisory and Simulation/Prediction types, can be used to monitor structural integrity and detect early signs of degradation. One critical application is the early detection of blade degradation, which is essential to ensure the reliability and longevity of wind turbines. Most traditional detection methods rely on direct physical interaction with the turbine, such as ultrasound measurements, vibration analysis, or drone-based visual inspections. However, these techniques are costly and time-consuming. An alternative approach involves integrating existing Supervisory Control and Data Acquisition (SCADA) systems, which help reduce both costs and operational complexity. SCADA systems collect real-time operational data from the wind turbine, including parameters such as rotor speed, generator output, temperature, and blade pitch angle. Analyzing this data allows for continuous performance assessment and anomaly detection. SCADA data offers the potential to be integrated into digital twin models, enabling the evaluation of blade degradation severity and supporting the development of predictive maintenance strategies.

1.3.1. Conventional Degradation Modeling Approaches

A widely used approach for degradation modeling in wind turbine blades focuses on fatigue analysis. Fatigue models are primarily aimed at predicting structural lifespan under cyclic loading by evaluating damage accumulation over time [28, 29]. These models are well-developed for composite materials and often include techniques such as S–N curve-based life estimation, stiffness degradation tracking, and progressive damage modeling. For instance, Rubiella et al. [28] provide an overview of fatigue modeling strategies at both the material and blade level, highlighting methods such as rainflow counting and damage equivalent loads used for lifespan prediction in operational environments.

However, despite their accuracy in forecasting failure or maintenance intervals, conventional fatigue models do not capture the aerodynamic implications of surface degradation. They are primarily concerned with internal damage mechanisms that affect structural integrity rather than aerodynamic efficiency. Shokrieh and Rafiee [29] note that while fatigue simulations can predict where and when failure may occur, they do not indicate how such damage influences performance metrics like the power coefficient (C_P) or optimal tip-speed ratio. This is a critical limitation, especially for control systems that rely on an optimal tip-speed ratio.

In contrast, a physics-based degradation model that dynamically estimates the evolving C_P–λ curve offers a more functional perspective. It allows real-time inference of aerodynamic performance changes from SCADA data, enabling more responsive and efficient turbine operation. This highlights a key gap in current modeling approaches: while fatigue models are excellent for long-term durability studies, they fall short in describing how degradation affects energy capture and operational states of turbines.

1.3.2. Degradation Modeling within the WSE-TSR Control Scheme

In this thesis, the Wind Speed Estimation - Tip Speed Ratio (WSE-TSR) control scheme is used, based on the work of Mulders et al. [30]. While not a digital twin itself, it can be embedded within one to enable real-time learning, degradation modeling, and adaptive control. In practice, this is achieved by using SCADA data (e.g., rotor speed, wind speed, and power output) to detect changes in turbine performance over time.

Modeling Degradation through a Multiplicative Framework

In [30], model uncertainty is integrated into the WSE-TSR control strategy through a multiplicative degradation function, which modifies the ideal aerodynamic rotor properties:

$$C_\mathrm{P}(\lambda) \triangleq \Gamma(\lambda) C_\mathrm{P}^\diamond(\lambda). \tag{1.4}$$

Here, $C_\mathrm{P}^\diamond(\lambda)$ denotes the controller's internal model of the ideal (undegraded) power coefficient curve, corresponding to nominal turbine performance. This model is typically derived from aerodynamic

simulations or manufacturer data. The degradation function $\Gamma(\lambda)$ represents the unknown deviation from the ideal power coefficient. Since $\Gamma(\lambda)$ is not directly measurable in real-world scenarios, an estimated degradation function is used within the WSE-TSR framework:

$$\hat{C}_P(\hat{\lambda}) \triangleq \hat{\Gamma}(\alpha, \hat{\lambda}) C_{\mathrm{P}}^{\diamond}(\hat{\lambda}), \tag{1.5}$$

where the degradation function is parameterized as:

$$\hat{\Gamma}(\alpha, \hat{\lambda}) \triangleq \alpha \gamma(\hat{\lambda}), \tag{1.6}$$

with $\gamma(\hat{\lambda}) : \mathbb{R} \to \mathbb{R}$ representing the predefined degradation profile that captures how degradation varies with the tip-speed ratio, and $\alpha \in \mathbb{R}^+$ being the magnitude scaling factor to reflect the actual severity of degradation. This formulation allows for the estimation of turbine degradation using only operational data such as rotor speed, generator torque, wind speed, and power output.

The WSE–TSR control scheme can be employed as a performance-aware control approach by incorporating the multiplicative degradation framework. The scheme consists of two components: the actual wind turbine and an internal estimator. The wind turbine, which experiences performance degradation, follows a degraded power coefficient curve as described in (1.4). The estimator uses a model-based representation of the turbine, incorporating a degradation function according to (1.5). If the degradation function $\hat{\Gamma}(\hat{\lambda})$ is correctly identified, the estimator and the actual turbine share the same power coefficient curve, and their dynamic behavior aligns. However, in practice, the true aerodynamic degradation is unknown, and the estimator typically assumes a different $C_{\mathrm{P}}(\lambda)$ curve than the real turbine. This mismatch results in a discrepancy between the outputs of the two systems. The degradation function must therefore be learned to compensate for this difference.

The WSE–TSR scheme allows for a learning algorithm that quantifies the misalignment between the estimator and the actual turbine. This is made possible by a key property of the controller: its internal wind speed estimate is influenced by the assumed power coefficient curve. When the internal model deviates from reality, the closed-loop system exhibits systematic errors, notably in the relationship between rotor speed and power output. The learning algorithm uses these observable deviations to update the internal aerodynamic model.

The structure of the WSE–TSR control scheme naturally leads to a convex minimization problem. The learning algorithm adjusts the internal model by minimizing the difference between the normalized derivatives of the real and estimated power coefficient curves:

$$\arg \min_{\hat{C}_P, \frac{\partial \hat{C}_P}{\partial \hat{\lambda}}} \left| \frac{1}{C_{\mathrm{P}}^{2/3}(\lambda)} \frac{\partial C_{\mathrm{P}}}{\partial \lambda} - \frac{1}{\hat{C}_P^{2/3}(\hat{\lambda})} \frac{\partial \hat{C}_P}{\partial \hat{\lambda}} \right|. \tag{1.7}$$

This cost function is derived from a frequency-domain analysis of the closed-loop system, where it is shown that a particular transfer function, relating control input perturbations to power estimation error, becomes minimized when the internal model matches the real aerodynamic behavior. Consequently, learning the correct power coefficient curve, and thereby aligning the behavior of the real wind turbine with that of the internal estimator, becomes equivalent to solving the convex minimization problem.

A full derivation of this learning framework, including the role of excitation signals and the structure of the internal model, is presented in Section 2.4.3. The method used in this thesis builds upon the work of Mulders et al. [30].

The degradation function introduced in (1.6) provides a simple representation of aerodynamic degradation by modeling it as a multiplicative scaling of the nominal power coefficient curve. In this formulation, the parameter $\alpha \in (0, 1]$ uniformly reduces the curve across all tip-speed ratios, while $\gamma(\hat{\lambda})$ defines a fixed shape. As a result, this approach effectively permits only a vertical downward shift of the $C_{\mathrm{P}}(\lambda)$ curve, assuming that the overall aerodynamic shape remains unchanged. While this simplification enables robust and computationally efficient learning, it cannot account for more complex aerodynamic changes such as peak shifts and curvature variations that may result from degradation. To move beyond this

limitation, it is crucial to exploit the information embedded in the structure and evolution of the $C_P(\lambda)$ curve itself. However, existing research does not account for how degradation mechanisms reshape this curve over time. A more refined model would treat $\Gamma(\lambda)$ as a flexible, shape-aware function capable of capturing localized aerodynamic changes rather than applying a global scaling. Integrating such information into the learning framework opens the door to more accurate degradation estimation, the identification of specific failures, and the development of control strategies that adapt to the turbine's evolving condition.

This presents a critical knowledge gap: degradation functions $\Gamma(\lambda)$ used in simulations and digital twin frameworks are often generic approximations rather than a physics-based framework of how real degradation affects turbine aerodynamics. To bridge this gap, a more detailed study of how different degradation scenarios influence the $C_P - \lambda$ curve is required. This includes investigating the shifts in optimal tip-speed ratio and changes in the maximum power coefficient. By integrating these aerodynamic insights into a learning-based framework, a more accurate degradation function $\Gamma(\lambda)$ can be formulated, leading to improved digital twin models and better control strategies for wind turbine operation.

1.4. Scope and Objectives

The efficiency and reliability of wind turbines are strongly influenced by the degradation of their aerodynamic components, particularly the rotor blades. Traditional modeling approaches often assume degradation manifests as a uniform reduction in the power coefficient (C_P), primarily modeled as a simple scalar shift. However, real-world degradation scenarios induce non-uniform, shape-altering transformations in the C_P–λ curve. These effects challenge existing digital twin frameworks and control strategies, such as the WSE-TSR scheme, which typically rely on simplified degradation models.

This thesis extends the current understanding by investigating how various blade degradation mechanisms affect the full shape of the C_P–λ curve, and how these insights can be used to build a more complex degradation modeling framework. The goal is to incorporate this extended model into the WSE-TSR scheme, allowing for more accurate estimation of the conditions of the wind turbine.

The main research question of this thesis is:

Can shape-altering changes in the C_P–λ curve caused by degradation be mathematically modeled in a framework and used to improve the WSE-TSR learning algorithm for performance tracking and operational optimization?

This question can be divided into the following subquestions, each of which is explained in more detail:

1. **What are the main blade degradation scenarios, and what is known from existing research about their impact on aerodynamic polars, wind turbine performance, and Annual Energy Production (AEP)?**

   Identify the possible degradation scenarios and find how their aerodynamic polars are affected by them. The change in aerodynamic properties change the aerodynamic performance of the wind turbine. The extent of this can be simulated in a BEM solver.

2. **How do these degradation scenarios influence the shape of the $C_P - \lambda$ curve across operational conditions?**

   Investigate how these degradation scenarios cause changes in the C_P–λ curve beyond a simple shift, such as flattening, peak displacement, or slope variation.

   The goal is to find a degradation function $\Gamma(\lambda)$, such that:

   $$C_P(\lambda) = \Gamma(\lambda) C_P^{\diamond}(\lambda), \tag{1.8}$$

   where $C_P^{\diamond}(\lambda)$ represents the ideal, undegraded coefficient. Using blade element momentum theory (BEMT), the effects of physical degradation can be simulated to extract the corresponding functional form of the degradation function $\Gamma(\lambda)$, enabling a physics-based representation of aerodynamic losses.

3. **How can this deeper understanding be used to formulate a more general physics-based mathematical degradation framework?**

   Develop a flexible degradation function $\Gamma(\lambda)$ that captures realistic effects from degradation scenarios. Here, $\Gamma(\lambda, \alpha)$ is the general degradation function that can take any shape. The current learning algorithm assumes that $\Gamma(\lambda, \alpha) = \alpha\gamma(\lambda) = \alpha$, with a real constant over all λ. However, degradation scenarios from BEMT analysis may show variations in $\gamma(\lambda)$ or non-uniform changes across λ. Can these non-uniform changes be captured in a mathematical framework?

4. **How is the proposed degradation framework implemented in the WSE-TSR control scheme, and is this implementation effective?**

   The WSE-TSR scheme is adapted to support complex, shape-altering degradation functions $\Gamma(\lambda)$. This modification enables the use of realistic $C_\mathrm{P} - \lambda$ curves derived from degradation scenarios simulated with a BEMT solver. Furthermore, the approach aims to estimate the severity of degradation associated with these realistic performance curves.

5. **Can the modified learning framework suggest operational adjustments or maintenance actions based on learned degradation profiles?**

   Explore whether the framework can recommend adaptive operational set points. After the implementation of the mathematical framework into the WSE-TSR, the degraded $C_\mathrm{P} - \lambda$ curve can be learned. This curve has potentially altered optimal operational set points, λ_opt. Using the severity of degradation, it may be possible to advise or predict maintenance.

1.5. Organization of the Document

This thesis is structured to provide a logical progression from theoretical foundations to practical implementation and evaluation. The document is organized as follows:

- **Chapter 2: Literature Review and Theoretical Background** — Introduces the fundamental principles of wind turbine aerodynamics, including actuator disk theory and Betz limit, blade element theory, and blade element momentum theory (BEMT). Based on current research and knowledge, it also discusses the causes and effects of blade degradation in wind turbines, especially leading edge erosion.
- **Chapter 3: Methodology** — Describes the degradation modeling framework used to simulate blade degradation and its impact on the $C_\mathrm{P} - \lambda$ curve. It introduces the internal degradation model, explains how degradation is parameterized, and details how the model is implemented into the WSE-TSR control scheme. Lastly, it introduces how degradation is learned using the proposed degradation framework in the WSE-TSR controls scheme.
- **Chapter 4: Results of Degradation Modeling Framework** — Presents simulation results demonstrating how different degradation scenarios influence turbine performance. It evaluates the accuracy of the proposed internal degradation model and its effect on aerodynamic efficiency.
- **Chapter 5: Results of Degradation Model Framework Implementation in the WSE-TSR Control Scheme** — Evaluates the performance of the WSE-TSR control scheme under degraded conditions. This includes the effectiveness of the learning mechanism in estimating degradation parameters.
- **Chapter 6: Discussion** — Reflects on the broader implications for wind turbine control, and compares the proposed method with existing approaches. Limitations and assumptions in the current work are examined.
- **Chapter 5: Conclusion** — Summarizes the main contributions of the thesis, revisits the research questions. It also discusses the potential and limitations of the current work and outlines opportunities for future research.

2

Literature Review & Theoretical Background

This chapter provides the theoretical background and a review of existing research related to wind turbine aerodynamics, blade degradation, and control strategies. It begins with the foundational aerodynamic models used to describe turbine performance, focusing on Actuator Disk Theory and Blade Element Momentum Theory (BEMT), which are applied to generate $C_P - \lambda$ curves for both ideal and degraded turbines. Next, it explores research on how degradation scenarios affect aerodynamic efficiency, particularly through changes in lift, drag coefficients, and the power coefficient (C_P). These effects are illustrated using both experimental data and simulation studies. The modeling of such degradation in OpenFAST is also addressed, with a focus on the AeroDyn module. Lastly, the chapter reviews relevant control approaches, emphasizing the Wind Speed Estimation–Tip-Speed Ratio (WSE–TSR) scheme, which forms the basis for the degradation learning approach used in this thesis. Collectively, these topics establish the necessary foundation for the design and evaluation of a physics-based framework for blade degradation modeling.

2.1. Wind Turbine Aerodynamics

This section explores the aerodynamics of wind turbines, focusing on the key theories that describe their performance. It begins with the Actuator Disk Theory, which simplifies the turbine to an idealized disk and introduces the Betz Limit, establishing the theoretical maximum efficiency of a wind turbine. The section then introduces Blade Element Theory, which refines this model by considering the forces acting on the airfoil of an individual blade segment. Finally, the Actuator Disk and Blade Element theories are combined in the Blade Element Momentum Theory (BEMT), providing theory for understanding and optimizing the aerodynamic performance of wind turbines.

2.1.1. Actuator Disk Theory and Betz Limit

Based on the momentum theory, Betz's work pioneered the use of the actuator disk model and defined the theoretical limit of velocity-driven turbines in open flow [31]. In this model, the turbine is simplified as an actuator disk in one dimension. The model has several assumptions:

(a) The flow is irrotational, steady, incompressible, homogenous, and inviscid.

(b) The disk represents an infinite number of blades, resulting in a uniformly distributed thrust over its surface.

(c) The wake is not rotational.

(e) The static pressures far upstream and far downstream of the disk are equal to the undisturbed environment static pressure.

In the model shown in Figure 2.1, U is the flow velocity, P is the pressure, P^+ is the pressure in front of the disk, and P^- is the pressure behind the disk. The subscript ∞ indicates the condition far upstream, d indicates the condition at the disk, and W indicates the condition far downstream.

As the wind passes through the rotor, the turbine extracts energy by exerting a force on the airflow,

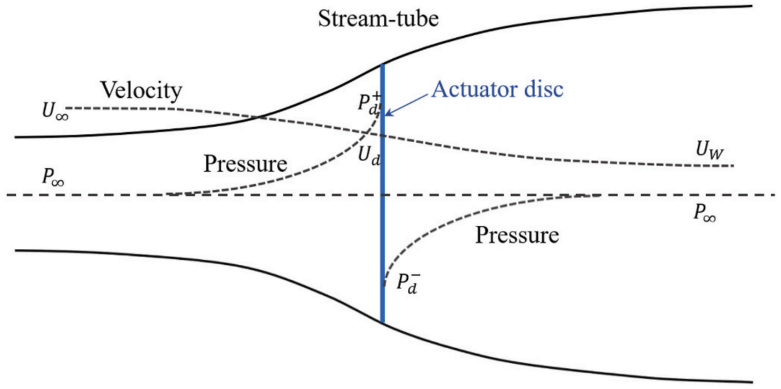

Figure 2.1: Side view of actuator disk and streamtube model. The figure illustrates the velocity and pressure distributions through an idealized wind turbine actuator disk. The streamtube narrows before reaching the disk due to flow acceleration and expands downstream as the velocity decreases. The actuator disk causes a discontinuous pressure drop, with P_d^+ representing the pressure immediately upstream and P_d^- the pressure immediately downstream. The velocities U_∞, U_d, and U_W represent the free-stream velocity, velocity at the disk, and velocity far downstream, respectively. [32]

slowing it down in the process. This force, known as thrust (T), acts in the direction of the incoming wind and plays a key role in determining both the structural loading of the turbine and the wake characteristics downstream.

To quantify this aerodynamic effect, the thrust coefficient (C_T) is introduced, representing the ratio of the thrust force acting on the rotor to the available wind force over the rotor's swept area:

$$C_T = \frac{T}{\frac{1}{2}\rho U_\infty^2 A}, \tag{2.1}$$

where:

- T is the thrust force on the rotor,
- ρ is the air density,
- U_∞ is the free-stream wind velocity,
- A is the rotor swept area.

While C_T describes the aerodynamic loading on the rotor, the power coefficient (C_P) quantifies how efficiently the turbine extracts energy from the wind:

$$C_\mathrm{P} = \frac{P}{\frac{1}{2}\rho U_\infty^3 A}, \tag{2.2}$$

where P is the mechanical power extracted by the rotor. The power coefficient is a crucial metric, as it directly reflects the turbine's efficiency relative to the available wind energy.

Both C_T and C_P are closely related to the axial induction factor (a), which describes the reduction in wind velocity as it passes through the rotor. The axial induction factor is defined as:

$$a = \frac{U_\infty - U_d}{U_\infty}. \tag{2.3}$$

These relationships form the foundation of momentum theory and are critical in understanding the limitations imposed by Betz's Law on the maximum achievable efficiency of a wind turbine. The

relationship between a, the power coefficient (C_{P}), and the thrust coefficient (C_T) of the turbine can be derived using the momentum theory [31].

To illustrate this derivation, the flow through the actuator disk is examined using the streamtube model shown in Figure 2.1. As shown in the figure, the streamtube expands from upstream to downstream. As the flow passes through the actuator disk, the actuator disk exerts a certain thrust force T on the flow, which results in a disturbed flow. In particular, this force will slow down the flow in the area covered by the disk. According to the mass conservation law, the mass flow rate remains constant everywhere:

$$\rho A_\infty U_\infty = \rho A_d U_d = \rho A_W U_W, \tag{2.4}$$

where ρ is the flow density, A_∞ is the cross-section of the streamtube far upstream, A_d is the cross-section at the disk, and A_W is the cross-section far downstream. Since the flow is slowed down by the disk, the cross-sectional area of the stream tube increases.

As mentioned above, the flow passing through the actuator disc is being slowed down, which leads to the momentum change rate of the flow expressed by:

$$\text{momentum change rate} = (U_\infty - U_W)\rho A_d U_d. \tag{2.5}$$

The force causing this momentum change rate is due to the pressure difference before and after the actuator disk:

$$(P_d^+ - P_d^-)A_d = (U_\infty - U_W)\rho A_d U_\infty (1-a). \tag{2.6}$$

Bernoulli's equation can be applied separately to the upstream and downstream sections of the flow. For the upstream region, this yields:

$$\frac{1}{2}\rho U_d^2 + P_d^+ = \frac{1}{2}\rho U_\infty^2 + P_\infty, \tag{2.7}$$

and for the downstream region:

$$\frac{1}{2}\rho U_d^2 + P_d^- = \frac{1}{2}\rho U_W^2 + P_\infty. \tag{2.8}$$

In the actuator disk model, it is assumed that the static pressure far upstream and far downstream is equal to the ambient pressure, i.e., $P_\infty = P_W$. This assumption is based on the fact that sufficiently far from the rotor, the wake has expanded, and any pressure variations induced by the turbine have dissipated. While this is a simplification, it provides a reasonable first-order approximation of wind turbine aerodynamics. In reality, wake recovery is a gradual process, and pressure differences may persist further downstream due to turbulence and atmospheric effects. However, for the purposes of momentum analysis, this assumption allows for a more straightforward derivation of key aerodynamic relationships.

Subtracting the equations for upstream and downstream gives:

$$(P_d^+ - P_d^-) = \frac{1}{2}\rho(U_\infty^2 - U_W^2). \tag{2.9}$$

Equating (2.6) and (2.9) yields:

$$\frac{1}{2}\rho(U_\infty^2 - U_W^2)A_d = (U_\infty - U_W)\rho A_d U_\infty (1-a). \tag{2.10}$$

From this relation, the downstream velocity can be expressed as:

$$U_W = (1-2a)U_\infty. \tag{2.11}$$

The thrust T acting on the actuator disk is given by:

$$T = (P_d^+ - P_d^-)A_d = 2\rho A_d U_\infty^2 a(1-a). \tag{2.12}$$

The thrust coefficient C_T, previously defined in Equation (2.1), can be combined with the thrust expression from Equation (2.12) to yield:

$$C_T = 4a(1-a). \tag{2.13}$$

The power P, defined as the work done by the thrust per second, is given by:

$$P = TU_d = 2\rho A_d U_\infty^3 a(1-a)^2. \tag{2.14}$$

Using the power coefficient C_P as defined earlier in Equation (2.2), and substituting the expression for P from Equation (2.14), we obtain:

$$C_\mathrm{P} = 4a(1-a)^2. \tag{2.15}$$

The power coefficient C_P reaches its theoretical maximum within the actuator disk model when the axial induction factor is $a = \frac{1}{3}$. This optimum can be found by differentiating Equation (2.15), and solving $\frac{dC_\mathrm{P}}{da} = 0$. Substituting $a = \frac{1}{3}$ into (2.15) yields a maximum C_P of $\frac{16}{27} \approx 0.593$, known as the Betz limit. This value represents the theoretical upper bound on the efficiency of any wind turbine operating in an ideal, inviscid flow.

The physical intuition behind the Betz limit lies in the trade-off between extracting energy and maintaining airflow. If the turbine does not slow the wind at all ($a = 0$), no energy is extracted. Conversely, if the wind is slowed too much ($a \to 1$), the flow stagnates, preventing air from passing through the rotor. In that case, the turbine acts as a solid barrier and generates no power. Maximum efficiency is achieved when the wind is slowed just enough to extract significant energy without fully obstructing the flow. Specifically, when the velocity at the rotor is reduced by one-third of the free-stream wind speed. This balance defines the optimal condition at $a = \frac{1}{3}$.

While the Betz limit sets a useful theoretical benchmark, it is based on several idealizing assumptions, including uniform flow, infinite blades, and no viscous losses [31]. Real wind turbines deviate from this ideal due to aerodynamic, mechanical, and electrical losses, as well as blade geometry effects and flow separation. As a result, practical values of C_P for modern turbines typically range between 0.4 and 0.5. To better capture these real-world complexities, the Blade Element Momentum Theory (BEMT) extends the actuator disk model by accounting for the local forces on individual blade elements and incorporating realistic aerodynamic behavior.

2.1.2. Blade Element Theory

While actuator disk theory provides a useful theoretical foundation, it oversimplifies the turbine by assuming uniform flow and infinite blades. To account for the actual geometry and aerodynamic behavior of real turbine blades, the momentum theory is combined with blade element theory, forming the basis of the BEMT.

In blade element theory, the turbine blade is divided into several individual segments, dr, where each segment experiences aerodynamic forces determined by its local flow conditions. This theory assumes that the forces acting on a blade segment can be calculated using the two-dimensional lift coefficient (C_L) and drag coefficient (C_D) of the blade airfoil, with the angle of attack determined by the incident resultant flow velocity. The three-dimensional effects and the velocity component in the spanwise direction of the blades are neglected [18]. A figure of the annular plane with the individual segments, dr, is shown in Figure 2.2.

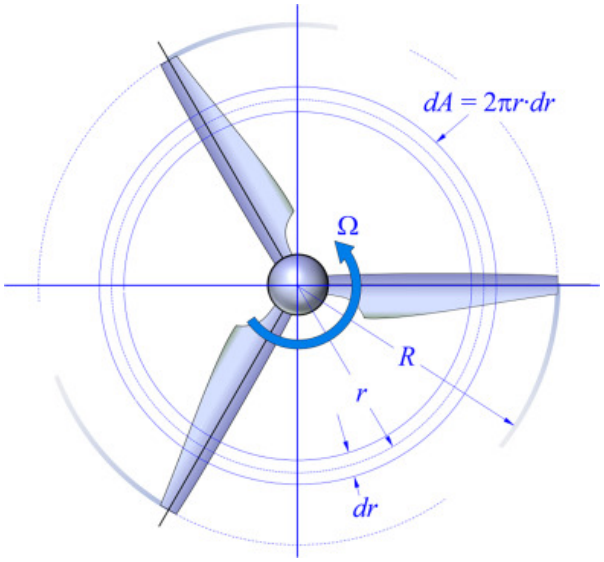

Figure 2.2: Annular plane used in blade element momentum theory. This figure illustrates the annular control volume used in the Blade Element Momentum (BEM) theory, where the rotor disk is divided into differential annular elements of width dr. Each element experiences aerodynamic forces based on local flow conditions, including wind velocity, rotational speed, and blade geometry. The differential area element is given by $dA = 2\pi r\, dr$. [33]

Blade element theory provides a framework to understand the aerodynamic forces acting on the blades of a wind turbine, by calculating these forces in each element of the blades. The aerodynamic behavior of a blade element is often analyzed using a velocity triangle. This triangle represents the combined effect of wind velocity $U_\infty(1-a)$, and the rotational velocity of the blade Ωr, where Ω is the rotor's angular velocity and r is the radial position of the blade element. A figure of such a velocity triangle can be seen in Figure 2.3a.

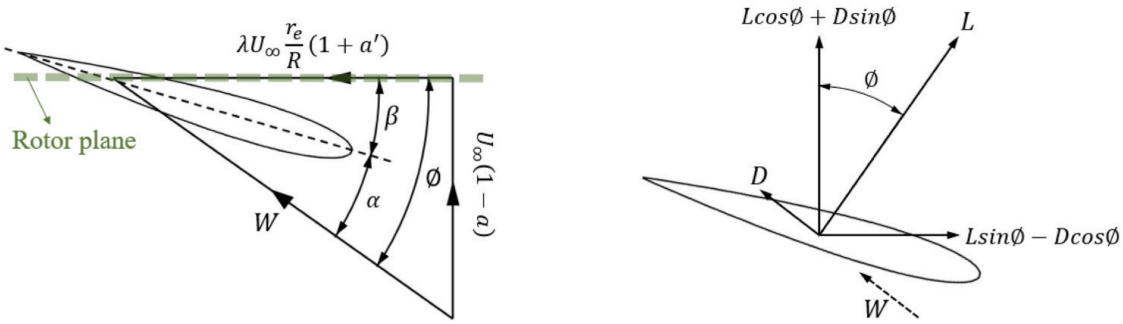

(a) Velocities relating to the wind turbine blade. [32]

(b) Forces acting on a single blade element. [32]

Figure 2.3: Comparison of velocity and force diagrams related to wind turbine blades.

In the velocity triangle, the apparent wind speed W seen by the blade element is the resultant of the rotational and wind-induced velocities. This triangle provides a basis for defining critical angles: the inflow angle ϕ, the angle of attack α, and the section pitch angle β. The inflow angle ϕ is the angle between W and the plane of rotation, while α is defined as $\phi - \beta$. The pitch angle β is a sum of the blade pitch and section twist angles.

The aerodynamic properties of a blade element depend on its shape, particularly its chord length, c. The chord length is the straight-line distance between the leading edge and trailing edge of the aerofoil section of the blade. It varies along the blade span, typically increasing near the root and decreasing towards the tip to optimize aerodynamic efficiency. A schematic of an aerofoil cross-section, illustrating

key geometric parameters such as the chord line, camber line, leading edge, trailing edge, and angle of attack α, is shown in Figure 2.4.

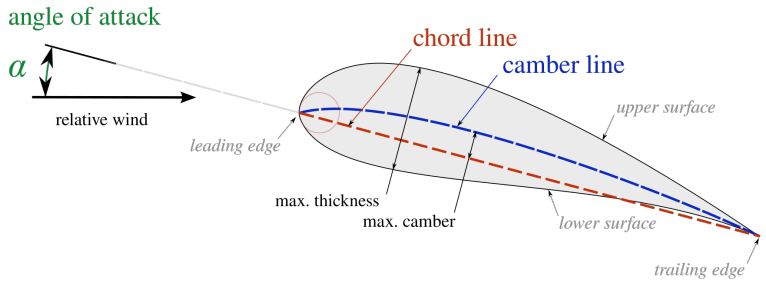

Figure 2.4: Diagram of an aerofoil section, showing the chord line, camber line, leading edge, trailing edge, and angle of attack α. [34]

The forces acting on an aerofoil section of a blade are determined from the velocity triangle. The apparent wind speed W produces a total aerodynamic force R, which can be decomposed into lift L and drag D forces. Lift acts perpendicular to W, while drag aligns with W. These can be calculated using

$$dL = \frac{1}{2} C_L \rho W^2 c \, dr, \quad dD = \frac{1}{2} C_D \rho W^2 c \, dr \tag{2.16}$$

where C_L and C_D are the lift and drag coefficients, ρ is air density, W is the apparent wind speed, c is the chord length, and dr is the blade element length. These coefficients depend on the angle of attack α, which is determined from the inflow angle ϕ and pitch angle β.

The total force can be resolved into components normal (F_n) and tangential (F_t) to the plane of rotation. The normal force contributes to thrust, while the tangential force drives the torque, enabling power generation. The inflow angle ϕ links L, D, F_n, and F_t through trigonometric relations. For an element of the blade with length dr, the lift (dL) and drag (dD) forces are expressed as:

$$dF_n = dL \cos\phi + dD \sin\phi, \quad dF_t = dL \sin\phi - dD \cos\phi, \tag{2.17}$$

Substituting dL and dD into the expression for dF_t, the tangential force is expressed as:

$$dF_t = \frac{1}{2} \rho W^2 c \, C_L \left(\sin\phi - \frac{\cos\phi}{C_L/C_D} \right) dr. \tag{2.18}$$

This equation demonstrates that a higher lift-to-drag ratio C_L/C_D increases the tangential force, enhancing the turbine's power output. Hence, maximizing C_L while minimizing C_D is critical for efficient energy conversion.

2.1.3. Blade Element Momentum Theory

Blade Element Momentum (BEM) theory combines the principles of blade element theory and momentum theory to analyze the forces on wind turbine blades. Blade element theory computes the aerodynamic force on individual blade elements, with the normal force element dF_n depending on lift and drag coefficients and the inflow angle ϕ, which is a function of the axial induction factor a. Momentum theory, on the other hand, calculates the thrust force on an actuator disk as a function of a.

The BEM theory equates the force per annular section derived from these two methods. For an annulus of thickness dr, the force element dF_n from a single blade is multiplied by the number of blades B. This

is equated to the thrust force from momentum theory, expressed using the annular area $2\pi r dr$. Solving this relationship provides an equation for the induction factor a.

The solution is typically iterative. Starting with an initial guess for a, dF_n is computed for each blade element. The resulting equation for a is then solved, and the solution is compared to the initial guess. This process repeats until a converges. Once convergence is achieved, the forces on each blade element are integrated to find the total forces and moments acting on the blade.

The BEM theory makes several assumptions:

- The global forces on the actuator disk are equal to the local forces on the blades.
- Interaction effects between adjacent blade elements are neglected.
- Radial flow over the blades is not considered, and three-dimensional effects are ignored as the theory relies on two-dimensional aerofoil data.
- The tangential force component in the plane of rotation is excluded, focusing solely on the normal component contributing to thrust.

Despite these limitations, BEM theory provides a robust framework for analyzing wind turbine blade performance, forming the basis for many aerodynamic design and optimization processes.

2.2. Blade Degradation in Wind Turbines

As discussed in Section 1.2.5, wind turbine performance is directly linked to the aerodynamic behavior of the blades, which is governed by the lift and drag coefficients. Degradation that alters these coefficients directly impacts the power coefficient C_{P} and overall energy production. The mechanisms and operational impact of blade degradation, such as those caused by rain, snow, icing, dirt, and bug accumulation, were introduced in Section 1.2.5. This section focuses specifically on leading-edge erosion (LEE), which forms the core of this thesis. Other degradation scenarios, including icing and blade cracks, have also been investigated to a limited extent for context and completeness. For clarity and to maintain focus, these additional cases are discussed separately in Appendix A.

2.2.1. Leading-Edge Erosion (LEE)

Leading-edge erosion (LEE) is one of the most studied forms of blade degradation. It is typically caused by particle impacts (e.g., rain, hail, sand) and results in material loss or surface roughness on the blade's leading edge.

CFD-Based Assessment of Leading-Edge Degradation Impact on AEP

Han et al. [35] investigated the aerodynamic effects of leading-edge contamination and erosion on wind turbine blade tip airfoils and the impact on AEP. The analysis was performed on the NACA 64-618 airfoil, this airfoil represents the blade tip region of the NREL 5-MW reference wind turbine.

The study did not use physically degraded blades or real-world experimental data. Instead, the effects of LEE were modeled computationally. The degradation geometries were simulated using CAD-based geometry manipulation. The modified geometries were then analyzed using the CFD software ANSYS Fluent. The simulations were performed under steady conditions to evaluate the aerodynamic coefficients (lift C_L and drag C_D) for both clean and degraded blade geometries. Finally, the altered aerodynamic coefficients were fed into the aeroelastic simulation tool FAST to estimate the impact on AEP. This allowed the authors to quantify the power losses due to multiple degrees of LEE.

The study found that LEE has a substantial negative impact on aerodynamic performance. Specifically, the lift coefficient (C_L) was observed to decrease by up to 53%, while the drag coefficient (C_D) increased by as much as 314% in the most severe cases. These aerodynamic degradations translated into a reduction in the turbine's AEP by approximately 2% to 3.7%, depending on the extent of the surface damage. Overall, the findings highlight the importance of maintaining blade surface quality to preserve turbine efficiency and energy output.

Erosion Effects Using Scanned Blade Profiles

Castorrini et al. [36] investigated the impact of blade erosion on wind turbine performance using a hybrid approach combining real-world data and computational simulations. The researchers recon-

structed erosion geometries of the airfoils by processing detailed laser scan data obtained from both field observations and laboratory experiments. These digitized surface profiles were then used to modify the CAD geometry of the turbine blade. Then these CAD models were used in CFD tools to assess changes in aerodynamic performance, including lift, drag, and pressure distribution. These results were subsequently used to estimate the corresponding loss in Annual Energy Production (AEP).

The study found that erosion has a substantial effect on aerodynamic efficiency. The lift-to-drag ratio decreased by 22%–38% in moderate erosion cases and up to 63% in severe erosion scenarios. These aerodynamic losses translated into reductions in AEP ranging from 1.4% for mild erosion, to 3.9% for moderate erosion, and up to 6.6% for the most severe degradation profiles.

Mitigation of Performance Loss via Tip-Speed Ratio Adaptation

Kumar and Rotea [37] investigated whether retuning control parameters, specifically the optimal tip-speed ratio (TSR), can help recover power or AEP losses caused by blade degradation before physical maintenance is performed. Rather than investigating the effects of this on the AEP, the paper focuses on how these conditions affect the relationship between power coefficient (C_P) and tip-speed ratio (λ). With this knowledge, it could be possible to retune control parameters to recover power and AEP losses before blade repairs are made.

In the paper, the data from [35] is used. Two cases were selected. The first case found that contamination reduced C_L by 30% and increased C_D by 150%. For the second case, erosion led to even more severe degradation, reducing C_L by 50% and increasing C_D by 300% [35]. To evaluate the impact of these aerodynamic changes, the authors modified the C_L and C_D curves in the OpenFAST simulation environment and generated new C_P–λ curves for the degraded blades. The change can be seen in Figure 2.5. For a clean blade, the optimal tip-speed ratio and maximum power coefficient are $\lambda_{opt} = 7.6$ and $C_P = 0.483$ (blue line in Figure 2.5b). With contamination, these values shift to $\lambda_{opt} = 8.2$ and $C_P = 0.431$ (red line in Figure 2.5b). Under erosion, the optimal tip-speed ratio increases to $\lambda_{opt} = 8.4$, while C_P drops significantly to 0.351 (green line in Figure 2.5b). The observed degradation not only reduces the maximum C_P but also shifts the optimal tip-speed ratio to higher values. This shift can be physically explained by how degradation affects the underlying lift and drag characteristics of the blade. A more detailed explanation is provided in the intermezzo box below.

Overall, the study demonstrates that adjusting the tip-speed ratio in response to blade degradation can recover up to 1.5% (contaminated blade) and 3.4% (eroded blade) of lost energy capture, aligning closely with the observed C_P improvements. This approach provides a valuable strategy to mitigate performance losses from blade degradation without requiring physical blade maintenance.

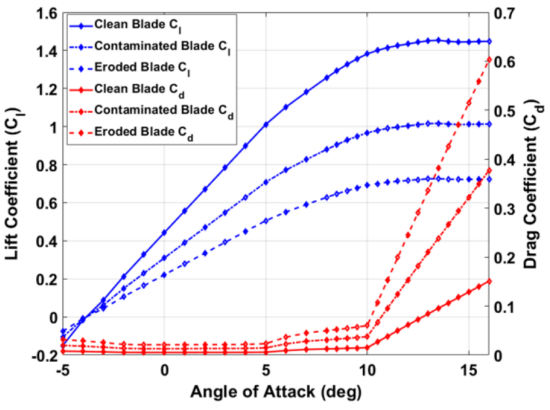

(a) Variation in lift (C_L) and drag (C_D) coefficients due to surface contamination and erosion for the NACA 64-618 airfoil. Contamination and erosion reduce lift and increase drag, degrading aerodynamic performance.[37]

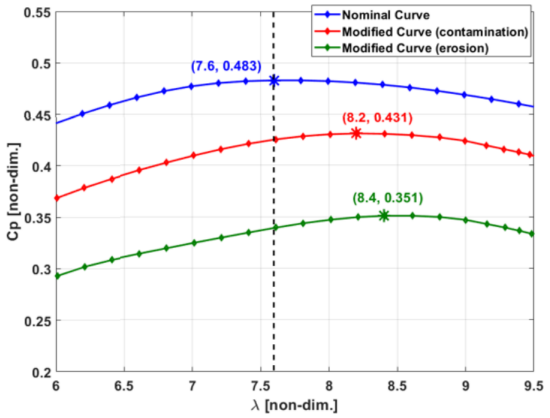

(b) Effect of blade surface degradation on the C_P–λ curve. Erosion and contamination reduce the maximum power coefficient and shift the optimal tip-speed ratio to higher values. [37]

Figure 2.5: Aerodynamic impact of blade contamination and erosion on turbine performance. Degraded airfoil characteristics lead to reduced aerodynamic efficiency and power capture, but adjusting the operational tip-speed ratio can partially mitigate these losses. [37]

Physical Insight: Change of Polars on C_P–λ Curve

It can be seen that when the blade becomes more eroded, the maximum power coefficient, C_P^{max}, decreases, and the optimal tip-speed ratio, λ_{opt}, increases. This decrease in the power coefficient can be explained as follows: Blade degradation alters the aerodynamic properties of the blades, reducing their ability to extract energy from the wind. Specifically, erosion reduces the lift generated by the blades, making them less effective at harnessing wind energy. Consequently, the lift coefficient, C_L, decreases. At the same time, degradation typically increases drag, which means that the drag coefficient, C_D, increases. This increase in drag leads to higher energy dissipation. As a consequence, the overall aerodynamic efficiency of the blades decreases, which is reflected in the reduction of the lift-to-drag ratio, $\frac{C_L}{C_D}$. A lower $\frac{C_L}{C_D}$ ratio reduces the tangential force, dF_t, as described by Equation (2.18). Since the tangential force is responsible for driving the turbine rotor, its reduction leads to a decrease in the extracted rotational energy and, ultimately, a lower power output. Therefore, as the aerodynamic efficiency of the blades diminishes due to increased drag and reduced lift, the turbine's overall ability to convert wind energy into power is reduced. This explains the observed decrease in the power coefficient under the same operating conditions.
The increasing optimal tip-speed ratio can be explained using the velocity triangle in Figure 2.3a. According to Equation (1.2), an increasing optimal value for the tip-speed ratio means that, under the same wind conditions, the rotor speed is higher for degraded blades. As concluded earlier, the more the blades degrade, the less energy they extract from the wind. Consequently, the induction factor, a, decreases. In the velocity triangle, the term $U_\infty(1-a)$ increases when a decreases. Since the wind conditions are assumed to remain unchanged, the inflow angle will also remain constant. From the velocity triangle, it follows that:

$$\tan\phi = \frac{W}{U_\infty(1-a)}. \tag{2.19}$$

To maintain the same inflow angle, while a decreases, the resultant velocity (W), which is the combination of the rotational and wind-induced velocities, must increase. It can therefore be concluded that the rotational velocity must increase, as the wind-induced velocity is assumed constant. The result is a higher optimal tip-speed ratio.

Detecting Power Loss Under Erosion and Turbulence Using Scaled Polars
Malik and Bak [38] investigated how LEE, turbulence intensity (TI), and time-averaging affect the detection of wind turbine power losses under operational conditions. Using the HAWC2 aeroelastic tool and a certified multibody model of a real offshore wind turbine, they represented erosion effects not through geometric blade modification but via scaled aerodynamic polars derived from wind tunnel data. These data, taken from Krog Kruse et al. [39], used sandpaper textures (P400 for moderate erosion and P40 for severe) applied to a NACA 633 airfoil to simulate realistic surface roughness due to rain-induced wear.

To integrate these aerodynamic effects into simulations, Malik and Bak adopted a simple yet effective strategy: modifying the lift and drag characteristics of the airfoil directly through multiplicative scaling. In their model, erosion was applied to the outer 15% of the blade by reducing the lift coefficient (C_L) by 10%, while increasing the drag coefficient (C_D) by factors of 1.5 (P400) and 2.0 (P40). As illustrated in Figure 2.6, this multiplicative approach effectively replicates the degraded aerodynamic performance observed in the measured data. The dotted lines in the figure represent the scaled clean airfoil data, closely matching the experimentally derived polars for the eroded cases. This confirms that moderate and severe erosion effects can be efficiently approximated by scaling the clean C_L and C_D values.

Simulation results showed that this degradation has a measurable impact on turbine performance. For instance, under 6% turbulence intensity, the power coefficient (C_P) dropped by approximately 6.5% at 7 m/s and 4.7% at 11 m/s for the severely eroded (P40) case. Overall, this study demonstrates a practical and accurate framework for incorporating erosion into turbine simulations by scaling aerodynamic coefficients.

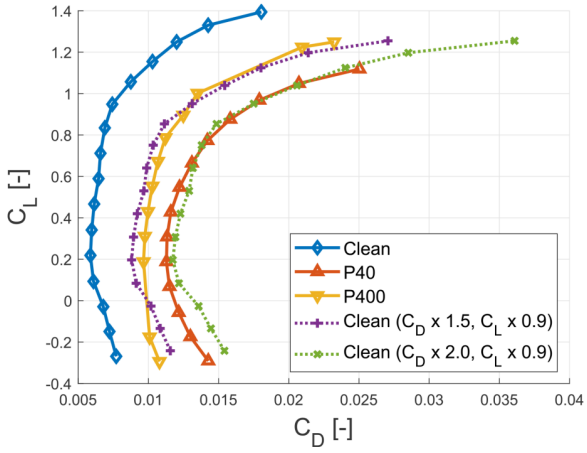

Figure 2.6: Aerodynamic degradation modeled using scaled airfoil polars. The figure illustrates the impact of simulated surface roughness on lift (C_L) and drag (C_D) coefficients. These scaled polars represent moderate and severe leading-edge erosion scenarios, enabling efficient simulation of aerodynamic performance loss in aeroelastic models. [38]

The literature consistently demonstrates that leading-edge erosion (LEE) significantly degrades aerodynamic performance by reducing lift, increasing drag, and ultimately lowering turbine power output. Several studies, including those by Han et al. and Castorrini et al., quantify this effect using CFD and real surface profiles, linking aerodynamic degradation to reductions in Annual Energy Production (AEP). Importantly, Kumar and Rotea show that LEE not only reduces the maximum power coefficient ($C_{\mathrm{P}}^{\max}$) but also shifts the optimal tip-speed ratio (λ_{opt}) to the right and alters the shape of the C_{P}–λ curve. These shape-altering effects underscore the need for degradation models that go beyond simple vertical scaling of performance.

Malik and Bak further support this need by demonstrating that erosion-induced aerodynamic losses can be effectively approximated through multiplicative scaling of lift and drag coefficients. This assumption is adopted in this thesis as the basis for simulating LEE within the degradation modeling framework. Collectively, these findings validate the approach of representing degradation through altered aerodynamic coefficients and justify the focus on modeling both shape and severity of performance loss.

2.3. Simulation of Blade Degradation Scenarios in BEMT

In the previous section, the aerodynamic effects of blade degradation were summarized based on findings from existing literature. Building on these insights, this section presents how degradation scenarios can be implemented within the OpenFAST simulation environment. Using the Blade Element Momentum Theory (BEMT) framework and the AeroDyn module, various damage mechanisms are modeled to investigate their impact on turbine aerodynamic performance in a consistent way. This section focuses on LEE, while other degradation scenarios and their implementation in the OpenFAST framework are discussed in Appendix A.

2.3.1. Overview of OpenFAST and AeroDyn

OpenFAST is an open-source, physics-based simulation tool developed by the National Renewable Energy Laboratory (NREL) for modeling the coupled dynamic response of wind turbines. It provides a framework for modules for different aspects of turbine physics, including ElastoDyn, ServoDyn, HydroDyn, and AeroDyn. In this study, only the aerodynamic behavior of the turbine is investigated in response to blade degradation. Therefore, the focus is limited to the AeroDyn module within OpenFAST. AeroDyn simulates the aerodynamic forces acting on the turbine blades based on airfoil characteristics, operational conditions, and blade geometry. It supports both steady and unsteady aerodynamic models.

The AeroDyn module in OpenFAST computes aerodynamic loads on wind turbine blades using BEMT. This theory was explained in Section 2.1.3. This aerodynamic load is normally passed to other modules in OpenFAST, unless the simulation is run in a decoupled aerodynamic mode. In this thesis, the latter

is the case. Degradation scenarios are implemented by modifying only the aerodynamic input files of AeroDyn.

2.4. Control Strategies

In modern wind turbines, especially in the partial-load region, control strategies play a crucial role in maximizing energy capture. The key control inputs in this region are the generator torque (or power) and the blade pitch angle. The blade pitch angle is typically held constant at its optimal value, β_{opt}, while the generator torque is actively adjusted based on the estimated wind conditions. The tip-speed ratio is affected by changes in rotor speed, which in turn are driven by the applied generator torque. Properly designed torque control allows the rotor speed (and thus the TSR)to remain close to its optimal value. Maintaining the TSR at this optimum is essential, as aerodynamic efficiency and power output are highly sensitive to deviations from the ideal operating point in the partial-load region.

2.4.1. Baseline $K\omega^2$ - Controller

The most commonly used baseline control strategy in the partial-load region is the $K\omega^2$ controller. This method sets the generator torque, T_g, proportional to the square of the rotor speed, ω_r:

$$T_g = K\omega_r^2. \tag{2.20}$$

This relation originates from the objective of operating the turbine at the optimal tip-speed ratio, λ_{opt}, which maximizes the power coefficient, $C_{\text{P}}^{\max}$. In the partial-load region, aerodynamic power is given by:

$$P = \frac{1}{2}\rho A C_{\text{P}}(\lambda) U^3, \tag{2.21}$$

where U is the wind speed and $A = \pi R^2$ is the rotor swept area. Since power is also related to torque by $P = Q\omega_r$, we can rearrange:

$$Q = \frac{P}{\omega_r} = \frac{1}{2}\rho A C_{\text{P}}(\lambda)\frac{U^3}{\omega_r}. \tag{2.22}$$

Now substituting the tip-speed ratio definition, $\lambda = \frac{\omega_r R}{U}$, to express U in terms of ω_r:

$$U = \frac{\omega_r R}{\lambda} \quad \Rightarrow \quad U^3 = \frac{\omega_r^3 R^3}{\lambda^3}, \tag{2.23}$$

we find:

$$Q = \frac{1}{2}\rho A R^3 \frac{C_{\text{P}}(\lambda)}{\lambda^3}\omega_r^2. \tag{2.24}$$

Assuming operation at λ_{opt} and $C_{\text{P}}^{\max}$, the entire coefficient becomes constant:

$$Q = K\omega_r^2. \tag{2.25}$$

Thus, the $K\omega^2$ control law directly enforces operation at the optimal aerodynamic point, maximizing energy extraction under ideal conditions.

This strategy is widely adopted due to its simplicity and its ability to function without requiring wind speed measurements, which are often inaccurate or unavailable due to nacelle flow distortion and spatial variation across the rotor disk. By regulating torque as a function of rotor speed alone, the controller ensures that the turbine tracks the optimal tip-speed ratio and achieves high efficiency in the partial-load regime. A schematic representation of the control framework is illustrated in Figure 2.7.

A limitation of this approach is that it assumes that K remains constant over time, which implies that the turbine's aerodynamic characteristics (e.g., C_{P} and λ_{opt}) do not change. In practice, these values may shift due to blade degradation, erosion, icing, or other factors, making the control law suboptimal

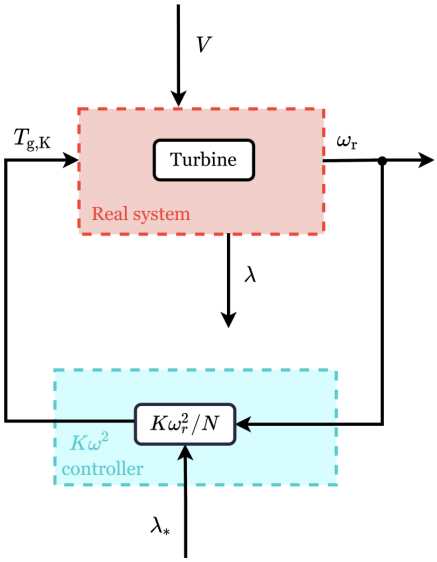

Figure 2.7: The block diagram of the $K\omega_r^2$ control framework is illustrated as follows. The red box highlights the wind turbine system with two inputs (the generator torque, $T_{g,K}$, and the wind speed, V) and two outputs (the rotational speed, ω_r, and the tip-speed ratio, λ). The measured ω_r and the optimal TSR, λ_{opt}, are used as inputs of the controller (cyan box) to compute $T_{g,K}$. [40]

unless recalibrated. To address this limitation, it becomes necessary to adopt a more adaptive control strategy. The Wind Speed Estimation–Tip Speed Ratio (WSE–TSR) control scheme, introduced in the next section, is designed to handle more complex wind turbine efficiency changes.

2.4.2. Wind Speed Estimator and Tip-Speed Ratio Tracking (WSE–TSR)

An advanced alternative to the $K\omega^2$ controller is the wind speed estimator and tip-speed ratio tracking (WSE–TSR) controller. This strategy estimates the rotor-effective wind speed (REWS) in real time and uses it to compute a rotor speed reference that ensures optimal operation. The closed-loop structure consists of an internal model-based wind speed estimator and a controller that minimizes the error between the reference rotor speed and the measured rotor speed. This is done by incorporating uncertainty in the model parameters [41].

Unlike the $K\omega^2$ method, the WSE–TSR controller does not rely on a fixed gain but adjusts its response based on real-time wind conditions and system dynamics. This makes it more adaptive but also more sensitive to model uncertainties and degradation in the turbine's aerodynamic properties.

2.4.3. Derivation of WSE-TSR Control Scheme

The control scheme is illustrated in Figure 2.8. This section describes the derivation of the control scheme and lists a few important properties of the control scheme. The control derivation follows from Mulders et al. [30].

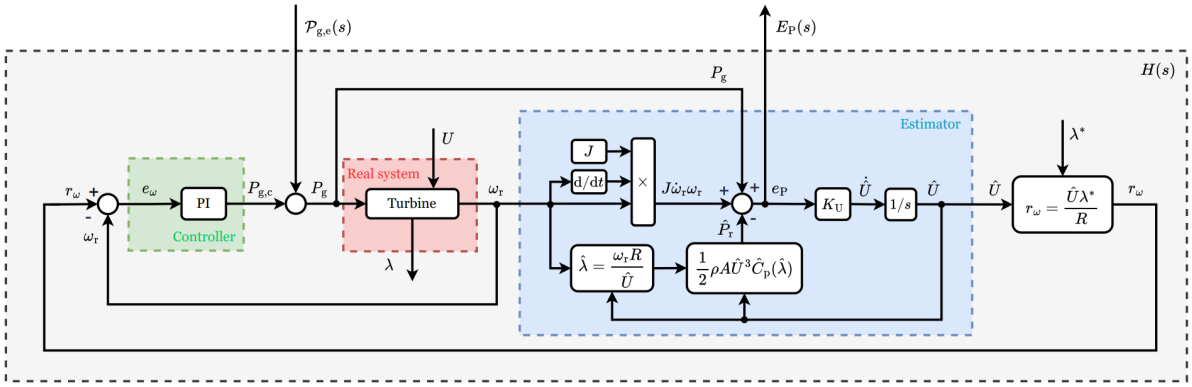

Figure 2.8: Block diagram of the partial-load wind turbine controller. The red box contains the wind turbine system (model), with generator power set point P_g as control input, rotor speed ω_r as output, and subject to a rotor-effective wind speed (REWS) U disturbance. The wind speed estimator is indicated in blue and relies on the system input and output, and a nonlinear internal model of the wind turbine to compute a rotor-effective wind speed estimate $\hat{U}$. This estimate is subsequently used to compute the rotor speed set point r_ω. The set point is used to calculate the error e_ω by subtracting the actual rotor speed ω_r, which is provided as input to the PI controller shown in green. The controller then determines the generator power set point P_g. [30]

The wind turbine is represented as a first-order system:

$$J\dot{\omega}_r(t)\omega_r(t) = P_r(t) - P_g(t), \tag{2.26}$$

where $J \in \mathbb{R}$ denotes the effective rotor inertia at the low-speed shaft (LSS), $\omega_r \in \mathbb{R}$ is the rotor speed, $\dot{\omega}_r(t)$ is the time derivative of the rotor speed, and $P_r \in \mathbb{R}$ and $P_g \in \mathbb{R}$ are the aerodynamic power and generator power set points, respectively. The aerodynamic power P_r is given by:

$$P_r = \frac{1}{2}\rho A U^3 C_P(\lambda), \tag{2.27}$$

where ρ is the air density, A is the rotor swept area, U is the rotor-effective wind speed, and $C_P(\lambda)$ is the rotor power coefficient, which is a function of the dimensionless tip-speed ratio λ defined in Equation (1.2).

The wind speed estimator is a dynamic variant of the commonly applied power balance equation [42], given by:

$$\dot{\hat{U}} = K_U e_P = K_U \left(P_g - \hat{P}_r + J\dot{\omega}_r\omega_r\right), \tag{2.28}$$

where the dot notation (e.g., $\dot{\hat{U}}$) denotes a time derivative, K_U is the estimator gain, and the estimated aerodynamic rotor power is defined as:

$$\hat{P}_r = \frac{1}{2}\rho A\hat{U}^3\hat{C}_P(\hat{\lambda}), \tag{2.29}$$

with $\hat{\lambda} = \frac{\omega_r R}{\hat{U}}$.

The TSR tracking controller is implemented as a proportional-integral (PI) controller:

$$\dot{P}_g = K_p\dot{e}_\omega + K_i e_\omega, \tag{2.30}$$

where again the dot represents the time derivative. The error $e_\omega = r_\omega - \omega_r$ is the difference between the rotor speed set point $r_\omega(t)$, obtained using Equation (1.2) with the desired TSR λ^* and the estimated wind speed $\hat{U}$, and the actual rotor speed ω_r.

Closed-Loop System Dynamics

To analyze the dynamics of the WSE–TSR control system, the closed-loop configuration with the turbine, estimator, and controller is represented in state-space form and linearized for frequency-domain analysis. The state vector is defined as $x = \begin{bmatrix}\omega_r & \hat{U}\end{bmatrix}^T$, the input vector $u = \begin{bmatrix}P_{g,e} & P_{g,c}\end{bmatrix}^T$ includes an excitation component and the control power contribution, and the output vector is $y = \begin{bmatrix}e_\omega & e_P\end{bmatrix}^T$, consisting of rotor speed and power errors.

The nonlinear state-space system is given by:

$$\dot{x} = f(x,u) = \begin{bmatrix} \frac{1}{J\omega_r}(P_r - P_g) \\ K_U\left(P_g - \hat{P}_r + J\dot{\omega}_r\omega_r\right) \end{bmatrix}. \tag{2.31}$$

Linearizing this system around a steady-state operating point yields the following system matrices:

$$A = \begin{bmatrix} \frac{T+\bar{P}_g/\bar{\omega}_r^2}{J} & 0 \\ K_U(Q-\hat{Q}) & -K_U V \end{bmatrix}, \quad B = \begin{bmatrix} -\frac{1}{J\bar{\omega}_r} & -\frac{1}{J\bar{\omega}_r} \\ 0 & 0 \end{bmatrix}, \tag{2.32}$$

$$C = \begin{bmatrix} 1 & -\frac{\lambda^*}{R} \\ Q-\hat{Q} & -V \end{bmatrix}, \quad D = 0, \tag{2.33}$$

where: $T = \partial\tau_r/\partial\omega_r$, $Q = \partial P_r/\partial\omega_r$, $\hat{Q} = \partial\hat{P}_r/\partial\omega_r$, $V = \partial\hat{P}_r/\partial\hat{U}$ and λ^* is the optimal tip-speed ratio.

The transfer function matrix $G(s)$ from input $u(s)$ to output $y(s)$ is obtained using the standard state-space formula:

$$G(s) = C(sI - A)^{-1}B + D = \begin{bmatrix} G_1(s) & G_1(s) \\ G_2(s) & G_2(s) \end{bmatrix}. \tag{2.34}$$

The individual components are:

$$G_1(s) = \frac{-1}{J\bar{\omega}_r}\left(\frac{s + K_U\left(V - \frac{\lambda^*}{R}(Q-\hat{Q})\right)}{\left(s - \frac{T+\bar{P}_g/\bar{\omega}_r^2}{J}\right)(s+K_UV)}\right), \tag{2.35}$$

$$G_2(s) = \frac{-1}{J\bar{\omega}_r}\left(\frac{(Q-\hat{Q})s}{\left(s - \frac{T+\bar{P}_g/\bar{\omega}_r^2}{J}\right)(s+K_UV)}\right). \tag{2.36}$$

The PI controller is defined in the Laplace domain as:

$$C(s) = \frac{P_{g,c}(s)}{E_\omega(s)} = \frac{K_P s + K_I}{s}. \tag{2.37}$$

In series with $G_1(s)$, the open-loop transfer function becomes:

$$L_1(s) = G_1(s)C(s) = \frac{-1\,(K_P s + K_i)\left(s + K_U\left(V - \frac{\lambda^*}{R}(Q-\hat{Q})\right)\right)}{J\bar{\omega}_r\, s\left(s - \frac{T+\bar{P}_g/\bar{\omega}_r^2}{J}\right)(s+K_UV)}. \tag{2.38}$$

Finally, the closed-loop transfer function from the excitation input $P_{g,e}(s)$ to the power error output $E_P(s)$ is:

$$H(s) = \frac{E_P(s)}{P_{g,e}(s)} = \frac{G_2(s)}{1 - L_1(s)}. \tag{2.39}$$

Or, written more compactly as:

$$H(s) = -\frac{(Q - \hat{Q})s^2}{J(\Xi_1 + \Xi_2)}, \tag{2.40}$$

in which $\Xi_1 = \bar{\omega}_r s \left(s - \frac{T + \bar{P}_g/\bar{\omega}_r^2}{J} \right) (s + K_U V)$ and $\Xi_2 = \left(\frac{1}{J} \right) (K_P s + K_i) \left(s + K_U \left(V - \frac{\lambda^*}{R}(Q - \hat{Q}) \right) \right)$

This transfer function $H(s)$ is of particular importance because its magnitude diminishes as the internal model becomes more accurate (i.e., as $Q \to \hat{Q}$). This property underpins the learning-based calibration methods discussed in later sections.

Convex Properties of the Transfer Function and Learning Potential

The closed-loop transfer function $H(s)$, defined in Equation (2.40), possesses a key property that enables model learning: its magnitude approaches zero as the internal model matches the true aerodynamic behavior of the wind turbine. This observation is foundational for the design of a learning algorithm that calibrates the internal model based on system excitation and closed-loop responses.

The DC gain of $H(s)$ is driven by the difference between the actual and estimated partial derivatives of the rotor power with respect to the rotor speed:

$$Q = \frac{\partial P_r}{\partial \omega_r}, \quad \hat{Q} = \frac{\partial \hat{P}_r}{\partial \omega_r}. \tag{2.41}$$

Using the chain rule and the aerodynamic power expression, this derivative can be written as:

$$Q = \frac{1}{2}\rho A R U^2 \frac{\partial C_{\mathrm{P}}}{\partial \lambda}, \quad \hat{Q} = \frac{1}{2}\rho A R \hat{U}^2 \frac{\partial \hat{C}_P}{\partial \hat{\lambda}}. \tag{2.42}$$

The difference term driving the gain of $H(s)$ becomes:

$$Q - \hat{Q} = \frac{1}{2}\rho A R \left(U^2 \frac{\partial C_{\mathrm{P}}}{\partial \lambda} - \hat{U}^2 \frac{\partial \hat{C}_P}{\partial \hat{\lambda}} \right). \tag{2.43}$$

Thus, the transfer function $H(s)$ satisfies:

$$H(s) \to 0 \quad \text{if and only if} \quad U^2 \frac{\partial C_{\mathrm{P}}}{\partial \lambda} = \hat{U}^2 \frac{\partial \hat{C}_P}{\partial \hat{\lambda}}. \tag{2.44}$$

This convex condition forms the basis of the learning problem, where the goal is to minimize the mismatch between real and estimated aerodynamic model parameters by solving the following optimization problem:

$$\arg \min_{\hat{C}_P, \frac{\partial \hat{C}_P}{\partial \hat{\lambda}}} \left\| \frac{1}{C_{\mathrm{P}}^{2/3}(\lambda)} \frac{\partial C_{\mathrm{P}}}{\partial \lambda} - \frac{1}{\hat{C}_P^{2/3}(\hat{\lambda})} \frac{\partial \hat{C}_P}{\partial \hat{\lambda}} \right\|. \tag{2.45}$$

In this formulation, the first term reflects real-world aerodynamic behavior, and the second term contains the estimated and modeled values.

Internal Model Uncertainty
To account for blade degradation, the aerodynamic rotor properties are modeled with an internal uncertainty [41]. The power coefficient $C_{\mathrm{P}}(\lambda)$ is expressed as:

$$C_{\mathrm{P}}(\lambda) = \Gamma(\lambda) {C_{\mathrm{P}}}^{\diamond}(\lambda), \tag{2.46}$$

where $\Gamma(\lambda)$ is an unknown degradation function, and $C_{\mathrm{P}}^{\diamond}(\lambda)$ represents the ideal aerodynamic rotor properties.

Similarly, the power coefficient used within the control system is given by:

$$\hat{C}_P(\hat{\lambda}) = \Gamma(\hat{\lambda}) C_{\mathrm{P}}^{\diamond}(\hat{\lambda}). \tag{2.47}$$

By equating the steady-state aerodynamic rotor power and its estimate ($P_r = \hat{P}_r$), substituting the uncertainty expressions (2.46) and (2.47), and rearranging, we obtain:

$$\hat{U}^3 = U^3 \frac{C_{\mathrm{P}}(\lambda)}{\hat{C}_P(\hat{\lambda})} = U^3 \frac{\Gamma(\lambda)}{\Gamma(\hat{\lambda})} \frac{C_{\mathrm{P}}^{\diamond}(\lambda)}{C_{\mathrm{P}}^{\diamond}(\hat{\lambda})}. \tag{2.48}$$

An accurate estimate of the actual rotor-effective wind speed requires $C_{\mathrm{P}}(\lambda) = \hat{C}_P(\hat{\lambda})$, implying consistency between $\Gamma(\lambda)$ and its estimate $\hat{\Gamma}(\hat{\lambda})$. The uncertainty calibration problem is therefore formulated as:

$$|\Gamma(\lambda) - \hat{\Gamma}(\hat{\lambda})| < \epsilon \quad \forall (\lambda = \hat{\lambda}) \subseteq \Lambda, \tag{2.49}$$

where Λ is the complete turbine operating range, and ϵ is a predetermined threshold.

Ill-Conditioning of the Learning Problem
Despite the structured form of the learning objective, the estimation problem is ill-conditioned. This issue arises from a lack of identifiability of the individual values of the aerodynamic power coefficient C_{P} and the wind speed U, since only their product can be identified.

The ill-conditioning of the WSE–TSR control scheme, as explained in [43], can be understood by examining the steady-state behavior of the system. In steady-state and under the WSE–TSR control scheme, the actual aerodynamic rotor power equals the internal model estimate:

$$P_r = \hat{P}_r \quad \Rightarrow \quad \frac{1}{2}\rho A U^3 C_{\mathrm{P}}(\lambda) = \frac{1}{2}\rho A \hat{U}^3 \hat{C}_P(\hat{\lambda}). \tag{2.50}$$

Canceling constant terms yields:

$$U^3 C_{\mathrm{P}}(\lambda) = \hat{U}^3 \hat{C}_P(\hat{\lambda}). \tag{2.51}$$

This equation implies that multiple combinations of U and C_{P} can result in the same power measurement. Therefore, if the internal model misrepresents C_{P}, the estimator compensates by converging to a biased wind speed $\hat{U}$ that still satisfies the power balance. This phenomenon is what defines the ill-conditioning of the WSE–TSR scheme.

To further illustrate this, Brandetti et al. [43] define the internal model uncertainty using a constant scaling factor γ, such that:

$$\hat{C}_P(\hat{\lambda}) = \gamma C_{\mathrm{P}}(\lambda), \quad \hat{U} = \frac{U}{\gamma^{1/3}}. \tag{2.52}$$

Substituting into the power balance confirms:

$$\hat{U}\hat{C}_P(\hat{\lambda}) = \left(\frac{U}{\gamma^{1/3}}\right)^3 \gamma C_{\mathrm{P}}(\lambda) = U^3 C_{\mathrm{P}}(\lambda), \tag{2.53}$$

which holds true regardless of the value of γ.

This formulation in Equation (2.51) reduces the problem to estimating a single degradation parameter, which can be tuned to minimize model mismatch during excitation-based learning experiments, while avoiding the unidentifiable separation of wind speed and aerodynamic losses.

Learning Limitations at the Optimal Power Coefficient

A critical limitation emerges when the turbine operates at the tip-speed ratio that maximizes the power coefficient, λ^*. At this point, the gradient of the power coefficient with respect to tip-speed ratio vanishes:

$$\frac{\partial \hat{C}_P}{\partial \hat{\lambda}}(\lambda^*) = 0, \tag{2.54}$$

which causes the learning algorithm to lose its sensitivity to model errors. This effect can be seen by examining the modified relation derived from substituting the model into the learning objective:

$$U^2 \left(\frac{\partial \Gamma}{\partial \lambda} C_P^\diamond(\lambda) + \Gamma(\lambda) \frac{\partial C_P^\diamond}{\partial \lambda}(\lambda) \right) = \hat{U}^2 \left(\frac{\partial \hat{\Gamma}}{\partial \lambda} C_P^\diamond(\lambda) + \hat{\Gamma}(\lambda) \frac{\partial C_P^\diamond}{\partial \lambda}(\lambda) \right), \tag{2.55}$$

where Γ and $\hat{\Gamma}$ represent the degradation profiles, and $C_P^\diamond$ is a normalized power coefficient term. Assuming constant degradation (i.e., $\Gamma = C$), the expression simplifies and under the flat-gradient condition becomes:

$$U^2 \Gamma(\lambda) \frac{\partial C_P^\diamond}{\partial \lambda}(\lambda) = 0, \tag{2.56}$$

which nullifies the learning signal and makes it impossible to identify any correction to $\hat{\Gamma}$. In other words, the learning mechanism stalls. To overcome this issue, the control strategy is temporarily modified during the learning phase by shifting the tip-speed ratio reference away from the maximum efficiency point λ^* to a nearby setpoint $\tilde{\lambda}^*$ where the gradient is non-zero:

$$\frac{\partial C_P^\diamond}{\partial \lambda}(\tilde{\lambda}^*) \neq 0. \tag{2.57}$$

This ensures that the difference term in the learning condition remains active:

$$U^2 \Gamma(\lambda) \frac{\partial C_P^\diamond}{\partial \lambda}(\lambda) = \hat{U}^2 \hat{\Gamma}(\tilde{\lambda}^*) \frac{\partial C_P^\diamond}{\partial \lambda}(\tilde{\lambda}^*). \tag{2.58}$$

By intentionally operating the turbine at a suboptimal setpoint $\tilde{\lambda}^*$, the learning algorithm regains observability of the degradation profile and can effectively calibrate the internal model. Once learning converges, the controller may revert to tracking λ^*.

3

Methodology

This chapter presents the methodology used to simulate and analyze aerodynamic degradation caused by leading-edge erosion (LEE) in wind turbine blades. The goal is to quantify its impact on aerodynamic performance, model its effect on the power coefficient curve, and develop an uncertainty model suitable for integration with the WSE–TSR control framework.

The approach begins with the simulation of LEE using the AeroDyn module within OpenFAST by modifying lift and drag coefficients. The resulting changes in the power coefficient (C_{P}) are then modeled using two different frameworks: a scalar-based degradation function and a physically motivated decoupled model that separates lift and drag effects. An excitation-based learning algorithm is subsequently introduced, enabling real-time identification of degradation severity. Finally, the integration of this degradation model into the WSE–TSR control scheme is described.

Other degradation mechanisms (e.g., twist deformation, material loss, and yaw misalignment) were also explored in preliminary simulations but fall outside the primary focus of this thesis. For completeness, a summary of those additional cases and their modeling is provided in Appendix A.

3.1. Leading Edge Erosion

In this work, the effects of leading-edge erosion (LEE) are simulated using the AeroDyn module in OpenFAST. The NREL 5-MW offshore baseline wind turbine is used as the reference model. To isolate the aerodynamic effects of LEE, simulations are performed in steady-state using the Blade Element Momentum Theory (BEMT) model.

To analyze the impact of erosion on turbine efficiency, the power coefficient C_{P} is computed as a function of tip-speed ratio (TSR), which is varied by adjusting the wind speed from 8 m/s to 18 m/s. The rotor speed is fixed at the rated value of 12.1 rpm, and the blade pitch angle is held constant at 0°, consistent with operation in the partial-load region. This assumption is consistent with prior work, such as [30], which also restricts analysis to the partial-load region. Under these conditions, the power coefficient C_{P} is purely a function of the tip-speed ratio, and optimal rotor speed tracking becomes the primary control objective. This makes the C_{P}–λ curve directly representative of aerodynamic efficiency, and thus highly sensitive to degradation effects such as erosion or surface roughness.

The aerodynamic properties of each blade section are defined using airfoil data files, which include lift and drag coefficients (C_L, C_D) as functions of angle of attack. The blade is discretized into 19 spanwise nodes, and 8 unique airfoil files are assigned along the span from root to tip.

To simulate LEE, these aerodynamic coefficient files are modified by scaling the C_L and C_D curves in the outer 20 m of the blade span. These changes are informed by experimental studies, specifically the work of Kruse et al. [39], as detailed in Section 2.2.1. Moderate and severe erosion cases are constructed by applying multiplicative scalings to the clean airfoil data. No changes are made to blade geometry (e.g., chord length or twist), allowing the isolated aerodynamic effects of LEE to be studied directly. Findings in [39], summarized in Figure 2.6, showed that:

- P400 (sandpaper representing moderate erosion): this data can be represented by reducing C_L by 10% and increasing C_D by 50%.
- P40 (sandpaper representing severe erosion): this data can be represented by reducing C_L by 10% and increasing C_D by 100%.

Based on these observations, two degradation cases were constructed for simulation in AeroDyn:

- Baseline (Clean): Original C_L and C_D from the NREL airfoil database.
- Moderate Degradation: C_L reduced by 10%, C_D increased by 50%.
- Severe Degradation: C_L reduced by 30%, C_D increased by 100%.

The moderate degradation case is motivated directly by the experimental study of Kruse et al. [39]. Their findings indicated a 10% reduction in lift and a 50% increase in drag, making this a representative case for moderate levels of leading-edge erosion. The severe degradation case intentionally goes beyond these experimental observations. While Kruse et al. found that both moderate (P400) and severe (P40) surface roughness led to a 10% reduction in lift, this multiplication factor in C_L degradation is not necessarily an upper bound. No conclusive evidence shows that more extreme erosion causes constant lift degradation. Therefore, to ensure that the proposed modeling framework can accommodate a wider range of degradation severities the severe case includes a larger reduction in C_L (30%) and a doubling of C_D.

Figure 3.1 compares the modified aerodynamic coefficients with those from experimental sandpaper cases. The alignment between the moderate degradation and P400 case validates the use of scaling approximations for simulating erosion effects in AeroDyn.

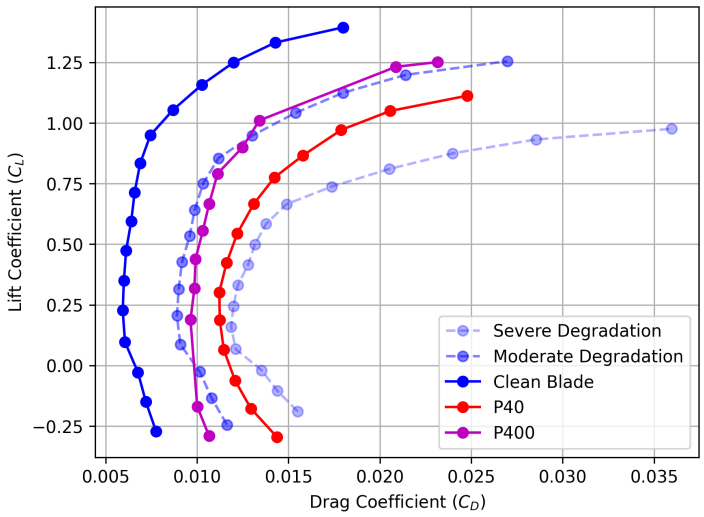

Figure 3.1: Comparison of lift (C_L) and drag (C_D) coefficients for clean, experimentally degraded (P400 and P40), and simulated degradation (moderate and severe) airfoils. Moderate degradation closely matches the P400 case from [39], while the severe case extends beyond P40 conditions to represent more extreme erosion.

3.2. Internal Uncertainty Model

As discussed in Section 1.3.2, the internal uncertainty model used in the WSE–TSR control strategy currently represents the state of the art for modeling aerodynamic degradation in real-time control. This model is based on an estimated degradation function:

$$\hat{C}_P(\hat{\lambda}) \triangleq \Gamma(\hat{\alpha}, \hat{\lambda}) C_{\mathrm{P}}^{\diamond}(\hat{\lambda}). \tag{3.1}$$

Here, $C_{\mathrm{P}}^{\diamond}(\hat{\lambda})$ is the power coefficient of the clean (undegraded) turbine, and $\hat{C}_P(\hat{\lambda})$ is the degraded $C_{\mathrm{P}}-\lambda$. The degradation function $\Gamma(\hat{\alpha}, \hat{\lambda})$ is used to capture aerodynamic degradation and is parameterized as:

$$\Gamma(\hat{\alpha}, \hat{\lambda}) \triangleq \alpha\gamma(\hat{\lambda}), \tag{3.2}$$

with $\gamma(\hat{\lambda}) : \mathbb{R} \to \mathbb{R}$ representing the degradation profile and $\alpha \in \mathbb{R}^+$ being the magnitude scaling factor.

For the degradation scenarios described in the previous section, the power coefficient curves have been simulated using AeroDyn as a function of tip-speed ratio λ. The degradation function $\Gamma(\lambda)$ can then be derived directly using Equation (3.1) by comparing the degraded and clean curves.

This scalar-based degradation formulation, where all scenarios are modeled using a common degradation profile $\gamma(\lambda)$ scaled by a degradation severity parameter α, is based on existing approaches in the literature (e.g., [30]). It represents the current learning framework within the WSE-TSR framework. The assumption of a shared shape $\gamma(\lambda)$ simplifies the learning process but may not accurately capture the true shape-altering changes in the power coefficient curve that result from degradation scenarios. In this chapter, this method is described as it forms the foundation of the current control strategy; however, its limitations will be addressed, and suggestions for improvement will be made in subsequent sections of this thesis.

The model can be applied by following the steps outlined below:

Step 1: Constructing the Degradation Function
First, the degradation function $\Gamma(\lambda)$ is defined as the relative performance loss due to aerodynamic degradation. This is calculated by normalizing the degraded power coefficient curve with respect to the clean (baseline) design curve:

$$\Gamma(\lambda) = \frac{\hat{C}_{\mathrm{P}}(\lambda)}{C_{\mathrm{P}}^{\diamond}(\lambda)} \tag{3.3}$$

Here:

- $\hat{C}_{\mathrm{P}}(\lambda)$ is the power coefficient under a given degradation state,
- $C_{\mathrm{P}}^{\diamond}(\lambda)$ is the power coefficient for the clean (undegraded) design,
- $\Gamma(\lambda)$ is the degradation function describing the power coefficient change at each tip-speed ratio.

Step 2: General Model Formulation
It is assumed that the degradation function $\Gamma(\lambda)$ can be modeled as a product of a scalar amplitude α and a common degradation profile $\gamma(\lambda)$:

$$\Gamma(\lambda) = \alpha\gamma(\lambda). \tag{3.4}$$

Here:

- $\gamma(\lambda)$ is a polynomial of degree n capturing the shape of the degradation across the tip-speed ratios,
- $\alpha \in \mathbb{R}^+$ is a positive scalar that reflects the severity of the degradation.

The degradation profile $\gamma(\lambda)$ is expressed as:

$$\gamma(\lambda) = \sum_{i=0}^{n} c_i \lambda^i, \tag{3.5}$$

where the coefficients $\mathbf{c} = \{c_0, c_1, \ldots, c_n\} \in \mathbb{R}^{n+1}$ are to be identified through curve fitting.

The choice of polynomial degree n is guided by the error propagation analysis presented in the intermezzo box below. There, it is shown that a fitting error in $\Gamma(\lambda)$ directly affects the estimated power coefficient $\hat{C}_P$. To ensure that this error remains within a predefined tolerance, the degree n is increased until the maximum fitting error in $\Gamma(\lambda)$ falls below the corresponding limit. This provides a principled way to select the minimal polynomial order needed for acceptable accuracy.

Intermezzo: Error for the Degradation Function

An error in the degradation function $\Gamma(\lambda)$ will directly translate into an error in the estimated power coefficient $\hat{C}_P(\lambda)$. To understand the impact of such model approximation errors, an error propagation analysis is performed. The model relation is given by:

$$\hat{C}_P(\lambda) = (1 - \Gamma(\lambda))\, C_P^{\diamond}(\lambda), \tag{3.6}$$

so a deviation $\Delta\Gamma$ introduces an error in the estimate of:

$$\Delta\hat{C}_P = -\Delta\Gamma \cdot C_P^{\diamond}. \tag{3.7}$$

To bound this propagated error, a predefined tolerance of $\Delta\hat{C}_P \leq 0.02$ is chosen. This threshold corresponds to a relative error of roughly 3.4% in the predicted power coefficient.
To obtain an estimate, the worst-case scenario where $C_P^{\diamond}$ is at its maximum is considered. The peak value of $C_P^{\diamond} = 0.59$ corresponds to the Betz limit, derived in Section 2.1.1.
Substituting into the propagation expression gives:

$$|\Delta\Gamma| \leq \frac{0.02}{0.59} \approx 0.034. \tag{3.8}$$

This value is used as the maximum permissible fitting error for the degradation function $\Gamma(\lambda)$ across the tip-speed ratio range. It serves as a quantitative guide when selecting the polynomial degree to approximate the degradation profiles without exceeding acceptable performance deviation.

Step 3: Fitting the Baseline Degradation Curve
To determine the polynomial coefficients, the model $\Gamma(\lambda) = \alpha\gamma(\lambda)$ is first fitted to one degradation dataset, the 'Severe Degradation' case. The fitting process minimizes the squared error between the observed degradation curve and the model:

$$\min_{\{c_i\}} \sum_{j=1}^{m} \left[\Gamma_{\text{Severe Degradation}}(\lambda_j) - \left(\sum_{i=0}^{n} c_i \lambda_j^i \right) \right]^2. \tag{3.9}$$

This step determines the polynomial shape $\gamma(\lambda)$ and sets $\alpha = 1$ for the severe degradation (reference) case.

Step 4: Reusing Polynomial Shape with Scalar Re-Fitting
Assuming the shape of the degradation profile $\gamma(\lambda)$ remains consistent across degradation levels, the coefficients $\{c_i\}$ obtained in Step 3 are fixed. Then, for another scenario (e.g., moderate degradation), the optimal scaling factor α is determined by solving:

$$\min_{\alpha} \sum_{j=1}^{m} \left[\Gamma_{\text{Moderate Degradation}}(\lambda_j) - (\alpha\gamma(\lambda_j)) \right]^2. \tag{3.10}$$

This approach provides a degradation model with a shared profile and a single scalar that captures the severity of degradation. The resulting α values offer an interpretable metric: lower values indicate greater aerodynamic degradation.

3.3. Proposed Degradation Framework

As introduced previously, the internal uncertainty model within the WSE-TSR control strategy aims to capture aerodynamic degradation in wind turbines by modeling deviations in the power coefficient curve.

In this thesis, a new degradation modeling framework is proposed that enhances physical interpretability by explicitly decoupling the aerodynamic effects of lift and drag. This new framework is developed as a contribution to existing work. By separating the contributions of lift decrease and drag increase, the model provides clearer physical insight into the origin and impact of degradation.

The model can be applied by following the steps outlined below:

Step 1: Constructing the Degradation Function Γ
The degradation function $\Gamma(\lambda)$ quantifies the relative reduction in aerodynamic performance due to degradation and is defined as:

$$\Gamma(\lambda) \triangleq \frac{C_{\mathrm{P}}^{\diamond}(\lambda) - \hat{C}_P(\lambda)}{C_{\mathrm{P}}^{\diamond}(\lambda)}, \tag{3.11}$$

where:

- $\hat{C}_P(\lambda)$: power coefficient of the degraded turbine,
- $C_{\mathrm{P}}^{\diamond}(\lambda)$: power coefficient of the clean (baseline) turbine,
- $\Gamma(\lambda) > 0$: indicates degradation, while $\Gamma(\lambda) = 0$ implies no degradation.

Rearranging Equation (3.11) gives the degraded power coefficient as:

$$\hat{C}_P(\lambda) = (1 - \Gamma(\lambda))\, C_{\mathrm{P}}^{\diamond}(\lambda). \tag{3.12}$$

Step 2: General Model Formulation – Decoupled Model
The proposed model decomposes the degradation function into two components, each associated with either lift or drag degradation:

$$\Gamma(\lambda) = k_1 \gamma_{C_L}(\lambda) + k_2 \gamma_{C_D}(\lambda), \tag{3.13}$$

where:

- $\gamma_{C_L}(\lambda)$: degradation profile due to lift reduction,
- $\gamma_{C_D}(\lambda)$: degradation profile due to drag increase,
- $k_1, k_2 \in \mathbb{R}^+$: sensitivity parameters quantifying the contribution of each aerodynamic effect.

This linear decomposition assumes that the aerodynamic performance loss due to degradation can be described as a superposition of two physically interpretable effects: lift reduction and drag increase. Although this assumption is introduced here as a modeling choice, it is not made arbitrarily. It is supported by a decoupling analysis performed in this work, which investigates how isolated changes in lift and drag individually affect the power coefficient curve. In this analysis, the lift coefficient was scaled by factors of 0.95, 0.90, 0.80, and 0.70, while keeping the drag coefficient fixed, to observe the effect of decreasing lift. Conversely, the drag coefficient was scaled by 1.1, 1.2, 1.5, and 2.0, while keeping the lift coefficient fixed, to isolate the effect of increasing drag. As shown in Section 4.3.1, simulations of isolated lift and drag degradation (introduced in Section 3.1) reveal a clear pattern: increasing drag primarily reduces the height of the C_P–λ curve, while reducing lift predominantly causes a rightward shift of the curve. This separation in impact enables the degradation function to be approximated as a superposition of the two independent contributions, one for lift and one for drag.

Step 3: Fitting the Baseline Degradation Curves $\gamma_{C_L}(\lambda)$ and $\gamma_{C_D}(\lambda)$
In this step, the individual degradation profiles corresponding to lift and drag degradation are identified. These profiles, $\gamma_{C_L}(\lambda)$ and $\gamma_{C_D}(\lambda)$, are modeled as polynomials of degree n, similar to the approach in the previous degradation framework.

Each profile is computed from a controlled degradation scenario by selectively altering the aerodynamic coefficients:

- For $\gamma_{C_L}(\lambda)$, the lift coefficient C_L was varied while keeping C_D constant. Specifically, C_L was scaled by $\{0.95, 0.90, 0.80, 0.70\} \times C_L^{\diamond}$.
- For $\gamma_{C_D}(\lambda)$, the drag coefficient C_D was varied while keeping C_L constant. Specifically, C_D was scaled by $\{1.1, 1.2, 1.5, 2.0\} \times C_D^{\diamond}$.

These values were chosen to span a realistic range of aerodynamic degradation due to leading edge erosion.

Both profiles are expressed as:

$$\gamma_{C_L}(\lambda) = \sum_{i=0}^{n} c_i^{(L)} \lambda^i, \quad \gamma_{C_D}(\lambda) = \sum_{i=0}^{n} c_i^{(D)} \lambda^i, \tag{3.14}$$

where the sets of coefficients $\mathbf{c}^{(L)} = \{c_0^{(L)}, c_1^{(L)}, \ldots, c_n^{(L)}\}$ and $\mathbf{c}^{(D)} = \{c_0^{(D)}, c_1^{(D)}, \ldots, c_n^{(D)}\}$ are fitted independently. Polynomials are particularly suitable here because the C_P–λ curves and their degradation profiles are generally smooth and continuous over the operational tip-speed ratio range.

The polynomial degree n is selected according to the error propagation bounds discussed in the intermezzo box in the previous section. Specifically, n is increased until the maximum fitting error in the degradation function $\Gamma(\lambda)$ remains below the threshold necessary to limit the propagated error in $\hat{C}_P(\lambda)$ to less than 0.02.

The polynomial coefficients are obtained by minimizing the squared error between the observed degradation and the polynomial approximation. For example, γ_{C_L} is determined by solving:

$$\min_{\{c_i^{(L)}\}} \sum_{j=1}^{m} \left[\Gamma_{\text{Lift-only}}(\lambda_j) - \left(\sum_{i=0}^{n} c_i^{(L)} \lambda_j^i \right) \right]^2 . \tag{3.15}$$

A similar optimization is used to fit γ_{C_D} based on the drag-only degradation scenario.

Step 4: Reusing Polynomail Shape with k_1- and k_2-Fitting
Once the polynomial shapes $\gamma_{C_L}(\lambda)$ and $\gamma_{C_D}(\lambda)$ are known, they can be reused to fit any observed degradation state by solving for optimal values of k_1 and k_2. This is done by minimizing the residual error between the observed degradation function $\Gamma_{\text{obs}}(\lambda)$ and the model:

$$\min_{k_1, k_2} \sum_{j=1}^{m} \left[\Gamma_{\text{obs}}(\lambda_j) - (k_1 \gamma_{C_L}(\lambda_j) + k_2 \gamma_{C_D}(\lambda_j)) \right]^2 . \tag{3.16}$$

This least-squares problem can be solved using standard optimization tools. Once k_1 and k_2 are identified, they provide an interpretation of how much each physical effect (lift or drag degradation) contributes to the overall performance loss.

This decoupled framework offers several key benefits. The first is that for undegraded states, the model parameters are $k_1 = 0$ and $k_2 = 0$. With increasing values, there is an indication of increasing degradation. Furthermore, the parameters k_1 and k_2 directly indicate the sensitivity of the turbine performance to lift and drag changes. Additionally, degradation mechanisms can be integrated into the model by extending the linear combination with new profiles and corresponding scaling parameters.

3.4. Sensitivity Parameters Analysis

To better understand how different combinations of aerodynamic degradation influence turbine performance, a range of degraded blade conditions is simulated. These conditions are based on the controlled degradation scenarios defined earlier in Step 3 of the proposed framework (Section 3.3), where lift and

drag coefficients are independently scaled to isolate their impact on the C_P–λ curve. For each pair of (C_L, C_D) scalars, a new power coefficient curve $C_{\mathrm{P}}(\lambda)$ is generated using AeroDyn simulations.

Each resulting $C_{\mathrm{P}}(\lambda)$ curve is then fitted using the decoupled degradation model described in Section 3.3, which expresses the degradation as:

$$\Gamma(\lambda) = k_1 \gamma_{C_L}(\lambda) + k_2 \gamma_{C_D}(\lambda), \tag{3.17}$$

with the corresponding degraded power coefficient:

$$\hat{C}_P(\lambda) = (1 - \Gamma(\lambda)) C_{\mathrm{P}}^{\diamond}(\lambda). \tag{3.18}$$

This fitting yields a unique pair of (k_1, k_2) values for each degraded (C_L, C_D) combination. As a result, the entire degradation space, originally parameterized in terms of lift and drag scalars, can now be re-expressed in terms of the sensitivity parameters (k_1, k_2). These parameters indicate how much lift- and drag-related degradation contributes to performance loss.

Severity Metrics

To quantify and visualize the severity of each degradation case, two metrics are introduced:

1. **Efficiency drop at peak power coefficient:**

   The first is based on comparing the maximum value of the power coefficient curve between the clean and degraded blades:

   $$\Delta C_{\mathrm{P}}^{\max} = \frac{C_{P,\max}^{\diamond} - C_{P,\max}}{C_{P,\max}^{\diamond}}. \tag{3.19}$$

   This metric provides a direct measure of overall efficiency loss and is particularly useful when interested in performance at optimal operating conditions.

2. **Efficiency drop at fixed operating point:**

   The second metric uses the value of C_{P} at the tip-speed ratio λ^* where the clean turbine achieves its maximum efficiency:

   $$\Delta C_{\mathrm{P}}^{\lambda^*} = \frac{C_{\mathrm{P}}^{\diamond}(\lambda^*) - C_{\mathrm{P}}(\lambda^*)}{C_{\mathrm{P}}^{\diamond}(\lambda^*)}. \tag{3.20}$$

   This approach shows the comparison at a fixed operating point. This could be the case when the operating point is not adjusted when degradation has occurred.

By calculating these metrics for each point in the (k_1, k_2) parameter space, a comprehensive severity map can be constructed. This map not only ranks different degradation scenarios but also provides physical insight into how lift and drag degradation contribute to turbine efficiency loss.

3.4.1. Modeling of Degradation Paths

The degradation map provides a static overview of potential degraded states. However, real wind turbines experience gradual and continuous degradation [44]. To capture this evolution, the turbine's progression can be modeled as a path through the (k_1, k_2) parameter space over time. This path begins at the clean state $(0, 0)$ and moves toward higher values of k_1 and k_2, indicating increased aerodynamic degradation.

If known degradation trajectories (e.g., from LEE cases such as P400 and P40) are available, they can be overlaid on the map. These trajectories serve as references. Statistical insights, such as the most probable path, variability across turbines, and bounds on performance loss, can be extracted. This approach turns the severity map into a dynamic framework, enabling better tracking, forecasting, and mitigation of turbine degradation over time. As more SCADA data, field observations, or structural diagnostics become available, more realistic and turbine-specific degradation paths can be identified and used to enhance the predictive accuracy of the model.

3.5. Degradation Parameter Learning

In order to implement the degradation framework in a control framework, the Wind Speed Estimation–Tip-Speed Ratio (WSE-TSR) control scheme is used. In the control scheme, it is possible to assign a degraded power coefficient curve to the wind turbine and implement the degradation framework in the estimator. However, the framework assumes a zero initial value for the degradation parameter α (in the old degradation framework described in Section 3.2) or for k_1 and k_2 (in the proposed degradation framework described in Section 3.3). This necessitates an adaptive learning mechanism that continuously refines the estimator based on closed-loop observations. In the following section, the algorithm for this is described.

3.5.1. Learning Algorithm

An excitation-based learning algorithm is proposed to calibrate the internal model of the wind turbine controller. The key idea is to excite the system and observe the transfer function magnitude response, which diminishes as the internal model becomes more accurate. The analysis for this was established in Section 2.4.3. The approach enables real-time adaptation of the degradation parameter by tracking how the system responds to the injected excitation. An overview of the algorithm is illustrated in Figure 3.2. The figure shows how the excitation signal is handled and how the resulting error is processed to update the model parameter. The remainder of this section describes each block of the diagram in more detail.

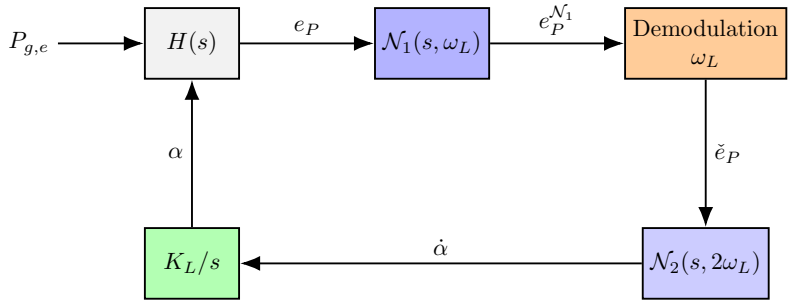

Figure 3.2: Block diagram of the excitation-based learning algorithm. A sinusoidal excitation is added to the generator power, and the resulting power error $e_P(t)$ is filtered, demodulated, and used to update the internal degradation parameter α.

A sinusoidal signal is introduced to the generator power, P_g. This is done to excite the closed-loop system dynamics at a specific frequency:

$$P_{g,e}(t) = A_{exc} \sin(\omega_L t). \tag{3.21}$$

Here, A_{exc} is the amplitude and ω_L is the excitation frequency. This input allows the system to generate measurable responses with information about model accuracy.

This excitation signal propagates through the wind turbine system. As output the error between the power of the wind turbine and the estimator is chosen. The resulting signal is again a sinusoidal whose amplitude and phase contain information about model mismatch:

$$e_P(t) = A_e \sin(\omega_L t + \psi_H). \tag{3.22}$$

The amplitude A_e is directly proportional to the magnitude of the transfer function $|H(j\omega_L)|$, and hence holds information of the model mismatch.

To isolate the excitation frequency component from the noisy system output, an (inverse) notch filter centered at ω_L is applied:

$$N_1(s, \omega_L) = \frac{Ks}{s^2 + 2\zeta\omega_L s + \omega_L^2}. \tag{3.23}$$

This filter helps eliminate unwanted frequency components and retains only the part of the signal that corresponds to the excitation frequency. As a result, there is a filtered signal $e_P^{N_1}$, with frequency ω_L.

Next, the filtered signal is demodulated to shift its frequency content from AC to DC (zero frequency), enabling reliable extraction of magnitude information. This is essential because the learning mechanism relies on integrating the signal over time to update the degradation parameter. To achieve demodulation, the signal of the power error after the notch filter is multiplied with a carrier. The signal after the notch filter can be written as $e_P^{N_1}(t) = A_e \sin(\omega_L t + \psi_H)$. The demodulation carrier is defined as $c(t) = \sin(\omega_L t + \psi_D)$. Here, ψ_D is a phase-offset tuning variable to compensate for the phase loss ψ_{HN} caused by dynamic operations and system delays.

$$\begin{aligned}
\breve{e}_P(t) &= e_P^{N_1}(t)c(t) && (3.24)\\
&= A_e \sin(\omega_L t + \psi_H)\,[\sin(\omega_L t + \psi_D)] && (3.25)\\
&= \frac{A_e}{2}\,[\cos(\psi_H - \psi_D) - \cos(2\omega_L t + \psi_H + \psi_D) && (3.26)\\
&= \frac{A_e}{2}\,[\cos(\psi_H - \psi_D) - \cos(2\omega_L t + \psi_H + \psi_D)] && (3.27)\\
&= \underbrace{\frac{A_e}{2}\,[\cos(\psi_H - \psi_D)]}_{\text{DC (0 Hz) component}} + \underbrace{\text{oscillating terms at } 2\omega_L}_{\text{filtered out}} && (3.28)
\end{aligned}$$

By multiplying the signal with in-phase and quadrature components, the demodulation extracts a DC component that is proportional to the original amplitude A_e.

To isolate the DC component from higher frequency terms, another notch filter is applied:

$$N_2(s, 2\omega_L) = \frac{s^2 + 2\zeta_1(2\omega_L)s + (2\omega_L)^2}{s^2 + 2\zeta_2(2\omega_L)s + (2\omega_L)^2}. \tag{3.29}$$

This removes the $2\omega_L$ harmonic and results in a clean DC signal that represents the system's frequency response magnitude.

The final step is to update the model's degradation scaling factor α using the demodulated DC signal:

$$\dot{\alpha}(t) = K_L \breve{e}_P^{(\mathrm{DC})}(t). \tag{3.30}$$

This real-time, gradient-like update adjusts the degradation parameter α to reduce the model mismatch observed in the system response, thereby improving model accuracy over time. The complete sequence of operations used for this learning procedure is summarized in Algorithm 1, which outlines the steps for excitation injection, signal filtering, demodulation, and parameter update.

Algorithm 1 Active Learning of Degradation Scaling α

1: **Initialize:** $\alpha \leftarrow 1.0$
2: **while** α not converged **do**
3: Inject excitation: $P_{g,e}(t) = A_P \sin(\omega_L t)$
4: Measure output $e_P(t)$ from system
5: Apply notch filter $N_1(s, \omega_L)$ to isolate excitation frequency
6: Demodulate filtered signal to extract DC component
7: Apply $N_2(s, 2\omega_L)$ to remove high-frequency content
8: Update α using: $\alpha \leftarrow \alpha + K_L \breve{e}_P^{(\mathrm{DC})} dt$

3.5.2. Learning Under a Path Constraint

Although the proposed degradation model framework includes two parameters, k_1 and k_2, the WSE-TSR scheme provides only scalar feedback based on the power tracking error at a single excitation frequency. This makes the full estimation of (k_1, k_2) an underdetermined problem: multiple combinations of these parameters can produce the same change in $\hat{C}_P(\lambda)$.

To resolve this, the learning is restricted to a predefined degradation path in the (k_1, k_2) space. This approach assumes that the turbine is undergoing a specific type of degradation, in this case, leading edge erosion, for which a degradation trajectory can be defined. Such a degradation path is determined by simulating multiple erosion cases and computing the corresponding k_1 and k_2 values using the decoupled degradation model. For instance, the erosion states created using P400 and P40 sandpaper have been simulated, and their corresponding (k_1, k_2) values have been identified. The identified points then form a represetative trajectory alogn with the learning algorithm can evolve. The k_1 and k_2 values are parametrized as:

$$k_1(\alpha), \quad k_2(\alpha),$$

by a scalar $\alpha \in [0, 1]$, representing degradation severity. The internal model becomes:

$$\Gamma(\lambda; \alpha) = k_1(\alpha)\gamma_{C_L}(\lambda) + k_2(\alpha)\gamma_{C_D}(\lambda),$$

and the WSE-TSR learning algorithm updates α directly:

$$\dot{\alpha}(t) = K_L \breve{e}_P^{(\mathrm{DC})}(t),$$

where $\breve{e}_P^{(\mathrm{DC})}(t)$ is the demodulated and filtered power tracking error.

This parameterization reduces the problem from two variables (k_1, k_2) to a single scalar parameter α, which captures the progression of degradation along a known trajectory. As a result, the learning algorithm only needs to estimate this scalar. With this reduction, the WSE-TSR control scheme can effectively adapt the internal model using scalar feedback, enabling real-time tracking of degradation severity.

4

Results of Degradation Modeling Framework

This chapter presents the results regarding the performance and implementation of the degradation modeling frameworks described in Sections 3.2 and 3.3. First, simulated LEE scenarios are used to quantify performance degradation through changes in the power coefficient curve. Then, two modeling frameworks are applied to these cases. Finally, LEE degradation is projected into a sensitivity parameter space defined by (k_1, k_2), where model severity and real-world cases (P400 and P40) are interpreted.

Additional degradation mechanisms, such as twist deformation, material loss, shaft tilt, and icing, were also simulated for comparison. These results are discussed in Appendix B.

4.1. Leading Edge Erosion Simulation Results

This section presents the simulation results for LEE using AeroDyn. The impact of surface erosion on aerodynamic performance is evaluated by analyzing the resulting power coefficient curves under clean, moderate, and severe degradation conditions.

The aerodynamic impact of LEE on wind turbine performance is illustrated by the resulting power coefficient curves under different degradation scenarios. As shown in Figure 4.1, three cases are simulated. The first contains the default values for the lift and drag coefficients. The second represents moderate degradation, simulated by decreasing the lift coefficient by 10% and increasing the drag coefficient by 50%. The last case represents severe degradation, the simulation has been done with a lift coefficient that is decreased by 30% and a drag coefficient that is increased by 100%. It can be seen that the leading edge erosion has a significant impact on the power coefficient and, with that also a significant impact on the aerodynamic performance of the wind turbine.

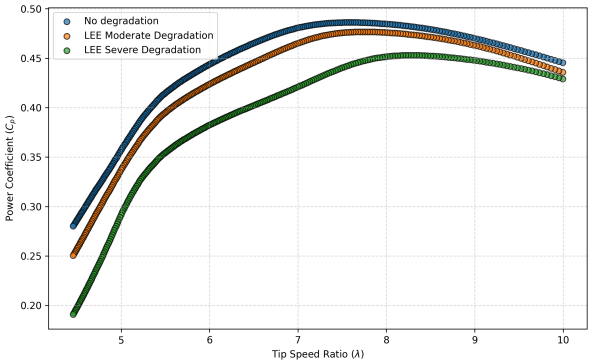

Figure 4.1: Simulated power coefficient (C_P) curves under leading edge erosion (LEE) conditions. The clean case uses default aerodynamic settings; moderate degradation involves a 10% reduction in C_L and 50% increase in C_D; severe degradation applies a 30% reduction in C_L and 100% increase in C_D. The results illustrate the aerodynamic efficiency loss due to surface erosion.

The wind turbine achieves its highest output when the power coefficient reaches its maximum value, denoted by $C_{P,\max}$. This peak occurs at a specific tip-speed ratio, λ_{opt}, which corresponds to the turbine's optimal operating point. Table 4.1 presents the values of $C_{P,\max}$ and λ_{opt} for each of the three degradation scenarios.

Table 4.1: Maximum power coefficient ($C_{P,\max}$) and corresponding optimal tip-speed ratio (λ_{opt}) for different LEE scenarios.

Scenario	$C_{P,\max}$	λ_{opt}
No Degradation	0.487	7.5
Moderate Degradation	0.478	7.8
Severe Degradation	0.456	8.4

The degradation results in a downward shift in the power coefficient, indicating a loss in aerodynamic efficiency. Additionally, the optimal tip speed ratio corresponding to the maximum power coefficient increases with the severity of erosion.

4.2. Internal Uncertainty Model

This section presents the results of applying the existing scalar degradation framework, introduced in Section 3.2, to the LEE degradation scenarios. This model captures aerodynamic degradation by defining a degradation function $\Gamma(\lambda)$, which quantifies the relative deviation of the degraded turbine's power coefficient from the clean reference case. Following the curve-fitting procedure described in Section 3.2, this degradation framework has been analyzed.

4.2.1. Degradation Functions

The degradation functions for the moderate and severe erosion cases are shown in Figure 4.2. These functions represent the normalized aerodynamic efficiency as a function of tip-speed ratio. A value of $\Gamma(\lambda) < 1$ indicates reduced performance compared to the undegraded (clean) turbine. The greater the deviation from 1.0, the more severe the aerodynamic degradation. The severe erosion case shows the most pronounced drop, emphasizing the performance impact of surface erosion.

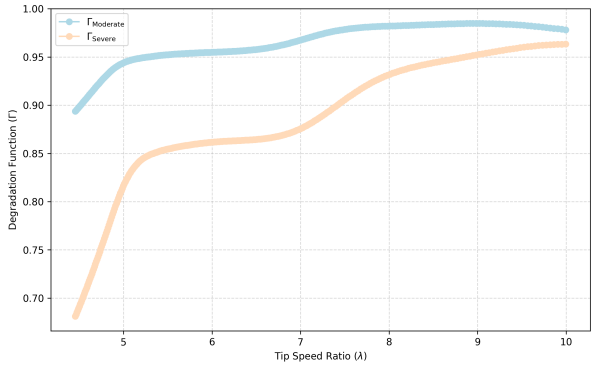

Figure 4.2: Normalized degradation functions $\Gamma(\lambda) = \hat{C}_P(\lambda)/C_{\text{P}}^{\diamond}(\lambda)$ for different levels of Leading Edge Erosion. These curves represent the relative aerodynamic performance loss across tip-speed ratios. Increased severity leads to a greater deviation below 1.0, indicating stronger degradation effects.

4.2.2. Fitting degradation functions

The methodology was applied to the degradation functions to extract a common degradation profile, denoted as $\gamma(\lambda)$. This profile was assumed to be shared across degradation states and scaled by a factor α. Two approaches were used:

- Fit a polynomial to the moderate degradation data as the baseline, then scale it to match the severe degradation case.

- Fit a polynomial to the severe degradation data as the baseline, then scale it to match the moderate degradation case.

A fifth-degree polynomial was found to sufficiently represent the baseline profiles in both cases. The corresponding coefficients are listed in Table 4.2.

Table 4.2: Fitted polynomial coefficients defining the degradation profile $\gamma(\lambda)$ for moderate and severe degradation baselines.

Coefficient	Moderate Degradation Baseline	Severe Degradation Baseline
c_0	9.5142	23.4613
c_1	-2.6729	-6.4221
c_2	0.3700	0.8662
c_3	-0.0252	-0.0575
c_4	0.00068	0.00150

Using the moderate degradation case as the baseline, the best fit to the severe degradation function was obtained by multiplying the polynomial $\gamma(\lambda)$ with $\alpha = 0.900$. This suggests that the severe degradation shares the same degradation shape but with a reduced magnitude. However, as shown in Figure 4.3a, the fit fails to capture the shape of the severe degradation function accurately.

Conversely, when the severe degradation case was used as the baseline, the best-fit scalar to match the moderate case was $\alpha = 1.107$. This implies that the moderate degradation profile is similar in shape but scaled up. In this case, the fit to the severe degradation data is good, while the match to the moderate degradation function is not accurate (see Figure 4.3b).

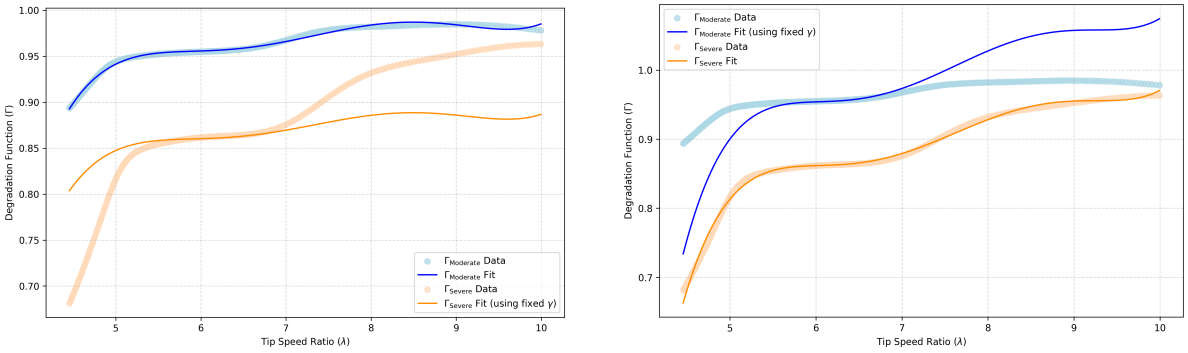

(a) Fitted degradation profile $\gamma(\lambda)$ using moderate degradation data as a baseline.

(b) Fitted degradation profile $\gamma(\lambda)$ using severe degradation data as a baseline.

Figure 4.3: Comparison of fitted degradation profiles using different baseline degradation levels.

4.2.3. Reconstruction of Power Coefficient Curve

To evaluate how well the scalar degradation model framework captures degradation, power coefficient curves were reconstructed using the estimated degradation function using:

$$\hat{C_{\mathrm{P}}}(\lambda) = \Gamma(\alpha, \lambda) C_{\mathrm{P}}^{\diamond} = \alpha\gamma(\lambda) C_{\mathrm{P}}^{\diamond}.$$

Figure 4.4 compares these reconstructed curves against the original simulated data for both the moderate and severe degradation cases. In both cases, the model captures a general vertical shift in the power coefficient curve. However, it fails to reproduce the horizontal shift in the tip-speed ratio at which the maximum power occurs. This mismatch highlights a key limitation of the scalar degradation framework: it assumes a fixed degradation shape $\gamma(\lambda)$ scaled by a single parameter α. While this approach simplifies model fitting and adaptation, it cannot fully represent more complex changes in the curve shape. This is particularly the case for the shift in the optimal operating point. As a result, relying solely on α for degradation modeling may be insufficient.

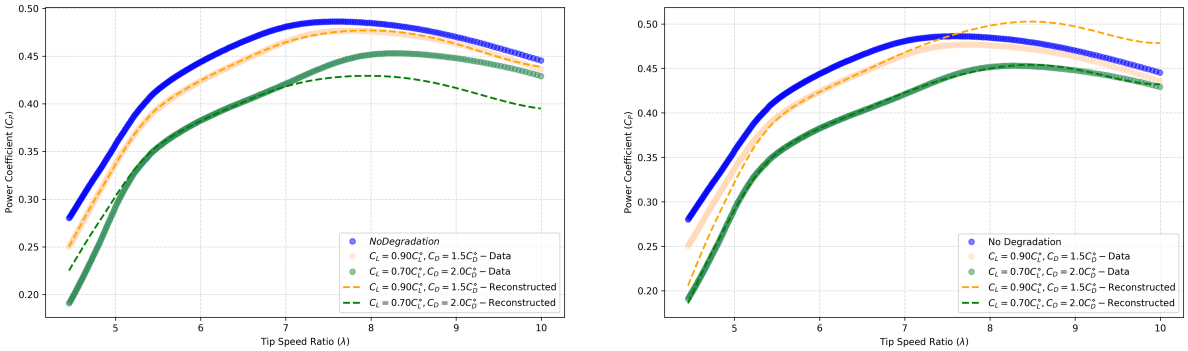

(a) Reconstructed vs. measured $C_P(\lambda)$ curve for moderate degradation.

(b) Reconstructed vs. measured $C_P(\lambda)$ curve for severe degradation.

Figure 4.4: Comparison of reconstructed and simulated $C_P(\lambda)$ curves under different degradation levels.

4.3. Proposed Internal Uncertainty Model

As demonstrated in the previous section, the scalar degradation model is able to capture vertical shifts in the power coefficient curve but fails to account for more complex changes, such as the horizontal shift in the optimal tip-speed ratio. To address this shortcoming, a new degradation modeling framework is proposed, as introduced in Section 3.3. Unlike the existing scalar model, this new proposed approach decomposes aerodynamic degradation into two physically interpretable components: lift and drag. By isolating these effects, the model provides a clearer understanding of how each aerodynamic parameter influences turbine performance. The analysis includes changes in the power coefficient curve, fitted degradation functions, and a reconstruction of the power coefficient curve using the model-based profiles.

4.3.1. Power Coefficient Change

To quantify the individual impact of lift and drag degradation on aerodynamic performance, a series of simulations were conducted in which either the lift coefficient (C_L) or the drag coefficient (C_D) was varied independently, while keeping the other fixed at its baseline (clean) value. Specifically, the lift coefficient was scaled by factors of $\{0.95, 0.90, 0.80, 0.70\} \times C_L^\diamond$, and the drag coefficient was scaled by $\{1.1, 1.2, 1.5, 2.0\} \times C_D^\diamond$. The resulting power coefficient curves are shown in Figures 4.5a and 4.5b.

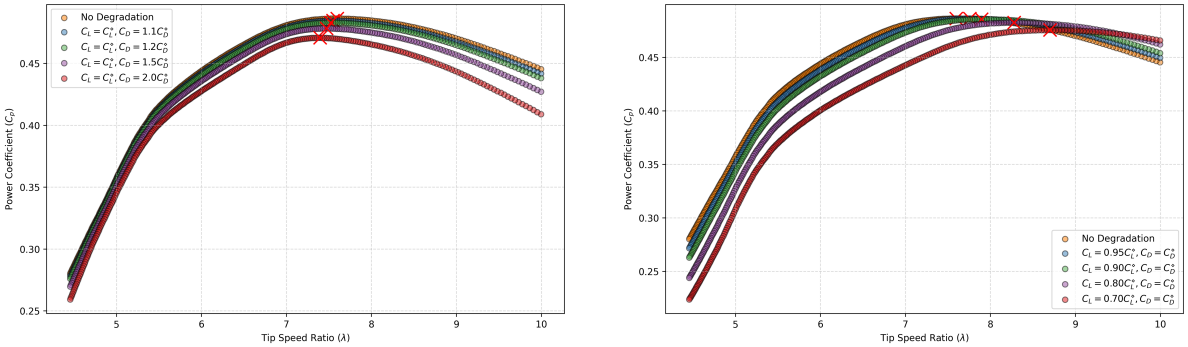

(a) Effect of increasing drag coefficient C_D on the power coefficient curve. The lift coefficient C_L is held constant at its clean reference value $C_L^\diamond$. Simulations were conducted for scaled drag values of 1.1, 1.2, 1.5, and 2.0 $\times$ $C_D^\diamond$. As drag increases, the power coefficient curve experiences a vertical drop, reflecting increased aerodynamic losses and reduced turbine efficiency.

(b) Effect of decreasing lift coefficient C_L on the power coefficient curve. The drag coefficient C_D is held constant at its clean reference value $C_D^\diamond$. Simulations were performed for scaled lift values of $0.95, 0.90, 0.80$, and $0.70 \times C_L^\diamond$. As lift decreases, the maximum power coefficient point shifts to higher tip-speed ratios, indicating a loss in aerodynamic effectiveness.

Figure 4.5: Impact of aerodynamic coefficient variations on the power coefficient curve. Changes in C_L and C_D independently influence $C_P(\lambda)$, providing insight into the sensitivity of power performance to aerodynamic degradation.

A clear distinction is observed in the degradation. Drag-induced degradation leads primarily to a vertical reduction in power coefficient. There is still a minor leftward shift in the tip speed ratio corresponding to the maximum value for the power coefficient. On the other hand, lift-induced degradation causes

a horizontal shift to the right, indicating that the turbine operates less efficiently and requires higher tip-speed ratios to achieve peak performance. Table 4.3 summarizes these effects, reporting the maximum power coefficient $C_{P,\max}$ and the corresponding optimal tip-speed ratio λ_{opt} for each degradation scenario.

Table 4.3: Maximum C_{P} and optimal λ values under independent variations of C_D and C_L.

Drag Coefficient Variation ($C_L = C_L^\diamond$)			**Lift Coefficient Variation** ($C_D = C_D^\diamond$)		
Condition	$C_{P,\max}$	λ_{opt}	**Condition**	$C_{P,\max}$	λ_{opt}
No Degradation	0.4862	7.599	No Degradation	0.4862	7.599
$C_D = 1.1C_D^\diamond$	0.4846	7.547	$C_L = 0.95C_L^\diamond$	0.4862	7.758
$C_D = 1.2C_D^\diamond$	0.4830	7.522	$C_L = 0.90C_L^\diamond$	0.4857	7.896
$C_D = 1.5C_D^\diamond$	0.4783	7.471	$C_L = 0.80C_L^\diamond$	0.4825	8.279
$C_D = 2.0C_D^\diamond$	0.4707	7.396	$C_L = 0.70C_L^\diamond$	0.4755	8.701

4.3.2. Degradation Functions

From the computed power coefficient curves, the degradation functions $\Gamma(\lambda)$ were calculated as the relative loss in performance, according to

$$\Gamma(\lambda) \triangleq \frac{C_{\text{P}}^\diamond(\lambda) - \hat{C}_P(\lambda)}{C_{\text{P}}^\diamond(\lambda)}.$$

These functions are shown in Figures 4.6a and 4.6b for drag and lift variations, respectively. As expected, the magnitude of $\Gamma(\lambda)$ increases with more severe degradation scenarios. That is, higher drag or lower lift leads to greater values of the degradation function, indicating larger relative aerodynamic performance losses. This trend confirms that $\Gamma(\lambda)$ behaves consistently with the modeled physical effects.

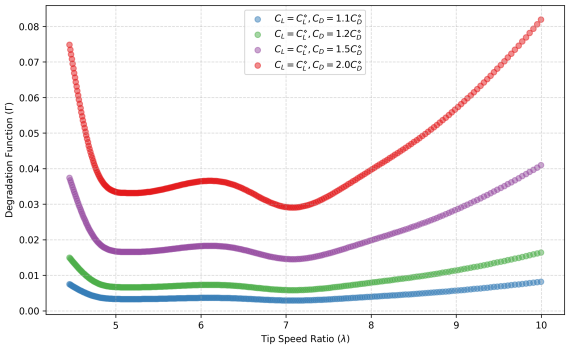

(a) Degradation functions $\Gamma(\lambda)$ resulting from isolated drag coefficient (C_D) increases. The degradation increases with higher C_D multipliers.

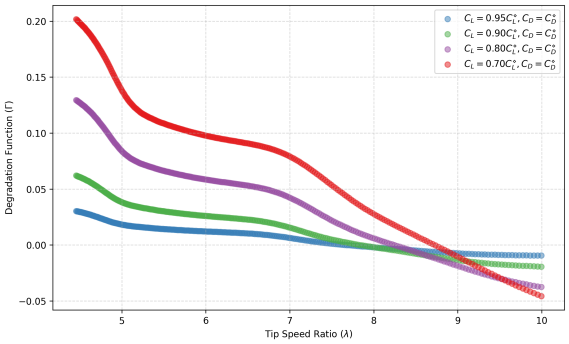

(b) Degradation functions $\Gamma(\lambda)$ resulting from isolated lift coefficient (C_L) reductions. The shape shifts significantly with increased reduction.

Figure 4.6: Degradation functions $\Gamma(\lambda)$ computed from changes in (a) drag and (b) lift coefficients. These represent the relative aerodynamic performance loss due to each isolated degradation.

4.3.3. Fitting Degradation Profiles

To construct the decoupled degradation model, polynomial profiles for lift- and drag-induced degradation were fitted to simulation data. The most extreme degradation scenarios were used to define the baseline polynomial shapes:

- For drag degradation, the profile $\gamma_{C_D}(\lambda)$ was obtained by fitting a polynomial to the degradation function from the case $C_D = 2.0\,C_D^\diamond$,
- For lift degradation, the profile $\gamma_{C_L}(\lambda)$ was fitted using the case $C_L = 0.7\,C_L^\diamond$.

These baseline degradation functions were modeled as polynomials of degree six, with the resulting coefficients summarized in Table 4.4. For all other degradation levels, these profiles were scaled by sensitivity parameters k_1 and k_2, enabling a compact and reusable formulation of the degradation function across varying severities.

Table 4.4: Fitted polynomial coefficients for degradation profiles from C_D and C_L change curves

Coefficient	C_D Variation	C_L Variation
c_0	47.2043	-8.2481
c_1	-40.7677	10.7410
c_2	14.4786	-4.9702
c_3	-2.7044	1.1297
c_4	0.2803	-0.1366
c_5	-0.0153	0.0084
c_6	0.00034	-0.00021

Figures 4.7a and 4.7b illustrate how the base profiles, when scaled with appropriate k_1 or k_2 values, align with the degradation functions obtained from other scenarios. The drag-induced degradation is captured with high accuracy across all levels. In the case of lift-induced degradation, the fit is slightly less accurate near the zero-crossing of the degradation function. The impact of this modeling error will be further evaluated in the following section, where reconstructed power coefficient curves are compared against the simulated data.

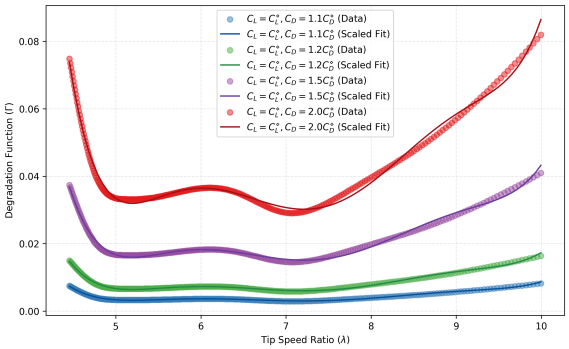

(a) Scaled degradation functions for drag coefficient variation. The base polynomial (from $C_D = 2.0\,C_D^\diamond$) is scaled using fitted sensitivity parameters k_2, showing excellent agreement across all levels of degradation.

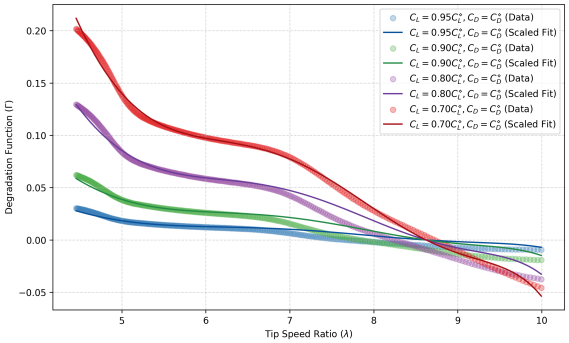

(b) Scaled degradation functions for lift coefficient variation. The base polynomial (from $C_L = 0.7\,C_L^\diamond$) is scaled using fitted sensitivity parameters k_1. The fit quality is slightly lower near the zero-crossing region but still captures the overall trend.

Figure 4.7: Fitted degradation functions for drag and lift variations using a base degradation profile scaled with sensitivity parameters. This supports the decoupled modeling approach where drag and lift degradation can be captured independently using fixed profiles and scalar multipliers.

The corresponding sensitivity parameters k_1 and k_2 used to scale the baseline profiles are listed in Table 4.5. These values quantify the severity of degradation relative to the baseline cases, where a value of 1.0 corresponds to the most severe simulated condition.

4.3.4. Reconstruct Decoupled Degradation

As introduced in Section 3.1, two representative degradation cases were simulated to model the effects of LEE: a moderate degradation scenario using $C_L = 0.9\,C_L^\diamond$ and $C_D = 1.5\,C_D^\diamond$, and a severe degradation case using $C_L = 0.7\,C_L^\diamond$ and $C_D = 2.0\,C_D^\diamond$. In this section, the corresponding power coefficient curves are reconstructed using the mathematical degradation model framework proposed. This involves combining the degradation profiles $\gamma_{C_L}(\lambda)$ and $\gamma_{C_D}(\lambda)$ with their associated sensitivity parameters k_1 and k_2, derived from the simulations for these specific lift and drag scaling.

Table 4.5: Fitted sensitivity parameters k_1 and k_2 for different levels of lift and drag degradation

Degradation Scenario	Parameter	Fitted Value
$C_L = 0.95\, C_L^{\diamond}$	k_1	0.1314
$C_L = 0.90\, C_L^{\diamond}$	k_1	0.2770
$C_L = 0.80\, C_L^{\diamond}$	k_1	0.6108
$C_L = 0.70\, C_L^{\diamond}$	k_1	1.0000
$C_D = 1.1\, C_D^{\diamond}$	k_2	0.1000
$C_D = 1.2\, C_D^{\diamond}$	k_2	0.2000
$C_D = 1.5\, C_D^{\diamond}$	k_2	0.5000
$C_D = 2.0\, C_D^{\diamond}$	k_2	1.0000

The degradation function is computed as:

$$\Gamma(\lambda) = k_1 \gamma_{C_L}(\lambda) + k_2 \gamma_{C_D}(\lambda), \tag{4.1}$$

and the reconstructed power coefficient curve is given by:

$$\hat{C}_P(\lambda) = (1 - \Gamma(\lambda))\, C_{\mathrm{P}}^{\diamond}(\lambda). \tag{4.2}$$

Figure 4.8a shows the combined degradation functions, and Figure 4.8b compares the reconstructed C_{P}–λ curves to the original simulated data. The model accurately captures both the downward shift in aerodynamic efficiency and the rightward shift in the optimal tip-speed ratio, validating the decoupled degradation model framework's ability to represent shape-altering effects in the $C_{\mathrm{P}}(\lambda)$ curve.

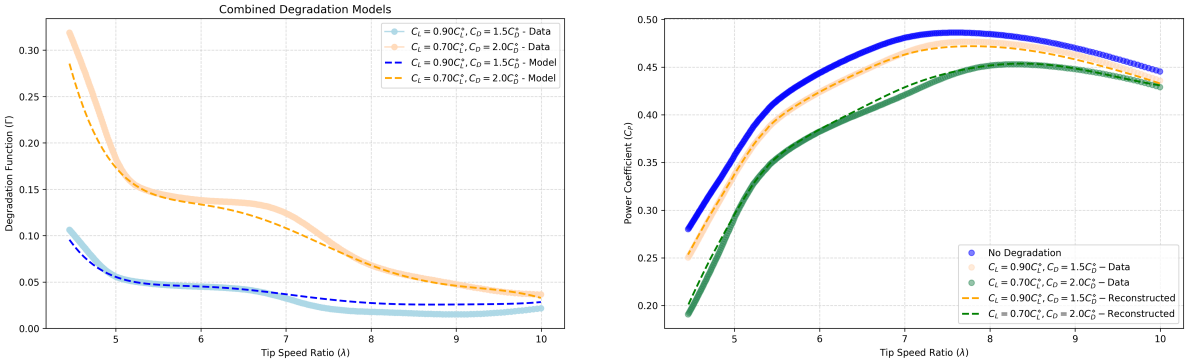

(a) Combined degradation function using decoupled γ_{C_D} and γ_{C_L} profiles.

(b) Reconstructed vs. measured $C_{\mathrm{P}}(\lambda)$ curves using decoupled internal model.

Figure 4.8: Evaluation of the decoupled internal model. The left plot shows the combined degradation function from separate γ_{C_D} and γ_{C_L} profiles, while the right plot compares reconstructed and simulated $C_{\mathrm{P}}(\lambda)$ curves.

4.4. Results Sensitivity Parameters Analysis

As discussed in Section 3.4, the sensitivity parameter space (k_1, k_2) provides an interpretable representation of aerodynamic degradation, separating the influence of lift and drag losses. To assess the impact of different combinations of degradation, two scalar performance metrics were introduced:

- $\Delta C_{\mathrm{P}}^{\lambda^*}$ metric – the relative power coefficient loss at the clean turbine's optimal tip-speed ratio λ^*,
- $\Delta C_{\mathrm{P}}^{\max}$ metric – the relative loss of the global maximum value of C_{P}.

These metrics were evaluated over a grid of (k_1, k_2) values, each corresponding to a specific degraded state as defined by the decoupled degradation model framework in Equations (4.1) and (4.2).

Figure 4.9a presents the $\Delta C_P^{\lambda^*}$ map, and Figure 4.9b shows the $\Delta C_P^{\max}$ map. Both plots reveal regions of equal performance loss, highlighted by iso-contours. This demonstrates that different combinations of k_1 and k_2 can yield the same aerodynamic efficiency loss.

This illustrates a key property of the decoupled model: power loss alone does not uniquely determine the (k_1, k_2) values. Without additional assumptions or prior knowledge, it is not possible to know the exact composition of lift and drag degradation from a scalar performance drop alone. To resolve this problem, physical degradation can be associated with characteristic paths in (k_1, k_2) space. These trajectories help constrain the estimation problem and will be further explored in the next section.

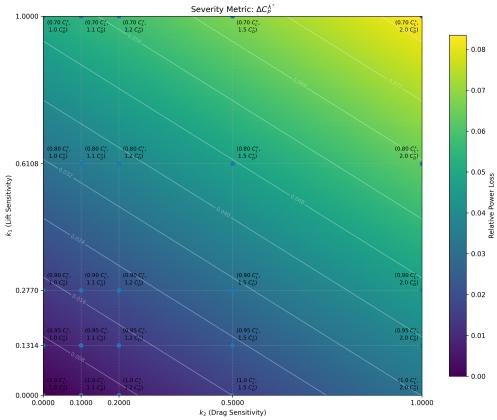

(a) Relative power coefficient loss $\Delta C_P^{\lambda^*}$ at the clean turbine's optimal tip-speed ratio λ^*. Each point in the (k_1, k_2) grid corresponds to a specific combination of lift and drag degradation. Contour lines highlight the non-uniqueness of degradation paths leading to the same performance drop.

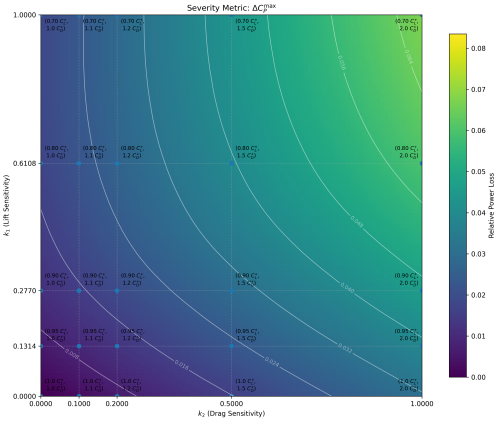

(b) Relative loss of maximum power coefficient $\Delta C_P^{\max}$ across the (k_1, k_2) space. The contours show that the same power loss can result from different aerodynamic degradation paths.

Figure 4.9: Maps of aerodynamic performance loss in (k_1, k_2) space using two scalar metrics: $\Delta C_P^{\lambda^*}$ and $\Delta C_P^{\max}$.

4.5. Real Degradation Cases

This section evaluates how well the proposed degradation model captures real-world aerodynamic degradation by applying it to a real LEE case. The degradation is projected onto the (k_1, k_2) sensitivity space to verify the model's ability to represent severity and progression in aerodynamic loss. Since the focus of this thesis is on LEE, the icing scenario is presented separately in Appendix B.2.

4.5.1. Leading Edge Erosion (LEE)

As introduced in Section 3.1, LEE is commonly modeled using controlled surface roughness such as P400 and P40 sandpaper. These cases were previously simulated with:

- P400: $C_L = 0.90\, C_L^{\diamond}, \quad C_D = 1.5\, C_D^{\diamond}$
- P40: $C_L = 0.90\, C_L^{\diamond}, \quad C_D = 2.0\, C_D^{\diamond}$

The corresponding degradation states were mapped to sensitivity parameters as:

$$\text{P400: } (k_1, k_2) = (0.277, 0.5), \quad \text{P40: } (k_1, k_2) = (0.277, 1.0).$$

These values were used to reconstruct the degradation functions and the corresponding $C_P - \lambda$ curves, shown in Figure 4.10a. The resulting curves align closely with the simulated curves, demonstrating that the decoupled degradation framework can correctly the curves for the P400 and P40 cases.

Figure 4.10b presents the corresponding error analysis. The maximum absolute error is small for both cases (P400: 0.00498 at TSR = 9.02, P40: 0.00542 at TSR = 4.54). The relative error remains below 0.02 across all points, with a slightly higher deviation observed for the more degraded P40 case.

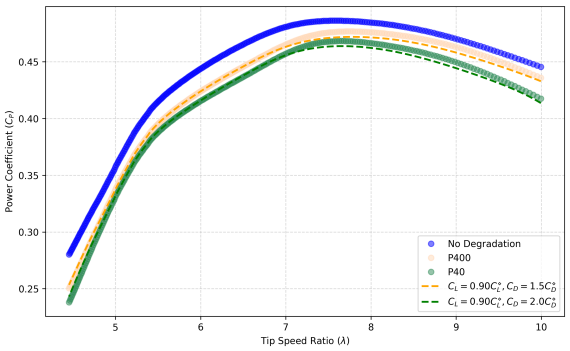

(a) Reconstructed vs. measured $C_P(\lambda)$ curves for P400 and P40 roughness levels.

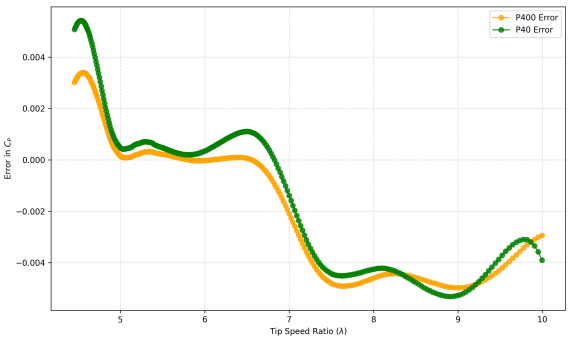

(b) Absolute reconstruction error for P400 and P40 cases.

Figure 4.10: Evaluation of the degradation model framework using LEE scenarios. Left: model-based reconstruction vs. simulation. Right: absolute and relative error profiles across λ.

4.5.2. Degradation Path in (k_1, k_2) Space

To visualize the progression of leading-edge erosion in terms of aerodynamic impact, the fitted (k_1, k_2) values corresponding to different surface roughness levels are plotted onto the degradation sensitivity maps. Figures 4.11a and 4.11b show the LEE trajectory overlaid on the (k_1, k_2) grid, highlighting its progression from clean to degraded states (e.g., P400 to P40).

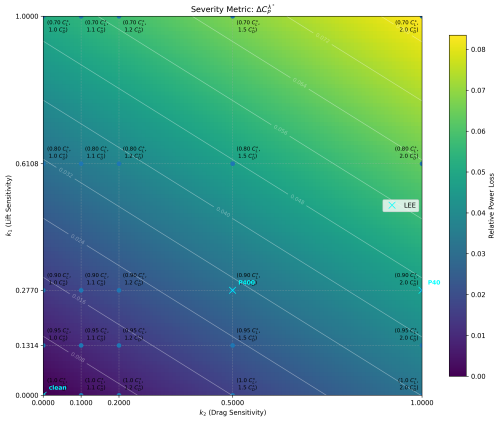

(a) Contour map of relative power loss at λ^* with the LEE trajectory overlaid. The path reflects a clear trend in increasing drag degradation (k_2) with erosion severity.

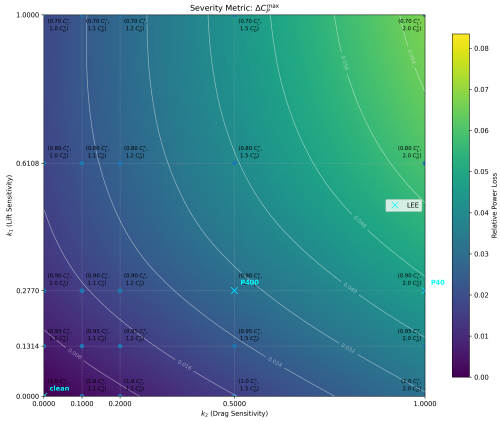

(b) Contour map of maximum relative C_P loss, showing how LEE maps to the (k_1, k_2) space.

Figure 4.11: Degradation path of leading-edge erosion (LEE) visualized on (k_1, k_2) sensitivity maps using two performance metrics: (a) $\Delta C_P^{\lambda^*}$ and (b) $\Delta C_P^{\max}$.

This analysis illustrates a key strength of the proposed modeling framework: degradation trajectories in the (k_1, k_2) space offer a compact and physically interpretable way to track turbine health. For LEE, the dominant direction of progression lies along the k_2 axis, reflecting a stronger contribution from drag degradation.

A complementary degradation path for icing is analyzed in Appendix B. While the scalar-fitting approach used there is less accurate for icing, it provides an illustrative comparison to LEE in terms of the trajectory.

5

Results of Degradation Model Framework Implementation in the WSE-TSR Control Scheme

This section presents the results of the implementation of the degradation model framework in WSE-TSR control scheme, focusing on the identification of optimal excitation parameters, baseline performance evaluation, and the learning behavior under simulated degradation scenarios. The analysis begins with the selection of the appropriate frequency for power excitation (P_g), followed by the tuning of damping parameters (ζ, ζ_1, and ζ_2). Subsequently, an evaluation with no degradation is performed to understand the undisturbed system response and extract the necessary phase offset for demodulation. After that, twin experiments are conducted to analyze the learning mechanism's ability to learn degradation parameters in controlled scenarios. Finally, the results of the learning parameter for the power coefficient curves for P40 and P400 from AeroDyn are presented.

Clarification on Degradation Scenarios Used
To avoid confusion, it is important to emphasize that two distinct types of degradation scenarios are used throughout this results chapter:

1. Simulation-Based Degradation: These scenarios use C_P–λ curves directly obtained from AeroDyn simulations for the P400 and P40 cases. These are based on realistic erosion effects (e.g., P400: $C_L = 0.9C_L^\diamond$, $C_D = 1.5C_D^\diamond$; P40: $C_L = 0.9C_L^\diamond$, $C_D = 2.0C_D^\diamond$). These curves are shown as data points in Figure 4.10a.

   The simulation-based power coefficient curves are used in Section 5.4, Section 5.5 (except Subsection 5.5.2), and Subsection 5.7.3.

2. Framework-Based Degradation: These scenarios use reconstructed C_P–λ curves generated by the internal degradation framework: with $\alpha = 0.5$ for the moderate (P400) case and $\alpha = 1.0$ for the severe (P40) case. These reconstructed curves are shown as dotted lines in Figure 4.10a.

   The framework-based power coefficient curves are used in -Subsection 5.5.2), Section 5.6 and Section 5.7 (except Subsection 5.7.3).

5.1. Learning Path Implementation in WSE-TSR Control Scheme

As discussed in Section 3.5.2, the WSE–TSR control scheme relies on scalar feedback, specifically a demodulated power error measured at a single excitation frequency. This setup presents an identifiability challenge: while the aerodynamic degradation model includes two parameters, k_1 and k_2, the available measurement data provides only a single scalar signal. As a result, estimating both k_1 and k_2 simultaneously from this limited information becomes an underdetermined problem.

To address this, the estimation is constrained to a predefined trajectory in the (k_1, k_2) parameter space. This trajectory reflects a physically plausible progression of LEE. By parameterizing this path using a scalar degradation severity parameter $\alpha \in [0, 1]$, the learning problem becomes well-posed: the two-parameter model is reduced to a single-parameter representation, which can be uniquely identified from scalar feedback.

The corresponding parameterization is expressed as:

$$k_1(\alpha), \quad k_2(\alpha) \quad \text{with} \quad \Gamma(\lambda; \alpha) = k_1(\alpha)\gamma_{C_L}(\lambda) + k_2(\alpha)\gamma_{C_D}(\lambda).$$

This formulation ensures that the WSE-TSR learning algorithm needs to estimate only a single scalar, α, which captures the severity of degradation along a known trajectory.

Because LEE is cumulative and irreversible, the degradation trajectory must be monotonic. That is, neither k_1 nor k_2 should decrease as degradation progresses. Physically, once aerodynamic degradation occurs, it cannot spontaneously reverse. This constraint is embedded into the assumed path, ensuring it reflects realistic turbine behavior over time.

The assumed degradation path is defined parametrically by:

$$k_2(\alpha) = \alpha, \quad k_1(\alpha) = \begin{cases} 0.277 \dfrac{1 - e^{-a\alpha}}{1 - e^{-a0.5}}, & \text{if } \alpha < 0.5 \\ 0.277, & \text{if } \alpha \geq 0.5 \end{cases}$$

where $a = 6$ is a shaping parameter that controls how rapidly the curve flattens. This ensures:

- $k_1(0) = 0$, $k_2(0) = 0$ (undegraded),
- $k_1(0.5) = 0.277$, $k_2(0.5) = 0.5$ (P400 erosion),
- $k_1(1.0) = 0.277$, $k_2(1.0) = 1.0$ (P40 erosion).

The degradation path is illustrated in Figure 5.1, overlaid on two severity maps computed in the (k_1, k_2) space. These maps are based on the severity metrics previously defined in Section 3.4.

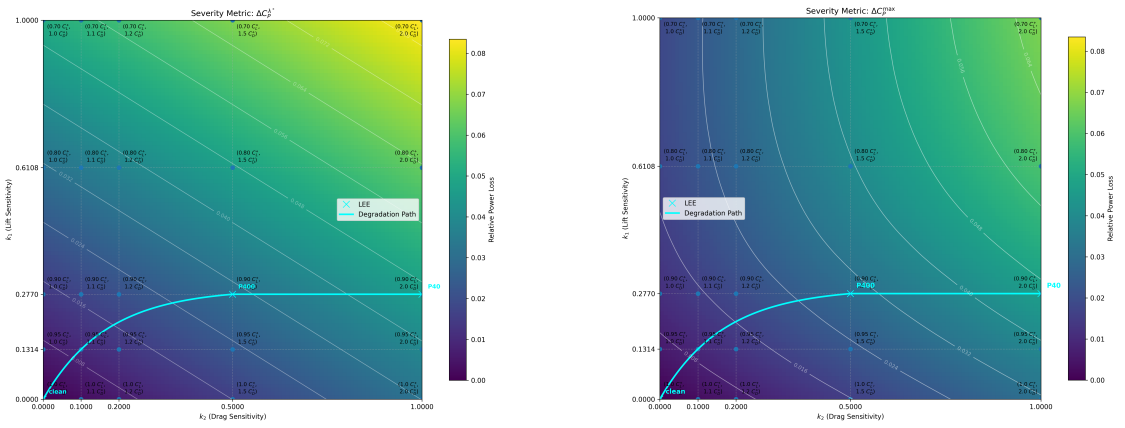

(a) Relative power loss computed at λ^* (optimal tip-speed ratio of the clean turbine).

(b) Maximum relative power loss across the $C_P(\lambda)$ curve.

Figure 5.1: Degradation severity maps in the (k_1, k_2) parameter space, evaluated using two metrics: (a) loss at fixed optimal tip-speed ratio, and (b) global maximum loss. The assumed degradation trajectory is overlaid, showing the evolution of leading edge erosion along the path from clean to severely degraded states.

In the results that follow, all references to learning the degradation parameter refer specifically to the scalar α that progresses along the constrained degradation path in (k_1, k_2)-space. By reducing the estimation problem to a single parameter and embedding prior knowledge of the degradation path, it is possible to implement the degradation model framework in the WSE-TSR control scheme.

5.2. Frequency Response

To identify the optimal frequency for excitation of the generator power (P_g), the system was subjected to a range of excitation frequencies. The amplitude of the output, represented by the power error, was measured to evaluate the system's sensitivity to each frequency.

The results indicate that the highest amplitude response for both P40 and P400 degradation scenarios occurred at a frequency of 0.0155 Hz. Specifically:

- P40 (Severe Degradation): Maximum amplitude of 36474.54 W at 0.0155 Hz.
- P400 (Moderate Degradation): Maximum amplitude of 30162.19 W at 0.0155 Hz.

Figure 5.2 shows the frequency response of the WSE-TSR system under the P40 degradation scenario. The plot shows the amplitude of E_p against excitation frequency, clearly illustrating a peak at 0.0155 Hz. Although this figure represents the P40 case, the frequency response for P400 is identical, with the same peak frequency observed.

Based on this analysis, the excitation frequency of 0.0155 Hz was selected as optimal for subsequent simulations, as it holds the maximum system response.

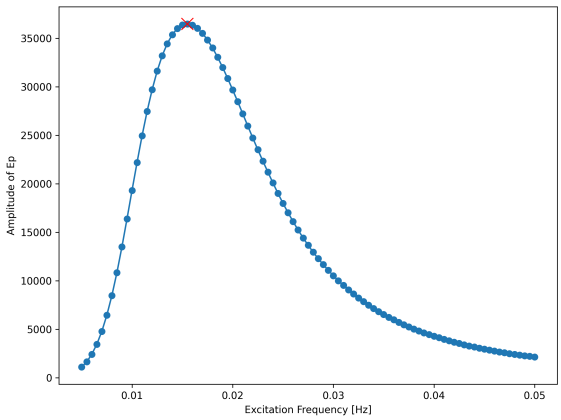

Figure 5.2: Frequency response of the WSE-TSR system for P40 degradation scenario. The highest response is observed at 0.0155, indicating the optimal excitation frequency for P_g.

5.3. Signal Filtering

The signal filtering process in the WSE–TSR control scheme involves two notch filters. As introduced in Section 3.5.1, the first notch filter isolates the excitation frequency ω_L, while the second filter suppresses the $2\omega_L$ harmonic that arises after demodulation. The effectiveness of both filters depends critically on the damping parameters ζ, ζ_1, and ζ_2, which shape their selectivity and convergence behavior.

Notch Filter 1: Isolating the Excitation Frequency

To isolate the response at the excitation frequency while filtering out all other frequency components, Notch Filter 1 is applied. As shown in Figure 5.3a, its magnitude response demonstrates how frequencies outside ω_L are attenuated. A lower damping ratio ζ sharpens this selectivity, creating a narrower passband.

However, increased frequency selectivity comes at the cost of slower convergence in the filtered signal. This trade-off is illustrated in Figure 5.3c, which shows the convergence behavior of the filtered signal over time for different values of ζ. At $\zeta = 0.01$, the system converges slowly due to filtered out information. In contrast, higher ζ values improve convergence. Additionally, Figure 5.3b shows that while the filter introduces some phase distortion, it remains minimal near the excitation frequency, ensuring that the phase of the signal is preserved.

To balance selectivity and responsiveness, a damping ratio of $\zeta = 0.05$ was selected for all subsequent experiments. This value provides a favorable trade-off between sufficient frequency filtering and convergence speed.

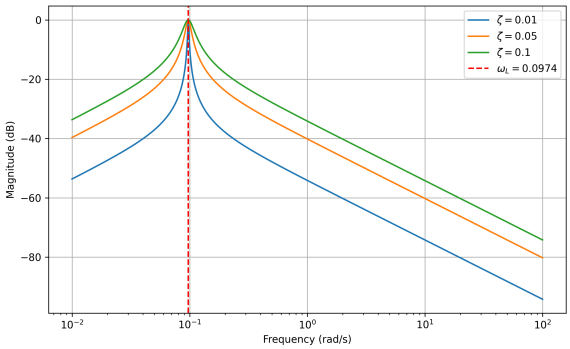

(a) Magnitude response of Notch Filter 1. Only the excitation frequency is kept, while surrounding frequencies are significantly attenuated. Lower ζ values increase attenuation strength but narrow the passband.

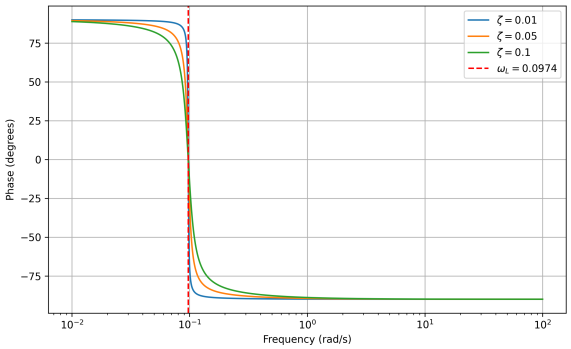

(b) Phase response of Notch Filter 1. Phase distortion remains minimal near the excitation frequency.

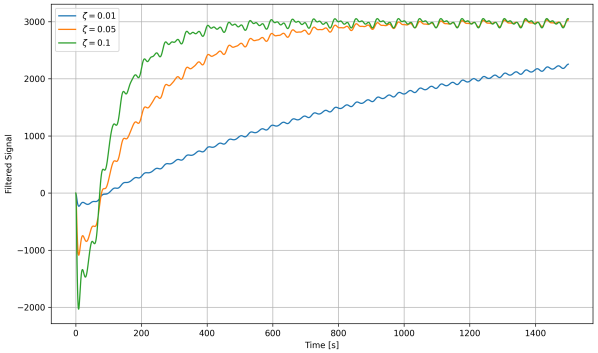

(c) Convergence behavior for various ζ values. $\zeta = 0.05$ offers a balance between filtering sharpness and speed.

Figure 5.3: Frequency and time-domain performance of Notch Filter 1 for different ζ values.

Notch Filter 2: Suppressing Harmonics After Demodulation

Following demodulation, a second notch filter is applied to remove the high-frequency component at $2\omega_L$. The effectiveness of the filtering depends on the selection of damping parameters ζ_1 and ζ_2. Lower values of ζ_1 result in sharper notching of the $2\omega_L$ frequency, while higher values of ζ_2 contribute to a broader suppression band. Through tuning, the parameters $\zeta_1 = 0.01$ and $\zeta_2 = 0.99$ were identified as the optimal combination. This setup ensures maximum attenuation of the unwanted component while preserving the remaining signal.

Figure 5.4c shows the time-domain impact of this configuration, clearly indicating that the $2\omega_L$ component is effectively removed. The filter's frequency-domain behavior is also shown in Figures 5.4a and 5.4b, where the notch depth and phase characteristics illustrate the frequency and phase behavior.

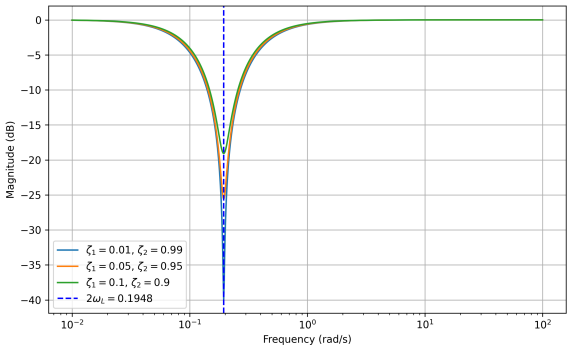

(a) Magnitude response of Notch Filter 2. Strong attenuation centered at $2\omega_L$.

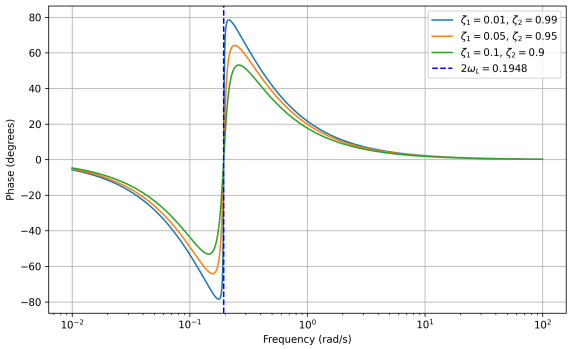

(b) Phase response of Notch Filter 2. Minimal phase distortion outside the notch region.

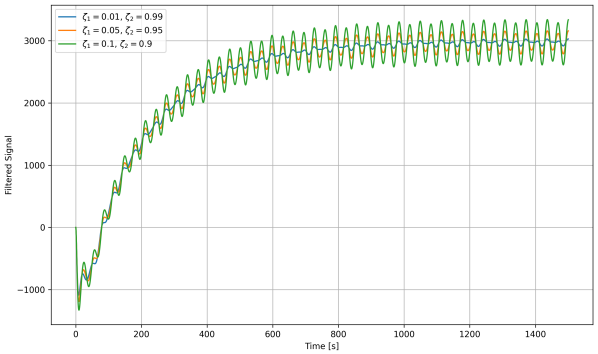

(c) Time-domain signal showing harmonic suppression for $\zeta_1 = 0.01$, $\zeta_2 = 0.99$.

Figure 5.4: Filtering performance of Notch Filter 2 for removing the $2\omega_L$ harmonic.

5.4. No Model Update

In this simulation scenario, the model is simulated without any updates to the degradation parameter α. Consequently, the power coefficient curves in the estimator remain fixed to the clean, undegraded state, while the actual wind turbine operates with the P40 and P400 degraded C_P curves.

The excitation mechanism introduced in the learning algorithm causes the tip-speed ratio to oscillate around the intended reference value. The reference value for the tip-speed ratio has been set to 9.5. This is shown in Figures 5.5a and 5.5b, where the tip-speed ratio behavior is plotted for the P40 and P400 scenarios, respectively.

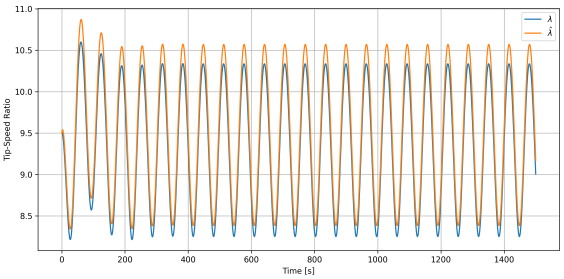

(a) TSR oscillation for P40 degradation scenario without model update. The estimator maintains a clean power coefficient curve, causing discrepancies in TSR tracking.

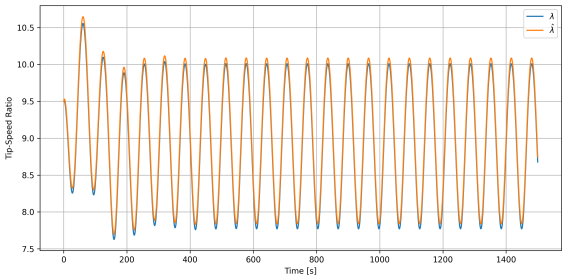

(b) TSR oscillation for P400 degradation scenario without model update. The estimator remains unchanged, resulting in TSR deviation.

Figure 5.5: Tip-speed ratio behavior under P40 and P400 degradation without model adaptation. The lack of model updates prevents the WSE from accurately tracking the optimal TSR.

The mismatch between the estimator (orange) and the real wind turbine (blue) can also be illustrated

in the power coefficient, shown in Figures 5.6a and 5.6b. The differences between the C_P curves of the wind turbine and the estimator are obvious. The wind turbine operates with degraded aerodynamic properties (P40 and P400), attaining values between 0.40 and 0.46, while the estimator assumes clean aerodynamic properties, attaining values between 0.43 and 0.48.

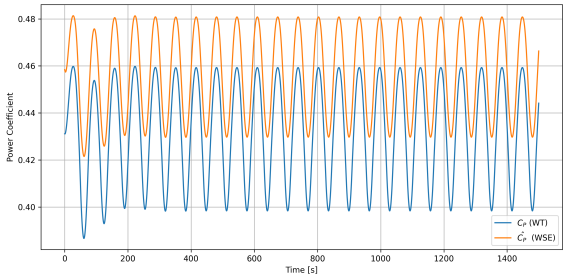

(a) Power coefficient curve for P40 degradation scenario without model update. The estimator assumes clean conditions, resulting in a gap between the estimated and real power coefficient.

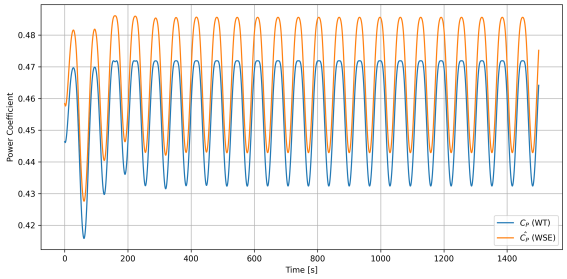

(b) Power coefficient curve for P400 degradation scenario without model update. The estimator assumes clean conditions, resulting in a gap between the estimated and real power coefficient.

Figure 5.6: Comparison of power coefficient curves for the P40 and P400 scenarios without model updates. The wind turbine operates under degraded conditions, while the estimator maintains a clean aerodynamic model, causing efficiency losses.

Ultimately, the goal is to use the learning algorithm to find the correct level of degradation using the degradation model framework. By estimating the degradation parameter α, the estimator can update its internal C_P curve to match the actual turbine behavior. This adaptation enables more accurate tracking of the true power coefficient, thereby closing the gap between the estimated $\hat{C}_P$ and the actual C_P observed in the figures above.

5.5. Phase Offset

(except

Accurate demodulation plays a central role in the learning algorithm described in Section 3.5.1, where the filtered signal is multiplied by a demodulation signal to extract the relevant DC component. The effectiveness of this process depends on the correct phase alignment between the excitation and the system response. While perfect phase alignment maximizes the learning signal and accelerates convergence, the algorithm remains robust to moderate inaccuracies in the chosen phase offset. If the offset is slightly off from the ideal value, the amplitude of the filtered signal decreases, resulting in smaller updates to the degradation parameter α. This leads to slower convergence, but the learning process still progresses in the correct direction and ultimately converges to the true degradation level. However, if the phase offset is poorly chosen, particularly if it results in a phase difference near 180°, the sign of the filtered signal can invert. In this case, the learning algorithm updates α in the wrong direction, causing divergence from the correct solution. Therefore, while the phase offset does not need to be perfectly tuned, it must be close enough to preserve the sign and general magnitude of the filtered signal to ensure consistent and reliable convergence.

To accurately compare the two signals and identify the ideal phase offset, both the original excitation and the filtered output are normalized. This normalization allows for a direct comparison of phase alignment, ensuring that the demodulation step is performed correctly.

Figures 5.7a and 5.7b illustrate the normalized excitation signal alongside the normalized output of the first notch filter for the P40 and P400 degradation scenarios, respectively. In the P40 case (Figure 5.7a), the two signals show close phase alignment, indicating that no additional phase offset is necessary. In contrast, the P400 case (Figure 5.7b) shows a noticeable phase shift between the signals, requiring a compensatory phase offset to align them and ensure effective demodulation.

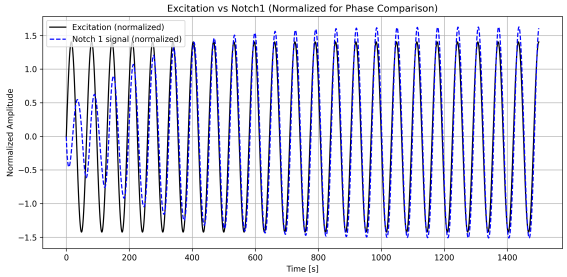

(a) Normalized excitation and notch-filtered signal for the P40 degradation scenario. The signals are closely aligned in phase, and no phase offset is required.

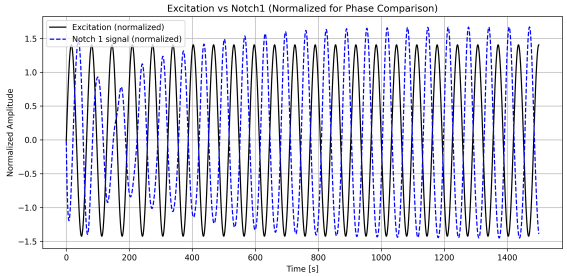

(b) Normalized excitation and notch-filtered signal for the P400 degradation scenario. A clear phase shift is present, indicating the need for a phase offset to align the signals.

Figure 5.7: Comparison of normalized excitation and notch-filtered signals for P40 and P400 scenarios. While the P40 case shows phase alignment, the P400 case exhibits a phase discrepancy introduced by system dynamics or filtering, requiring compensation via a phase offset.

To compensate for the observed phase misalignment in the P400 case, a phase offset of $-165°$ is introduced. This correction brings the filtered signal back in phase with the excitation, optimizing the demodulation process and improving the amplitude response. Figure 5.8 displays the result of this phase correction, showing proper alignment between the excitation and the notch-filtered output.

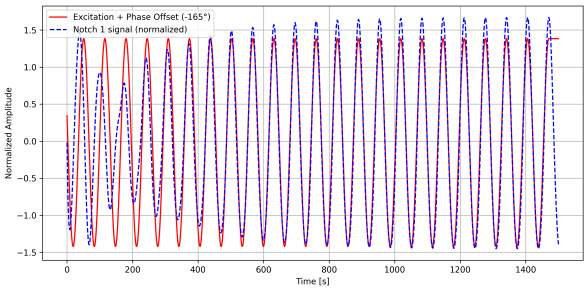

Figure 5.8: Phase-shifted signal after notch filtering for the P400 scenario. A phase offset of $-165°$ ensures proper alignment for optimal signal comparison.

5.5.1. Phase Offset in Learning Process

As the learning algorithm iteratively updates the degradation parameter α, it is important to ensure that the phase difference between the excitation signal and the system response remains consistent. If the phase response of the system changes significantly with varying α, the demodulation phase offset must be adjusted accordingly to maintain accurate signal extraction.

In the initial undegraded case ($\alpha = 0$), the optimal demodulation phase offset for P400 was found to be approximately $-165°$, due to the inherent phase shift introduced by the system dynamics. However, as the degradation parameter increases, the phase characteristics of the system change. For instance, when $\alpha = 0.1$, the phase difference between the excitation signal and the filtered output is nearly zero, as shown in Figure 5.9. This marks a significant shift from the earlier case and indicates that the excitation and response signals are now aligned in phase. Because of this alignment, a phase offset of $0°$ becomes optimal for subsequent experiments where $\alpha \geq 0.1$.

To avoid the complexity of dynamically adjusting the phase offset during learning, the estimation of α for the P400 scenario is initialized at $\alpha = 0.2$ rather than from 0. This ensures that the system response begins within a region where the excitation and output signals are already aligned in phase, allowing a fixed phase offset of $0°$ to be used throughout the learning process. Figure 5.9 confirms that a phase offset of $0°$ is optimal for the updated model.

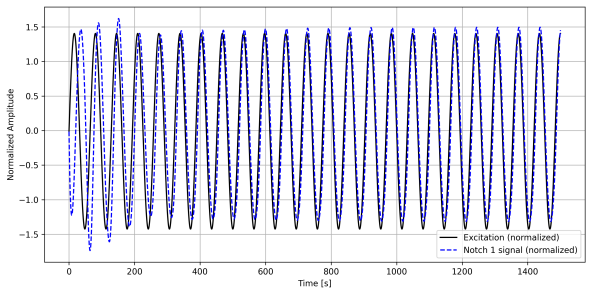

Figure 5.9: Normalized excitation and notch-filtered signal for the P400 scenario with $\alpha = 0.1$. The phase alignment is observed, eliminating the need for a phase offset.

5.5.2. Phase Offset for Reconstructed Curves from Degradation Framework Curves

It became clear that the phase offset can change when changing the degradation parameter α. It can be concluded that for different curves the phase offset can change. Therefore, in this section the real wind turbine is supplied with the constructed curves from the degradation model framework.

The same results where found as in the P40 and P400 curves from AeroDyn. For P40, there was no phase shift found. Therefore a phase offset of 0° works. This is visible in Figure 5.10, where the normalized signals are in phase.

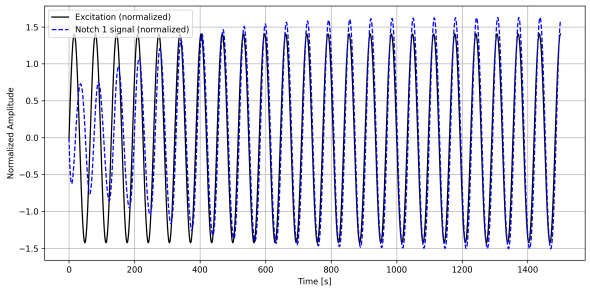

Figure 5.10: Normalized excitation and notch-filtered signal for the AeroDyn-generated P40 scenario. The signals are aligned in phase, confirming that a phase offset of 0° is sufficient.

For the model-generated P400 scenario, the behavior is initially less clear. When $\alpha = 0$, the signal appears to oscillate at twice the intended frequency, making phase alignment with the excitation challenging. This is plotted in Figure 5.11a.

To address this, the α value was increased to 0.1, and the signal alignment was re-evaluated. As shown in Figure 5.11b, the excitation and notch-filtered signals are now in phase. This demonstrates that for the AeroDyn P400 scenario, starting with an α of 0.2 and applying a phase offset of 0° is sufficient for optimal filtering and demodulation.

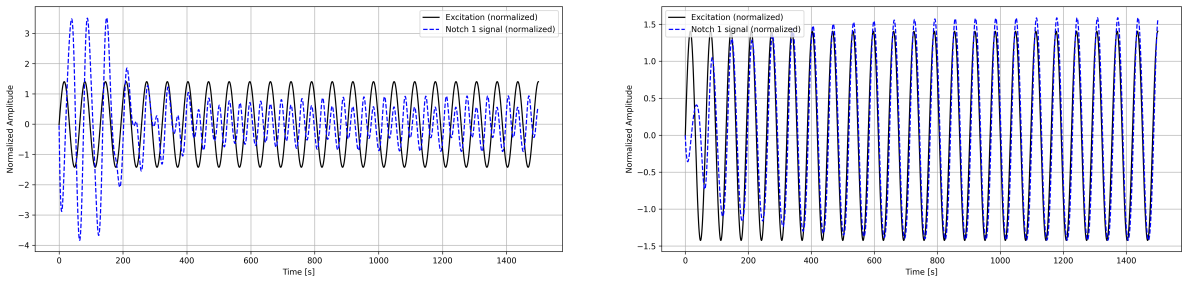

(a) Normalized excitation and notch-filtered signal for AeroDyn P400 with $\alpha = 0$. The signal oscillates at twice the expected frequency, disrupting demodulation.

(b) Normalized excitation and notch-filtered signal for AeroDyn P400 with $\alpha = 0.1$. The signals are in phase, allowing for proper demodulation.

Figure 5.11: Comparison of excitation and notch-filtered signals for AeroDyn P400. At $\alpha = 0$, the signal is unclear, while increasing α to 0.1 aligns the signals.

5.6. Twin Experiment

To validate the accuracy of the degradation model framework and its integration into the WSE-TSR control scheme, a twin experiment was performed. The goal of this experiment is to isolate and assess estimator performance when both the wind turbine and estimator operate with the same power coefficient curves.

From the previously derived degradation mapping, it was found that:

- For the case of $C_L = 0.9$ and $C_D = 1.5$ (P400 reconstruction), an α value of 0.5 is required to match the power coefficient curve.
- For the case of $C_L = 0.9$ and $C_D = 2.0$ (P40 reconstruction), an α value of 1.0 is necessary.

The twin experiment is simulated accordingly. In the wind turbine, the reconstructed C_P curve from the degradation model framework is used. Simultaneously, the estimator assumes a clean power coefficient curve, with the degradation framework applied with $\alpha = 0.5$ for P400 and $\alpha = 1.0$ for P40. This setup allows for comparison between the wind turbine and the estimator when the same exact curves are used.

5.6.1. Analysis for P400 (Moderate Degradation)

For the P400 scenario, the estimation results are shown in Figure 5.12. The estimator's ability to track the TSR is evaluated in Figure 5.12a, where it can be seen that the estimated TSR $\hat{\lambda}$ closely follows the actual turbine TSR, with minimal deviation.

Figure 5.12b shows the time evolution of the power coefficient C_P. The estimator maintains a trajectory that closely follows the turbine's output, with only slight variations over time.

Finally, Figure 5.12c plots the power coefficient against the tip-speed ratio for both the turbine and the estimator. Although the curves align well, a slight overestimation of C_P by the estimator is observed. This deviation is quantified with a maximum absolute difference of 0.0032 and a maximum relative error of 0.0077. These differences are small but indicate a model mismatch likely due to assumptions in linearization or estimator dynamics.

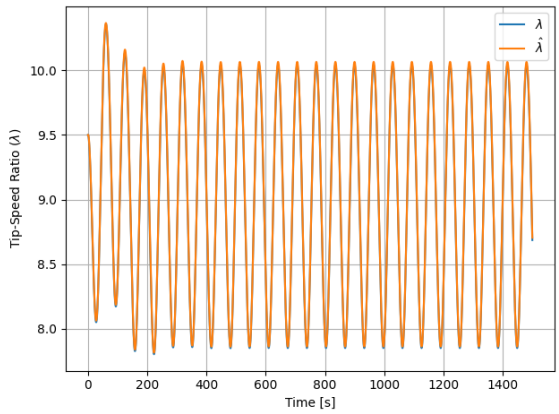

(a) Comparison of actual and estimated TSR for the P40 degradation scenario. The estimator closely tracks the turbine's TSR, demonstrating accurate dynamic estimation performance.

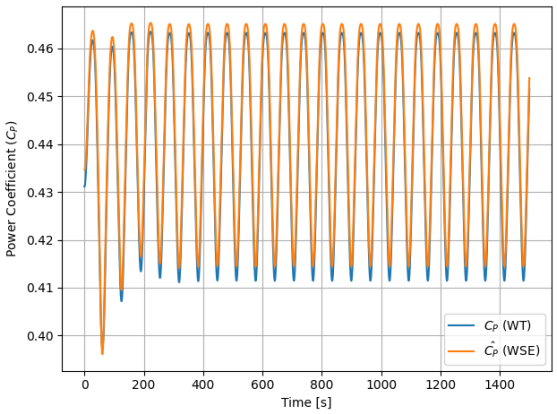

(b) Time evolution of the power coefficient C_P for the P40 case. The estimator replicates the turbine's power extraction characteristics with minimal deviation.

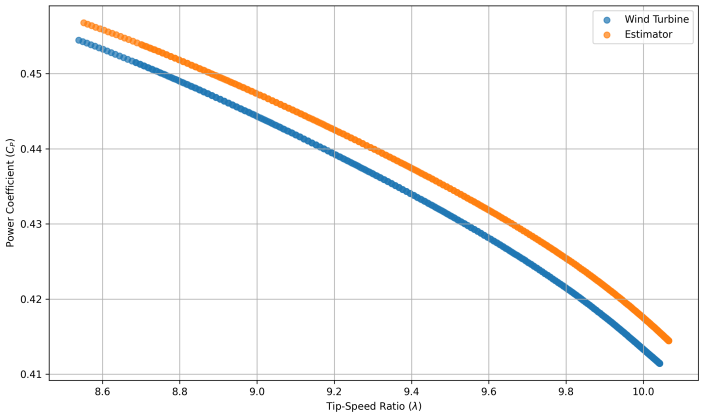

(c) Scatter plot of C_P versus TSR for the P40 scenario. The estimator's curve aligns well with the turbine's, though a slight overestimation of C_P is observed.

Figure 5.12: Twin experiment results for the P40 degradation case. All plots demonstrate strong agreement between the wind turbine and the estimator, which uses the same degradation model internally. The estimator accurately tracks TSR and power coefficient dynamics, with only minor discrepancies.

5.6.2. Analysis for P40 (Severe Degradation)

For the P40 scenario, the estimator's performance is evaluated in the twin experiment using a degradation parameter of $\alpha = 0.5$. The comparison between the actual turbine TSR (λ) and the estimated TSR ($\hat{\lambda}$) is shown in Figure 5.13a. As in the P400 case, the estimator tracks the turbine's dynamics closely, though minor deviations are still visible.

Figure 5.13b presents the time evolution of the power coefficient C_P. The mismatch between the turbine and estimator outputs remains small. The relationship between C_P and TSR is illustrated in Figure 5.13c. The estimator slightly overestimates the power coefficient across the operating range. The maximum absolute difference is 0.0021, with a maximum relative error of 0.0049. These discrepancies, though small, underscore the sensitivity of power coefficient estimation to even minor tracking deviations.

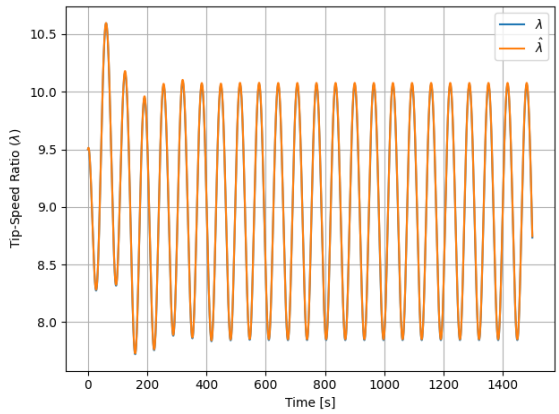

(a) Comparison of actual and estimated TSR for the P400 degradation scenario. The estimator shows strong tracking performance, with only slight deviations from the turbine's dynamics.

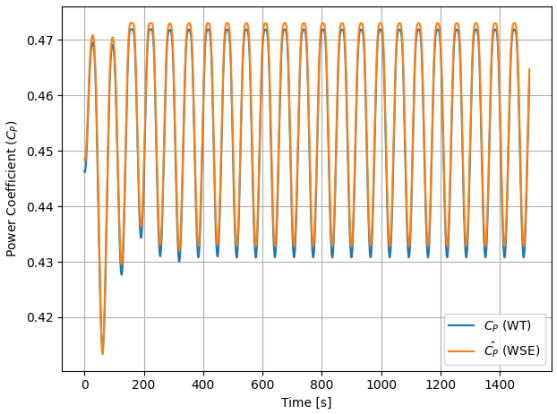

(b) Time evolution of the power coefficient C_P for P400. The estimator closely follows the turbine's performance, with minor discrepancies.

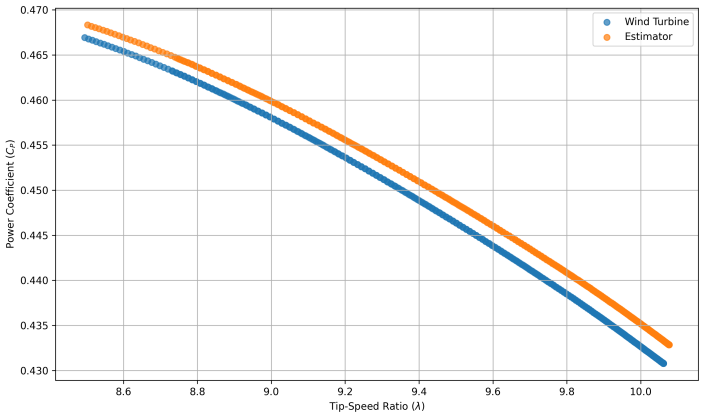

(c) C_P vs. TSR for the P400 case. The estimator slightly overestimates the power coefficient across the range.

Figure 5.13: Twin experiment results for the P400 (moderate degradation) scenario. The estimator successfully tracks the turbine's behavior, both in TSR and power coefficient dynamics. Small overestimations of C_P are observed.

The results show that while the estimator performs well in tracking the TSR and C_P values, slight overestimations are present in both cases. These differences are small, but highlight the importance of accurate internal model calibration.

5.7. Error Analysis Twin Experiment

The previous result showed a slight mismatch between the power coefficient (C_P) curves applied in the Wind Turbine and the Estimator. These differences cause the signal filtering to generate a non-zero update for the degradation coefficient, α. The magnitude of this update is influenced by the amplitude of the excitation applied to the system. Larger excitation amplitudes produce greater output amplitudes, which in turn affect the observed update error. In Figures 5.14a and 5.14b the development of the amplitude of signal after filtering is shown. It can be seen that for none of the applied excitation amplitudes a zero update for the degradation coefficient has been found.

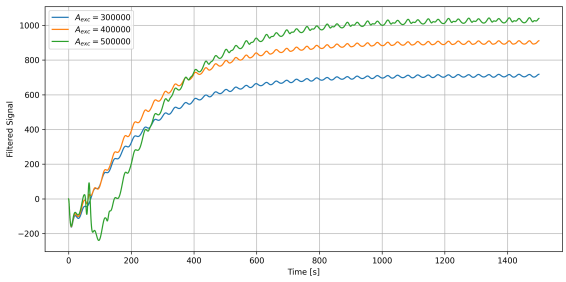

(a) Amplitude of the filtered learning signal for P40 across different excitation amplitudes. Despite matched degradation settings, the signal remains non-zero.

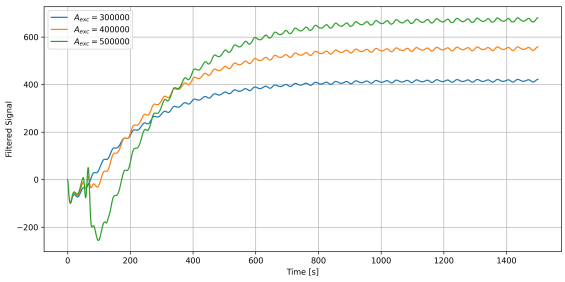

(b) Filtered signal amplitude for P400 as a function of excitation amplitude. The non-zero output confirms persistent error in the estimation, even in the twin setup.

Figure 5.14: Filtered signal output for various excitation amplitudes in the twin experiment. Although the turbine and estimator share the same power coefficient curve, a small estimation error remains, causing non-zero signal output.

To enable consistent comparison across excitation amplitudes, the filtered signal is normalized with respect to the input amplitude. This normalization highlights the relative estimation error, independent of signal strength. The normalized error is plotted in Figures 5.15a and 5.15b. In both cases, the normalized filtered signal decreases with increasing excitation amplitude. This suggests that larger excitations relatively show less error by spanning a wider range of tip-speed ratio values.

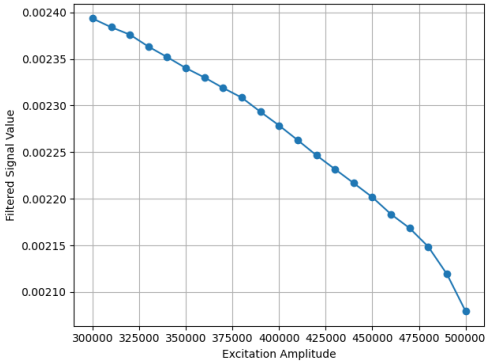

(a) Normalized filtered signal amplitude vs excitation level for P40. Higher excitation leads to lower relative error, likely due to improved coverage of the relevant TSR range.

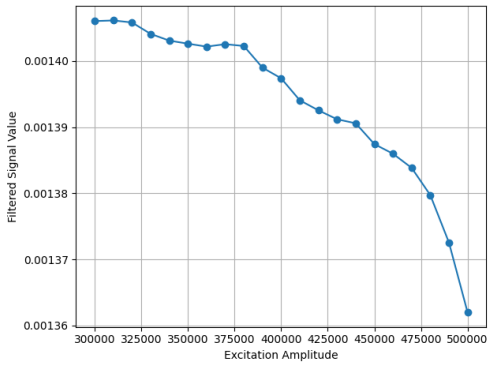

(b) Normalized filtered signal for P400 as a function of excitation amplitude. Higher excitation leads to lower relative error, likely due to improved coverage of the relevant TSR range.

Figure 5.15: Normalized filtered signal amplitude against excitation amplitude for P40 and P400 cases.

5.7.1. Output Signal Adjustment

The filtered error signal serves as the learning input for updating the degradation parameter α. While the signal itself is small, it represents the residual mismatch between the estimator and the turbine model. When this signal is multiplied by the excitation amplitude and the learning gain, it determines the rate at which α is updated.

In the twin experiment setup, where turbine and estimator use identical degradation curves, any non-zero signal represents a systematic estimation error rather than a real degradation change. The largest observed update error, after scaling, is approximately 0.00034. Though this appears negligible, the integration of this small value over time can result in a significant drift in the estimated α, leading to overcompensation or misidentification of the degradation level.

To stabilize α and ensure reliable degradation tracking, a correction step is necessary: the systematic offset in the filtered signal must be subtracted before integration.

Tables 5.1 and 5.2 summarize the filtered signal values and their resulting contribution to the α update error across various excitation amplitudes for the P40 and P400 scenarios, respectively.

Table 5.1: Filtered signal error values for P40 scenario.

Amplitude	Filtered Signal	Error (Filtered × Amplitude)
300000	0.002393	718.02
350000	0.002340	819.09
400000	0.002279	911.46
450000	0.002202	990.75
500000	0.002079	1039.66

Table 5.2: Filtered signal error values for P400 scenario.

Amplitude	Filtered Signal	Error (Filtered × Amplitude)
300000	0.001406	421.81
350000	0.001403	490.92
400000	0.001397	558.95
450000	0.001387	624.34
500000	0.001362	680.99

The error analysis shows that proper correction must be applied during degradation parameter learning to prevent wrong α estimation.

5.7.2. Degradation Parameter Learning for Moderate and Severe Degradation

The learning algorithm is now applied to dynamically learn the degradation parameter α for both the P40 and P400 scenarios. Unlike in the twin experiment, where α was fixed, the value is now learned in real time through integration of the filtered power error signal. To ensure consistent initial conditions, the system is stabilized for the first 1500 time steps. It can be seen in Figures 5.14a and 5.14b, that in the first few timesteps the value of the filter output is not stable. After stabilization, the learning algorithm is activated to adjust the value of α using the output of the filter. Three excitation amplitudes 300000 W, 400000 W, and 500000 W, are used as representative inputs. The results below demonstrate how the learning algorithm responds under these conditions.

The results for the P40 scenario are shown in Figure 5.16. In Figure 5.16a, the evolution of α over time is plotted. It shows that the algorithm successfully converges to a value of exactly 1.0, indicating a perfect match with the expected degradation parameter. The value of the filtered signal, presented in Figure 5.16b, confirms this convergence as it approaches the error (filtered x amplitude) from Table 5.1. Finally, the C_{P} curve in Figure 5.16c shows convergence to the curves that were also found in the twin experiment, see Figure 5.12b.

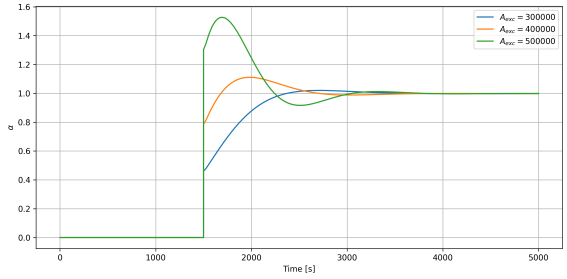

(a) Evolution of degradation parameter α over time for the P40 case. The value converges precisely to 1.0, indicating perfect reconstruction of the degradation state.

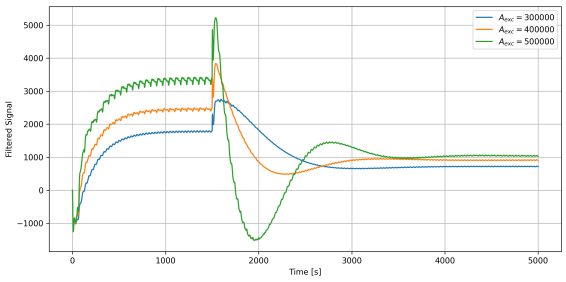

(b) Filtered signal over time for P40. The signal stabilizes fist, after which the learning algorithm updates the parameter and aligns with the predicted systematic error.

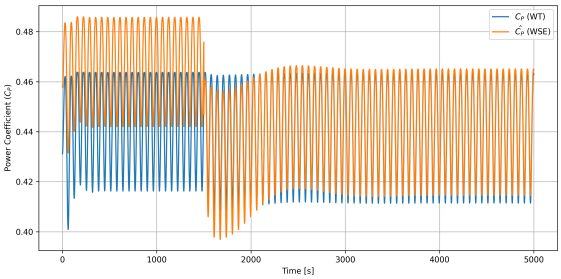

(c) Comparison of power coefficient (C_P) evolution in the estimator and turbine for the P40 case. The estimator's curve matches the turbine's degraded performance as expected.

Figure 5.16: Learning performance for the P40 scenario. The estimator accurately identifies $\alpha = 1.0$, confirming effective degradation tracking and alignment with the turbine's C_P curve.

The same learning process was applied to the P400 scenario. As seen in Figure 5.17, the evolution of α in Figure 5.17a converges to 0.5, which aligns with the expected value for the P400 degradation case. The combined filtered signal plot (Figure 5.17b) confirms this. The C_P curve comparison in Figure 5.17c shows the same overlap between the WSE and the wind turbine model as was found previously for fixed $\alpha = 0.5$, see Figure 5.13b.

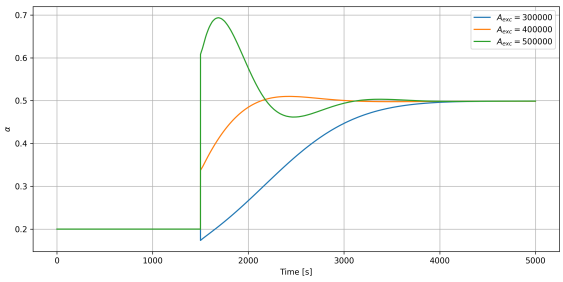

(a) Convergence of α for P400. The algorithm accurately learns the degradation parameter, converging to $\alpha = 0.5$.

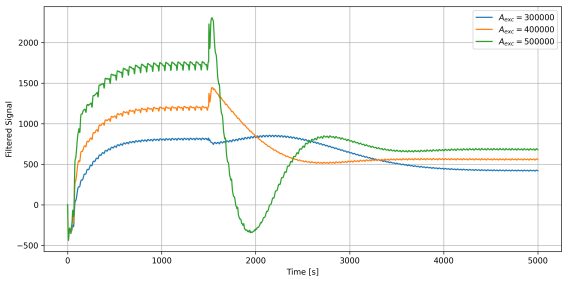

(b) Filtered signal over time for P400. The signal stabilizes first, after which the learning algorithm updates the parameter and aligns with the predicted systematic error.

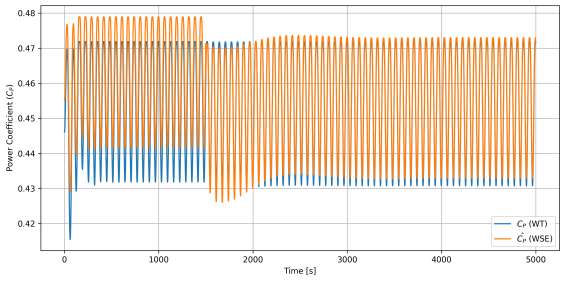

(c) Comparison of power coefficient (C_P) evolution in the estimator and turbine for the P400 case. The estimator's curve matches the turbine's degraded performance as expected.

Figure 5.17: Learning performance for the P40 scenario. The estimator accurately identifies $\alpha = 1.0$, confirming effective degradation tracking and alignment with the turbine's C_P curve.

These results demonstrate that the learning algorithm successfully identifies the correct degradation parameter α in both severe (P40) and moderate (P400) cases. The convergence of α, alignment of the filtered signal with expected error levels, and consistency in C_P curves all confirm the effectiveness of the WSE-TSR learning strategy. This validates the overall degradation model framework implementation in the WSE-TSR control scheme.

5.7.3. Degradation Parameter Learning for P40 and P400 scenarios

In this section, the learning algorithm is applied to optimize the degradation parameter α using real degraded power coefficient curves from AeroDyn rather than those reconstructed by the degradation model. Specifically, the AeroDyn-generated C_P curves for the P40 and P400 cases are directly supplied to the wind turbine. To compensate for the systematic bias observed in the twin experiments, the corresponding error correction values (filtered signal offsets) are incorporated into the learning process. Other settings such as the used excitation amplitudes are the same as in the twin experiment.

The results of the learning process can be seen in Figure 5.18 for the P40 case and in Figure 5.19 for P400. Unlike in the twin experiment, the learned α values no longer converge consistently across different excitation amplitudes. These results indicate that excitation amplitude now affects the learned value of α, suggesting that the system error is influenced by the specific aerodynamic shape of used C_P curves, which differ slightly from the model-based reconstructions in the twin experiment. A more comprehensive investigation of the error is required to properly calibrate the learning of the degradation parameter for shape-altered C_P curves.

In Figure 5.18a, another result that can be seen is that the learned alpha is bigger than the alpha found for the reconstructed CP curve for the P40 case ($\alpha = 1.0$). Partly, a discrepancy from this value is possible since there is a difference between the reconstructed CP curve and the AeroDyn CP curve. However, consistent with Figures 4.10a and 4.10b, the learned degradation parameter value should be lower than 1.0, because the value of C_P is larger in the region of the attained TSR values.

The resulting mismatch is also evident in the C_P curves (Figure 5.18c), where the estimator's C_P curve

deviates from the turbine's C_P curve.

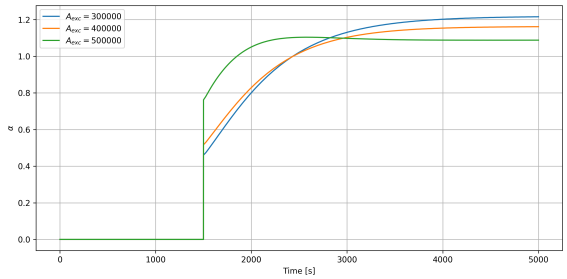

(a) Time evolution of α for the P40 learning case. The learned value varies with excitation amplitude and exceeds the expected value of 1.0.

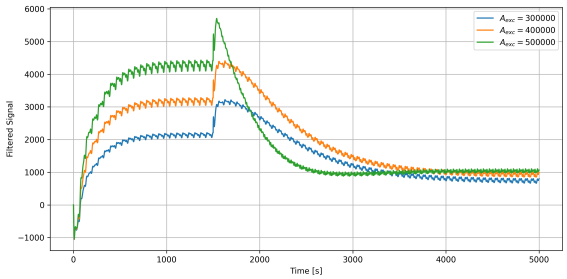

(b) Filtered signal convergence for P40. Different excitation amplitudes result in different converged signal values.

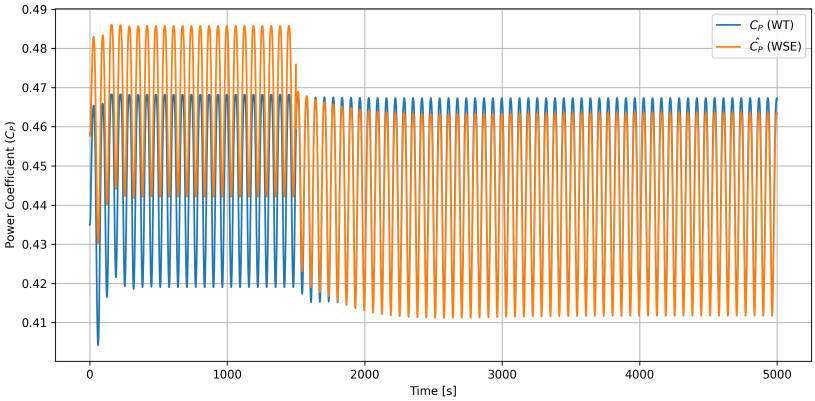

(c) Power coefficient comparison for P40 learning. The estimator's values deviate from the turbine's actual C_P, showing the impact of α overestimation.

Figure 5.18: Learning results for P40 using the real degraded C_P curve from AeroDyn. The learned α values are higher than expected and depend on excitation amplitude, highlighting sensitivity to model differences.

A similar pattern is observed in the P400 learning results (Figure 5.19). In Figure 5.19a, α does not consistently converge to the expected value of 0.5. Instead, it varies with the amplitude of the excitation signal. The filtered signal (Figure 5.19b) again shows different convergence values across amplitudes. Finally, the C_P values in Figure 5.19c illustrates that the estimator diverges from the turbine's C_P values.

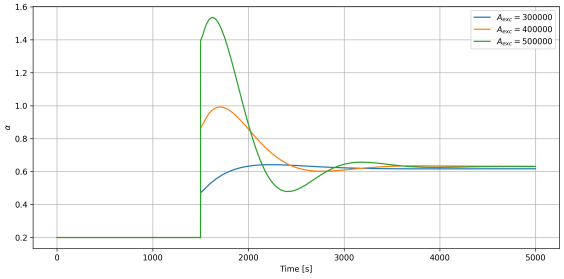

(a) Evolution of α for P400 learning. The final value depends on excitation amplitude and deviates from the expected value of 0.5.

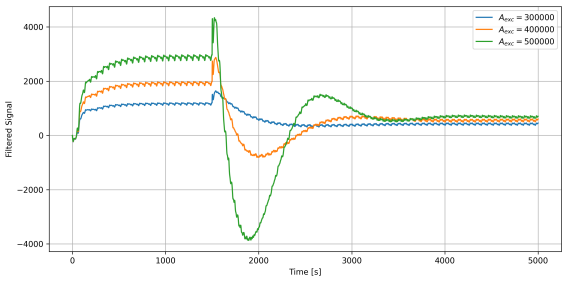

(b) Filtered signal for P400. Convergence value depends on the excitation amplitude used.

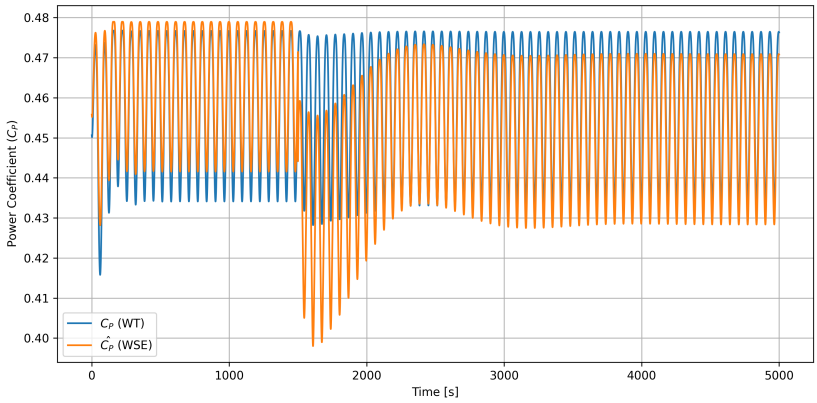

(c) Comparison of C_P values for P400 learning. Discrepancy between turbine and estimator outputs highlights model misalignment.

Figure 5.19: Learning results for P400 using the real degraded C_P curve from AeroDyn. The learned α values are higher than expected and depend on excitation amplitude, highlighting sensitivity to model differences.

These results demonstrate that when realistic C_P curves are used, the learning algorithm becomes sensitive to excitation amplitude and internal model mismatch. The results indicate that, when using real degraded C_P curves, the learning algorithm shows a sensitivity to the excitation amplitude. Unlike the twin experiment where the estimator and wind turbine are calibrated perfectly, real-time learning demonstrates variable convergence, suggesting the internal model error is affected by the specific C_P curve that is used. Further research into these errors is necessary to optimize the algorithm for real-time degradation tracking of α.

6

Discussion

This chapter critically reflects on the assumptions, methods, and limitations of the degradation modeling framework and its integration into the WSE–TSR control scheme.

6.1. Discussion Degradation Framework

This section discusses limitations of the degradation modeling framework developed in this thesis. It reflects on the aerodynamic fidelity of the AeroDyn-based simulations, the simplifications involved in treating isolated degradation mechanisms, and the interpretation of specific degradation cases such as icing.

6.1.1. Fidelity of the AeroDyn Simulations

Throughout this thesis, $C_\mathrm{P} - \lambda$ curves generated by AeroDyn simulations were treated as "true" representations of wind turbine performance. However, it must be acknowledged that AeroDyn itself is a model based on Blade Element Momentum Theory (BEMT), which inherently relies on simplifying assumptions. While BEMT does not fully resolve three-dimensional flow effects, AeroDyn incorporates correction models (e.g., Prandtl's tip loss) to mitigate these limitations.

In the context of the simulations conducted for this thesis, enabling or disabling Prandtl's tip loss correction had no noticeable effect on the resulting $C_\mathrm{P} - \lambda$ curves. This indicates that, under the specific operating conditions and blade parameters studied here, tip loss effects were not dominant. However, in real-world scenarios, tip loss corrections can significantly impact aerodynamic predictions and should not be disregarded.

Given its widespread use in both academic and industrial settings and its demonstrated ability to replicate empirical behavior with high accuracy under a wide range of operational conditions, AeroDyn can reasonably be considered a high-fidelity aerodynamic model. Nonetheless, future validation against experimental data could further enhance the confidence in the capabilities of the degradation model framework.

6.1.2. Limitations of Isolated Degradation Scenarios

The thesis treated each degradation scenario in isolation. While this was useful for analytical clarity and sensitivity studies, in real-world scenarios, degradation rarely occurs in isolation. Blade damage due to erosion may simultaneously occur with icing, for example.

Hence, the current framework does not capture the interactions between different degradation scenarios. A next step could be to extend the model to handle multiple degradation mechanisms simultaneously, potentially using data from real degrading wind turbines.

6.1.3. Interpretation of Icing Degradation

In the icing scenario discussed in Appendix A, geometries of airfoils along the blade span were used as a baseline for progressive ice development over time. While this does not capture the true time-dependent growth of ice on a rotating blade, the underlying physical change to the polars remains valid. Since blade tip regions experience higher tangential velocities, they are more prone to severe icing, leading to aerodynamic characteristics.

Therefore, using segmented blade data as a proxy for icing development is a reasonable assumption that captures the main aerodynamic impact. However, the research could have benefited from a more structured degradation dataset consisting of clearly defined degradation levels, each with corresponding and known lift (C_L) and drag (C_D) coefficients.

6.2. Integration into the WSE-TSR Control Scheme

This section evaluates how the degradation modeling framework was integrated into the WSE-TSR control scheme. It discusses the limitations arising from the learning formulation, particularly the impact of linearization, ill-conditioning, and the scalar representation of degradation. Additionally, it reflects on how the assumed degradation path in parameter space may affect accuracy.

6.2.1. Sources of Learning Error in the WSE-TSR control SCheme

The WSE–TSR control scheme, introduced in Section 2.4.3, presents a promising foundation for real-time aerodynamic model adaptation. It incorporates a wind speed estimator and tip-speed ratio tracking controller, whose learning behavior is analytically understood through the derivation of the closed-loop transfer function $H(s)$. This transfer function, derived via linearization of the nonlinear system dynamics, forms the basis of the learning mechanism: as $H(s) \to 0$, the internal model is assumed to match the actual turbine behavior.

However, several important considerations arise from this formulation:

Effect of Linearization on Learning. The derivation of $H(s)$ assumes a linearization of the nonlinear closed-loop dynamics around a steady operating point. This can introduce approximation errors. In particular, the claim that $H(s) \to 0$ implies perfect model agreement holds only within the local linear regime. Deviations from this linearization due to model nonlinearities or noise may reduce the accuracy of this implication. This is particularly relevant during high-dynamic events or under strong nonlinearity introduced by degradation.

Impact of Ill-Conditioning and Scalar Degradation Modeling. As discussed in [43], the WSE–TSR framework is inherently ill-conditioned because the control scheme couples the unknowns U and $C_P(\lambda)$ into a single measurable quantity. Without additional constraints, the system cannot uniquely separate the effects of changes in wind speed and changes in aerodynamic efficiency. To avoid this ill-conditioning, a scalar degradation model was employed in this thesis, where the estimated power coefficient is defined as:

$$\hat{C}_P(\hat{\lambda}) = C_P^{\diamond}(\hat{\lambda})(1 - \hat{\Gamma}(\alpha, \hat{\lambda})).$$

In particular, the degradation shape $\hat{\Gamma}(\alpha, \lambda)$ used in this thesis is composed of two basis functions related to lift and drag degradation:

$$k_1(\alpha), \quad k_2(\alpha) \quad \text{with} \quad \Gamma(\alpha, \lambda) = k_1(\alpha)\, \gamma_{C_L}(\lambda) + k_2(\alpha)\, \gamma_{C_D}(\lambda).$$

Here, $\gamma_{C_L}(\lambda)$ and $\gamma_{C_D}(\lambda)$ represent degradation effects on lift and drag, respectively, while the scalar-dependent weights $k_1(\alpha)$ and $k_2(\alpha)$ determine the relative contribution of each mechanism. This formulation introduces a known degradation shape $\hat{\Gamma}(\alpha, \hat{\lambda})$ and learns the scalar parameter α. Under this structure, the ill-conditioning is effectively avoided, as the power coefficient is no longer an independent unknown but instead parameterized by a single scalar. However, this simplification introduces its own limitations.

Dependence on Estimated Tip-Speed Ratio. A remaining sensitivity lies in the dependence of the learned degradation shape on the estimated tip-speed ratio $\hat{\lambda}$. Since $\hat{C}_P(\hat{\lambda})$ is evaluated at this estimated point, any bias in wind speed estimation directly affects $\hat{\lambda}$, and thus the evaluation of the degradation function. This feedback loop introduces a possible source of modeling error: even if the assumed degradation shape is accurate, incorrect $\hat{\lambda}$ values may lead to incorrect estimation of α.

6.2.2. Assumed Degradation Path in (k_1, k_2)-Space

The degradation path in the (k_1, k_2) parameter space was assumed to be monotonically increasing. This assumption aligns with the physical expectation that degradation effects, particularly those affecting lift and drag, tend to intensify over time. However, this path was not derived from real turbine degradation data and therefore lacks empirical proof. The selected path serves as a synthetic trajectory to test the learning algorithm, but realistic degradation paths could be formed with data from degrated wind turbines.

7

Conclusion

The goal of this thesis was to investigate whether shape-altering changes in the C_P–λ curve, caused by aerodynamic degradation of wind turbine blades, could be captured in a mathematical degradation model framework and subsequently used to enhance the WSE-TSR learning control scheme for performance tracking and operational optimization. To achieve this, five key subquestions were formulated. In this conclusion, each of these questions is addressed, based on the theoretical background, methodology, and results discussed throughout the thesis.

1. What are the main blade degradation scenarios, and what is known from existing research about their impact on aerodynamic polars, wind turbine performance, and Annual Energy Production (AEP)?

This question was addressed in Section 2.2, where a detailed literature review was conducted to identify key blade degradation scenarios and evaluate their aerodynamic and operational impact. The focus was on LEE, as this is well-documented in existing research and offers quantifiable effects on aerodynamic performance.

For LEE, four major studies were reviewed. The first showed that severe erosion could reduce the lift coefficient by up to 53% and increase the drag coefficient by 314%, leading to a loss of Annual Energy Production (AEP) of around 2–3.7%. The second study confirmed these findings, reporting up to a 63% reduction in the lift-to-drag ratio and AEP losses reaching 6.6% in severe cases. The third study investigated how degradation affects the C_P–λ curve, showing that contamination and erosion shifted the optimal tip-speed ratio from 7.6 (clean blade) to 8.2 and 8.4 respectively, while the maximum power coefficient decreased from 0.483 to 0.431 and 0.351. This highlights the dual effect of LEE: reducing efficiency and altering optimal control parameters. A fourth paper introduced a simplified method for simulating erosion by applying multiplicative scaling to the clean aerodynamic polars: for moderate erosion (P400), the lift coefficient was reduced by 10% and the drag coefficient increased by 50%; for severe erosion (P40), the lift coefficient was reduced by 10% and the drag coefficient doubled. This approach proved effective for modeling degradation in simulations.

Additional simulations were performed for general degradation scenarios such as icing, blade cracks, blade twist, chord reduction, and shaft tilt (see Section A and B). The icing case served as a comparative aerodynamic benchmark alongside the LEE case. For the other scenarios, either no significant aerodynamic impact was observed in the simulation, or existing literature indicated that their primary effects are structural rather than aerodynamic. As a result, these cases were not considered central to the aerodynamic degradation framework developed in this thesis.

In addition to reviewing literature, this thesis also examined degradation scenarios through simulations (see Section 4.1 and Section 3.1) that altered aerodynamic polars directly. For the LEE case, modifying the polars resulted in a shift of the optimal tip-speed ratio from 7.5 (clean blade) to 7.8 (moderate degradation) and 8.4 (severe degradation), while the maximum power coefficient dropped from 0.487 to 0.456. These results confirm that degradation not only reduces efficiency but also affects the location

of peak performance.

Overall, these findings establish that different degradation mechanisms have distinct and quantifiable effects on the aerodynamic performance of wind turbines. They alter not only the magnitude but also the shape and peak location of the C_P–λ curve, thereby impacting both energy production and control strategy effectiveness.

2. How do these degradation scenarios influence the shape of the $C_P - \lambda$ curve across operational conditions?

From the literature review and simulation studies, it became evident that LEE influences the shape of the C_P–λ curve in a non-uniform manner. Specifically, multiple studies reported a shift in the optimal tip-speed ratio (λ_{opt}) toward higher values as a result of degradation.

These insights from literature were confirmed through simulations using AeroDyn (see Sections 3.1 and 3.4). By increasing drag coefficients and reducing lift coefficients in the input polars, a consistent pattern emerged: the C_P–λ curve not only decreased in peak magnitude but also shifted to the right. This indicated that a change in lift and drag coefficients alters both the efficiency and the optimal operational point of the turbine.

To further investigate the underlying mechanisms, lift and drag were modified independently in the simulations (see Section 3.3) to find their individual contributions to the change in the $C_P - \lambda$ curve. The results showed that reducing lift primarily caused a horizontal shift of the peak C_P toward higher λ values. In contrast, increasing drag mainly lowered the entire curve vertically without significantly affecting the peak position. This decomposition provides a clearer understanding of how specific aerodynamic degradations translate into measurable performance changes: lift reduction is responsible for shifting the peak, while drag increase causes the overall loss in efficiency.

In summary, both the literature and simulation results confirm that degradation does not simply reduce the C_P curve uniformly but introduces shape-altering effects. These include a decrease in peak value, a rightward shift in λ_{opt}, and changes to the curve's slope and curvature. Such transformations underscore the inadequacy of scalar degradation and justify the need for a more flexible, shape-altering framework to capture the operational consequences of blade degradation.

3. How can this deeper understanding be used to formulate a more general physics-based mathematical degradation framework?

Building on the insight that lift and drag affect the C_P–λ curve in fundamentally different ways, a generalized, physics-based degradation model was proposed in Section 3.3. The formulation decouples the degradation function $\Gamma(\lambda)$ into separate contributions from lift and drag effects, leading to the following model:

$$\Gamma(\lambda) = k_1 \gamma_{C_L}(\lambda) + k_2 \gamma_{C_D}(\lambda)$$

Here, $\gamma_{C_L}(\lambda)$ and $\gamma_{C_D}(\lambda)$ represent degradation profiles resulting from isolated changes in lift and drag, respectively. The sensitivity parameters k_1 and k_2 quantify the magnitude of each effect. This formulation is based on simulation results (see Sections 3.3 and 4.3) that show lift reductions and drag increases influence the C_P–λ curve in different ways: drag mainly shifts the curve downward, while lift causes a horizontal shift. When these effects are combined linearly, the resulting C_P curve closely matches the degraded curves observed in AeroDyn, suggesting that a linear combination of lift- and drag-related effects can effectively reproduce real aerodynamic degradation behavior.

The model was validated against AeroDyn simulation results (see Section 4.5.1) by applying known multiplicative changes to the C_L and C_D coefficients. Specifically, for the LEE scenarios based on sandpaper roughness P400 and P40, the lift was scaled by a factor of 0.9, while the drag was increased by factors of 1.5 and 2.0, respectively. These polars were used to generate C_P–λ curves in AeroDyn, which were then compared to those predicted by the mathematical degradation framework. The results showed a strong agreement, confirming that the sensitivity parameters k_1 and k_2 could be accurately learned and used to reconstruct the full degraded power coefficient curve.

However, the icing case presented a limitation. In this scenario, the degradation in C_L and C_D could not be approximated by a simple multiplicative scaling. Instead, the aerodynamic changes were more complex and spatially varying, particularly along the blade span. As a result, the current framework, which assumes multiplicative adjustments, will not represent reality. To extend the framework's applicability to such cases, an updated version of the framework is needed, one that incorporates shifts in C_L and C_D rather than scaling. The overall structure of the model, however, remains useful, as it allows for a flexible representation of degradation using predefined component profiles and tunable sensitivity parameters.

In summary, the proposed degradation model provides a physically interpretable and efficient way to model complex C_P–λ changes under LEE.

4. How is the proposed degradation framework implemented in the WSE-TSR control scheme, and is this implementation effective?

While the extended degradation model introduces two sensitivity parameters, k_1 and k_2, the WSE-TSR control scheme provides only scalar feedback, derived from the power tracking error. This makes the direct estimation of both parameters an underdetermined problem, as multiple (k_1, k_2) pairs can result in identical estimated $C_P(\lambda)$ behavior. To address this limitation, a parametric degradation path was introduced (see Section 3.5.2) to reduce the two-dimensional sensitivity space to a one-dimensional curve controlled by a single degradation parameter α. This parameterization enables learning within the WSE-TSR scheme. The path is defined as:

$$k_1(\alpha) = \begin{cases} 0.277 \dfrac{1 - e^{-a\alpha}}{1 - e^{-a0.5}}, & \text{if } \alpha < 0.5 \\ 0.277, & \text{if } \alpha \geq 0.5 \end{cases}$$

with shaping parameter $a = 6$. This parametrization ensures that the undegraded state corresponds to $\alpha = 0$, the moderate (P400) erosion case corresponds to $\alpha = 0.5$ with $(k_1, k_2) = (0.277, 0.5)$, and the severe (P40) erosion case corresponds to $\alpha = 1.0$ with $(k_1, k_2) = (0.277, 1.0)$.

To validate this integration, a twin experiment was conducted (see Section 3.5.1). After performing optimal frequency response and phase offset analysis, the system was shown to be capable of learning the correct degradation levels for both the moderate and severe cases. However, an error was observed in the learning behavior for the P400 and P40 cases. This error may be attributed to subtle variations in the C_P–λ curves, which affect the system's power tracking error and thereby influence the learning outcome. An error offset correction, extracted from the twin experiment, was applied to enable learning.

To validate the integration of the degradation framework into the WSE-TSR control scheme, a twin experiment was conducted (see Section 3.5.1). The system was shown to successfully track and identify degradation for both moderate (P400) and severe (P40) erosion scenarios. However, a consistent learning error was observed. As discussed in Section 6.2.1, this deviation may stem from two possible factors: approximation errors introduced by the linearization used in the derivation of the learning transfer function $H(s)$; and a sensitivity to errors in the estimated tip-speed ratio $\hat{\lambda}$, which directly affects the evaluation of the degradation function.

To account for this systematic deviation, an offset correction based on the twin experiment results was applied to enable accurate degradation learning under the scalar degradation model. While this workaround is effective for the tested cases, it highlights the importance of robust estimation methods and potential future improvements in control system structure or design.

Overall, the results demonstrate that the proposed degradation path and learning scheme enable the WSE-TSR controller to adapt its internal model to degraded conditions with a reasonable degree of accuracy. Nonetheless, discrepancies between degradation profiles and their impact on the control signal suggest that further research is needed. In particular, the sensitivity of the learning algorithm to differences in C_P curves and associated error dynamics must be better understood to improve robustness and generalization of the framework.

5. Can the modified learning framework suggest operational adjustments or maintenance actions based on learned degradation profiles?

While maintenance recommendations were not the central focus of this work, the framework showed potential in this direction. As demonstrated in Section 3.5.2 and Section 3.5.1, the learned degradation profiles shifted the optimal operating point (λ_{opt}), suggesting the need for dynamic adjustment of the setpoint to mitigate performance loss. In scenarios where degradation becomes severe, the control system could flag a significant deviation from nominal behavior, potentially triggering maintenance alerts. Thus, the proposed framework lays the foundation for future integration of maintenance strategies and operational setpoint adaptation, enabling predictive maintenance and lifetime extension.

Main Conclusion

This thesis set out to answer the following central research question:

Can shape-altering changes in the C_P–λ curve, caused by aerodynamic blade degradation, be captured in a mathematical model and used to enhance the WSE-TSR control scheme for real-time performance tracking and operational optimization?

The first part of this question, whether shape-altering degradation can be modeled, was researched through a combination of literature review and simulation-based analysis. It was shown that the degradation scenario, LEE, induces non-uniform changes in the aerodynamic polars, including lift reduction, drag increase, and shifts in the optimal tip-speed ratio. These effects alter both the efficiency and the shape of the C_{P}–λ curve. Based on this understanding, a new physics-based degradation framework was developed that decouples lift and drag effects and models their contributions separately. Using two sensitivity parameters, the framework accurately reconstructs degraded performance curves.

The second part of the research question, whether this model can enhance the WSE-TSR control scheme, was addressed by integrating the framework into the controller via a one-dimensional degradation trajectory. Since the full model includes two parameters, k_1 and k_2, but the WSE-TSR control scheme provides only scalar feedback, a dimensionality reduction was necessary. A physically motivated degradation path was therefore introduced, allowing the model to be parameterized by a single scalar α. This transformation enabled the learning of degradation severity from the control system. Simulation results showed that the learning algorithm could adapt to degraded conditions and update the internal turbine model. However, to ensure convergence to the correct value, a small offset correction had to be subtracted from the learned signal. This highlights a limitation of the current scalar feedback loop and points to future improvements in the learning algorithm integration in the WSE-TSR control system.

7.1. Future Work

While the current thesis proposes a framework for modeling and learning aerodynamic degradation in wind turbines, several open questions and opportunities for further research remain. These are outlined below.

7.1.1. Ill-conditioning

A promising extension is to have simultaneous learning of both degradation sensitivity parameters k_1 and k_2. Currently, the framework uses a predefined degradation path in the (k_1, k_2)-space. Allowing these parameters to be learned adaptively could provide a more flexible approach. One potential method for achieving this is to modify the WSE-TSR control scheme such that the degradation path is no longer fixed. This could involve including an additional optimization algorithm that updates k_1 and k_2 based on feedback from the system dynamics.

As discussed earlier, the WSE–TSR control scheme is ill-conditioned, particularly when attempting to estimate multiple degradation parameters simultaneously. In this thesis, the problem was mitigated by reducing the model complexity to a single-parameter degradation function. However, this simplification limits the flexibility of the framework. To address this, future research should investigate the ill-conditioning problem in greater depth. One promising direction is to incorporate not only amplitude but also phase information into the learning algorithm. Specifically, analyzing the phase difference

between the input excitation and the output power error signal of the closed-loop system may provide additional information that helps to better separate the effects of wind speed from those of aerodynamic degradation.

7.1.2. Analysis of Twin Experiment Error

In the twin experiment, a residual error was consistently observed in the output signal, even though the estimator and the simulated wind turbine used the same internal model structure and power coefficient curve. Because both systems share identical C_P–λ curves, this error cannot stem from shape mismatches in the aerodynamic model. Instead, the source of the discrepancy likely lies in the estimator dynamics or numerical effects.

This finding underscores an important challenge: even in idealized conditions, the learning mechanism does not perfectly converge. This raises the question of how the system would behave when the true C_P curve is not known, as is the case in real-world applications. Future research should investigate how sensitive the learning algorithm is to modeling mismatches, such as deviations in the shape of $C_P(\lambda)$ or errors in the wind speed estimation. Controlled studies where shape errors are introduced could help identify which features of the curve are most critical to accurate learning.

Moreover, the current learning method relies primarily on amplitude information in the power tracking error signal. As proposed in Section 7.1.1, future extensions could explore the use of phase differences between the input excitation and the system response. These phase dynamics may carry independent information that helps distinguish between different types of model uncertainty.

7.1.3. Experimental Validation and Real-World Data

Applying the proposed framework to real-world wind turbine data or wind tunnel experiments is an important step for assessing its practical relevance and robustness. However, such validation efforts are only meaningful if the ill-conditioning problem outlined in Section 7.1.1 is first resolved. Specifically, the current limitation of using only scalar feedback restricts the learning algorithm's ability to independently estimate multiple degradation parameters. Without addressing this constraint, it would be difficult to interpret learning results from field data or experimental setups, where the true degradation state is unknown. Therefore, experimental validation should be pursued after the learning scheme has been extended by introducing additional feedback, to ensure that the degradation parameters can be reliably identified in real-world conditions.

7.1.4. Research on Blade Pitch

This thesis focuses on how blade degradation affects the C_P–λ curve, specifically within the partial-load region. However, degradation may also impact turbine performance across the full operational spectrum, including variations in blade pitch angle. This would require extending the analysis from a two-dimensional $C_P(\lambda)$ curve to a three-dimensional $C_P(\lambda, \beta)$ surface. Such an approach would allow investigation of degradation effects not only in the partial-load regime but also in transition and full-load regions. Future work could explore how degradation shifts this surface and how these changes could be captured in control strategies that adjust both tip-speed ratio and pitch.

7.1.5. Inverse Mapping from Scalar Metrics

In the current implementation, the two-dimensional degradation model defined by (k_1, k_2) is reduced to a single scalar parameter α by constraining the learning to a predefined path. However, if two independent scalar degradation metrics are available, such as the relative power loss at the optimal tip-speed ratio and the global maximum power loss, then the current degradation state could be identified by finding the intersection point of their corresponding iso-contours in the (k_1, k_2) space.

This concept builds directly on the analysis presented in Section 4.4, where contour maps of the metrics $\Delta C_P^{\lambda^*}$ and $\Delta C_P^{\max}$ were constructed (see Figures 4.9a and 4.9b). These results showed that while multiple degradation combinations can yield the same scalar performance loss, the intersection of iso-contours from two metrics could uniquely identify the current degradation state. Future research could explore this inverse mapping approach as a way to enable two-parameter learning without relying on a predefined degradation trajectory.

References

[1] Deltares. *Offshore Wind Energy*. Accessed: 2025-07-02. 2025. URL: https://www.deltares.nl/en/expertise/areas-of-expertise/energy-transition/offshore-wind-energy.

[2] J. Smith and M. Lee. "The Potential of Offshore Wind Energy in the Renewable Energy Transition". In: *Renewable Energy Journal* 45 (2020), pp. 123–135.

[3] Author(s). *Offshore Wind Energy Potential: Unlocking Stronger and More Consistent Winds*. Available online at https://example.com/offshore_potential.

[4] WindEurope. *Wind Energy Market Outlook - H1 2024*. Tech. rep. Accessed: 2025-01-22. WindEurope, June 2024. URL: https://windeurope.org/intelligence-platform/product/latest-wind-energy-data-for-europe-autumn-2024/#overview.

[5] Intergovernmental Panel on Climate Change. *Climate Change 2022: Impacts, Adaptation, and Vulnerability*. Available online at https://www.ipcc.ch/report/ar6/wg2/.

[6] P. Johnson and K. Andersson. "Bigger Turbines, Bigger Benefits: Optimizing Offshore Wind Power". In: *Energy Science Reviews* 78 (2021), pp. 67–89.

[7] T. Miller and L. Brown. "Land-Use Constraints and the Shift to Offshore Wind Farms". In: *Environmental Research Letters* 14 (2019), pp. 105–118.

[8] R. Kumar and A. Singh. "The Economics of Offshore Wind Turbine Operation and Maintenance". In: *Energy Economics* 39 (2021), pp. 15–29.

[9] R. Taylor and J. Green. "Digital Twins: The Next Frontier in Wind Energy Management". In: *Renewable Systems Innovations* 11 (2020), pp. 201–220.

[10] Michiel Zaaijer and Axelle Viré. *Introduction to Wind Turbines: Physics and Technology*. Unpublished manuscript. Version 13/09/2024. 2024.

[11] María Isabel Blanco. "The economics of wind energy". In: *Renewable and Sustainable Energy Reviews* 13.6 (2009), pp. 1372–1382. ISSN: 1364-0321. DOI: https://doi.org/10.1016/j.rser.2008.09.004. URL: https://www.sciencedirect.com/science/article/pii/S1364032108001299.

[12] Alicia Key, Owen Roberts, and Annika Eberle. "Scaling trends for balance-of-system costs at land-based wind power plants: Opportunities for innovations in foundation and erection". In: *Wind Engineering* 46 (Dec. 2021), p. 0309524X2110602. DOI: 10.1177/0309524X211060234.

[13] Engineers Guidebook. *Understanding Wind Energy: Benefits and Challenges*. Accessed: 2025-05-21. Aug. 2024. URL: https://engineersguidebook.com/understanding-wind-energy-benefits/.

[14] Guoyan Cao, Karolos Grigoriadis, and Yaw Nyanteh. "LPV Control for the Full Region Operation of a Wind Turbine Integrated with Synchronous Generator". In: *TheScientificWorldJournal* 2015 (Mar. 2015), p. 638120. DOI: 10.1155/2015/638120.

[15] G. Lazzerini et al. "COFLEX: A novel set point optimiser and feedforward-feedback control scheme for large flexible wind turbines". In: *Wind Energy Science Discussions* 2024 (2024), pp. 1–35. DOI: 10.5194/wes-2024-151. URL: https://wes.copernicus.org/preprints/wes-2024-151/.

[16] Jason H. Laks, Lucy Y. Pao, and Alan D. Wright. "Control of wind turbines: Past, present, and future". In: (2009), pp. 2096–2103. DOI: 10.1109/ACC.2009.5160590.

[17] Mohammad Ebrahimpour et al. "Numerical Investigation of the Savonius Vertical Axis Wind Turbine and Evaluation of the Effect of the Overlap Parameter in Both Horizontal and Vertical Directions on Its Performance". In: *Symmetry* 11.6 (2019). ISSN: 2073-8994. DOI: 10.3390/sym11060821. URL: https://www.mdpi.com/2073-8994/11/6/821.

[18] Eduardo Menezes, Alex Araújo, and Nadège Bouchonneau. "A review on wind turbine control and its associated methods". In: *Journal of Cleaner Production* 174 (Oct. 2017). DOI: 10.1016/j.jclepro.2017.10.297.

[19] Jacob Aho et al. "Controlling Wind Energy for Utility Grid Reliability". In: *Mechanical engineering (New York, N.Y.: 1919)* 135 (Sept. 2013). DOI: 10.1115/1.2013-SEP-4.

[20] Juan Leon et al. "Power Fluctuations In High Installation Density Offshore Wind Fleets". In: *Wind Energy Science* (Aug. 2020). DOI: 10.5194/wes-2020-95.

[21] Erich Hau. *Wind Turbines: Fundamentals, Technologies, Application, Economics*. 2nd. Berlin: Springer, 2005. ISBN: 9783540251225.

[22] Sandia National Laboratories and Stanford University. "Impact of blade leading edge erosion on wind turbine performance". In: *Wind Energy Science* 5.2 (2020), pp. 549–564. DOI: 10.5194/wes-5-549-2020.

[23] Ying Du et al. "Damage detection techniques for wind turbine blades: A review". In: *Mechanical Systems and Signal Processing* 141 (2020), p. 106445. ISSN: 0888-3270. DOI: https://doi.org/10.1016/j.ymssp.2019.106445. URL: https://www.sciencedirect.com/science/article/pii/S0888327019306661.

[24] Leon Mishnaevsky. "A review on wind turbine blade erosion and its effects on energy yield and maintenance costs". In: *Renewable and Sustainable Energy Reviews* 168 (2023), p. 113016. DOI: 10.1016/j.rser.2022.113016.

[25] Leon Mishnaevsky Jr. et al. "Leading edge erosion of wind turbine blades: Understanding, prevention and protection". In: *Renewable Energy* (2021), pp. -–-. DOI: DOI_HERE.

[26] Sunil Gupta, Ravi S. Iyer, and Sanjeev Kumar. *Digital Twins: Advancements in Theory, Implementation, and Applications*. Springer, 2023.

[27] Yuriy Marykovskiy et al. *Extended Taxonomy of Digital Twins*. Available at Zenodo: https://doi.org/10.5281/zenodo 2023. DOI: 10.5281/zenodo.8021787.

[28] Clemence Rubiella, Cyrus A. Hessabi, and Arash Soleiman-Fallah. "State of the art in fatigue modelling of composite wind turbine blades". In: *International Journal of Fatigue* 117 (2018), pp. 230–245.

[29] Mahmood M. Shokrieh and Roham Rafiee. "Simulation of fatigue failure in a full composite wind turbine blade". In: *Composite Structures* 74.3 (2006), pp. 332–342.

[30] Sebastiaan Mulders et al. "A learning algorithm for the calibration of internal model uncertainties in advanced wind turbine controllers: A wind speed measurement-free approach". In: *Proceedings of the 2023 American Control Conference (ACC)*. May 2023, pp. 1486–1492. DOI: 10.23919/ACC55779.2023.10156125.

[31] Ing. Dr. A. Betz. "The Maximum of the Theoretically Possible Exploitation of Wind by Means of a Wind Motor". In: *Wind Engineering* 37.4 (2013), pp. 441–446. DOI: 10.1260/0309-524X.37.4.441.

[32] Hongbo Hou et al. "Actuator disk theory and blade element momentum theory for the force-driven turbine". In: *Ocean Engineering* 285 (2023), p. 115488. ISSN: 0029-8018. DOI: https://doi.org/10.1016/j.oceaneng.2023.115488. URL: https://www.sciencedirect.com/science/article/pii/S0029801823018723.

[33] John B. Woodward and Tor Erik Andersen. "Marine Engines". In: *Encyclopedia of Physical Science and Technology (Third Edition)*. Ed. by Robert A. Meyers. Third Edition. New York: Academic Press, 2003, pp. 121–132. ISBN: 978-0-12-227410-7. DOI: https://doi.org/10.1016/B0-12-227410-5/00404-X. URL: https://www.sciencedirect.com/science/article/pii/B012227410500404X.

[34] Olivier Cleynen. *Figure from Wikimedia Commons*. Own work, CC0. 2014. URL: https://commons.wikimedia.org/w/index.php?curid=16100403.

[35] Woobeom Han, Jonghwa Kim, and Bumsuk Kim. “Effects of contamination and erosion at the leading edge of blade tip airfoils on the annual energy production of wind turbines”. In: *Renewable Energy* 115 (2018), pp. 817–823. ISSN: 0960-1481. DOI: https://doi.org/10.1016/j.renene.2017.09.002. URL: https://www.sciencedirect.com/science/article/pii/S0960148117308649.

[36] Alessio Castorrini, Andrea Ortolani, and M. Sergio Campobasso. “Assessing the progression of wind turbine energy yield losses due to blade erosion by resolving damage geometries from lab tests and field observations”. In: *Renewable Energy* 218 (2023), p. 119256. ISSN: 0960-1481. DOI: https://doi.org/10.1016/j.renene.2023.119256. URL: https://www.sciencedirect.com/science/article/pii/S0960148123011710.

[37] Devesh Kumar and Mario Rotea. “Brief communication: Real-time estimation of the optimal tip-speed ratio for controlling wind turbines with degraded blades”. In: *Wind Energy Science* 9 (Nov. 2024), pp. 2133–2146. DOI: 10.5194/wes-9-2133-2024.

[38] T. H. Malik and C. Bak. “Challenges in detecting wind turbine power loss: the effects of blade erosion, turbulence, and time averaging”. In: *Wind Energy Science* 10.1 (2025), pp. 227–243. DOI: 10.5194/wes-10-227-2025. URL: https://wes.copernicus.org/articles/10/227/2025/.

[39] Emil Krog Kruse, Christian Bak, and Anders Smærup Olsen. “Wind tunnel experiments on a NACA 633-418 airfoil with different types of leading edge roughness”. English. In: *Wind Energy* 24.11 (2021), pp. 1263–1274. ISSN: 1095-4244. DOI: 10.1002/we.2630.

[40] Livia Brandetti et al. “Analysis and multi-objective optimisation of model-based wind turbine controllers”. In: *Wind Energy Science* (June 2023). DOI: 10.5194/wes-2023-66.

[41] Sebastiaan Mulders et al. “An iterative data-driven learning algorithm for calibration of the internal model in advanced wind turbine controllers”. In: *IFAC-PapersOnLine* 56 (Jan. 2023), pp. 8406–8413. DOI: 10.1016/j.ifacol.2023.10.1035.

[42] Yichao Liu et al. “The Immersion and Invariance Wind Speed Estimator Revisited and New Results”. In: *IEEE Control Systems Letters* 6 (2022), pp. 361–366. DOI: 10.1109/LCSYS.2021.3076040. URL: https://doi.org/10.1109/LCSYS.2021.3076040.

[43] L. Brandetti et al. “On the ill-conditioning of the combined wind speed estimator and tip-speed ratio tracking control scheme”. English. In: *Journal of Physics: Conference Series* 2265.3 (2022). 2022 Science of Making Torque from Wind, TORQUE 2022 ; Conference date: 01-06-2022 Through 03-06-2022. ISSN: 1742-6588. DOI: 10.1088/1742-6596/2265/3/032085.

[44] Dr. Donatella Zappalá. *Personal communication*. 2025.

[45] Raja Imran, Dil Hussain, and Mohsen Soltani. “An Experimental Analysis of the Effect of Icing on Wind Turbine Rotor Blades”. In: (May 2016). DOI: 10.1109/TDC.2016.7520041.

[46] Liangquan Hu et al. “Wind Turbines Ice Distribution and Load Response under Icing Conditions”. In: *Renewable Energy* 113 (May 2017). DOI: 10.1016/j.renene.2017.05.059.

[47] Ece Sagol, Marcelo Reggio, and Adrian Ilinca. “Issues concerning roughness on wind turbine blades”. In: *Renewable and Sustainable Energy Reviews* 23 (2013), pp. 514–525. ISSN: 1364-0321. DOI: https://doi.org/10.1016/j.rser.2013.02.034. URL: https://www.sciencedirect.com/science/article/pii/S1364032113001366.

[48] Oscar Castro and Kim Branner. “Effect of tunneling cracks on structural property degradation of wind turbine blades”. In: *Composite Structures* 268 (2021), p. 113914. ISSN: 0263-8223. DOI: https://doi.org/10.1016/j.compstruct.2021.113914. URL: https://www.sciencedirect.com/science/article/pii/S0263822321003743.

[49] Leon Mishnaevsky Jr. “Root Causes and Mechanisms of Failure of Wind Turbine Blades: Overview”. In: *Materials* 15 (Apr. 2022), p. 2959. DOI: 10.3390/ma15092959.

[50] Xiaobo Chen et al. “Experimental investigation on structural collapse of a large composite wind turbine blade under combined bending and torsion”. In: *Composite Structures* 160 (2017), pp. 435–445. DOI: 10.1016/j.compstruct.2016.10.095.

[51] P. S. Veers. “Yaw misalignment in wind turbines”. In: *Wind Energy* (2003). White paper summary cited in Sereema yaw misalignment study.

[52] M. J. Churchfield et al. “Detailed Field Test of Yaw-Based Wake Steering”. In: *NREL Technical Report NREL/TP-5J00-66734* (2018). DOI: 10.2172/1475806.

A

Literature Review and Methodology General Degradation

This chapter investigates general blade degradation mechanisms beyond LEE, focusing on their aerodynamic impacts. Section A.1.1 reviews studies on icing and blade cracks. Section A.2.1 explains how these degradations are modeled in a BEMT framework like OpenFAST. Section A.3.1 outlines the methodology used to simulate the degradation scenarios. Together, these sections provide the foundation for evaluating broader degradation effects on wind turbine performance. The results of this methodology can be seen in Appendix B.

A.1. Overview Additional Blade Degradation in Wind Turbines

A.1.1. Icing Scenarios

Wind turbines operating in cold climates are significantly affected by atmospheric icing, which alters the aerodynamic properties of the blades, leading to changes in power output. This section examines how icing affects wind turbine aerodynamics, using studies conducted on the NREL Phase VI and NREL 5MW wind turbines, as well as wind tunnel experiments on the NACA 633-418 airfoil.

Experimental Wind Tunnel Study on the Effect of Icing on Wind Turbine Blades

Imran et al. [45] performed an experimental study to investigate the aerodynamic consequences of ice accumulation on wind turbine rotor blades. The study focuses on how different ice distributions affect the power coefficient (C_P). Three small-scale 3D-printed models of the NREL 5 MW rotor blade were created and tested in a wind tunnel. All models used the same blade profile but differed in their surface condition: a clean blade, a blade with ice on the leading edge, and a blade with ice on both the leading and trailing edges. The experiments measured lift, drag, and moment coefficients using strain gauges, and these aerodynamic coefficients were put into WT-Perf (a BEM simulation tool) to generate C_P curves for comparison. In addition to the clean and iced blade configurations, a reference curve was included. This curve represents the ideal aerodynamic performance from design documentation and is used as a baseline for comparison with the experimentally obtained results.

The results, shown in Figure A.1, highlight the aerodynamic changes associated with different icing scenarios. Figure A.1a shows the C_P curves for all three blade configurations at a pitch angle of 4 degrees. At this low pitch angle, the presence of ice on both the leading and trailing edges leads to a significant reduction in C_P and causes a distinct change in the shape of the curve compared to the clean and leading-edge-only cases. This altered shape indicates a change in aerodynamic behavior that would likely affect control performance.

In contrast, Figure A.1b illustrates C_P behavior at a higher pitch angle of 16 degrees. Here, the overall C_P values are reduced for the iced blades, but the general shape of the curves remains similar to the clean case. However, a noticeable shift of the peak C_P toward higher tip-speed ratios is observed,

particularly in the lead-and-trail-edge ice configuration. This shift suggests that turbine operation under icing conditions would require retuning to maintain efficiency.

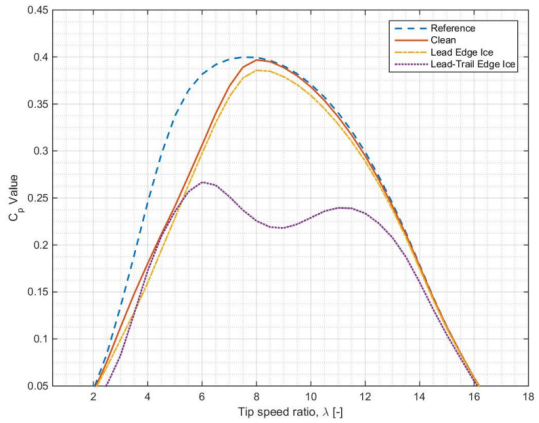

(a) C_P vs. λ for clean and iced blades at 4° pitch. Severe icing (leading and trailing edge) significantly reduces C_P and alters the curve shape. [45]

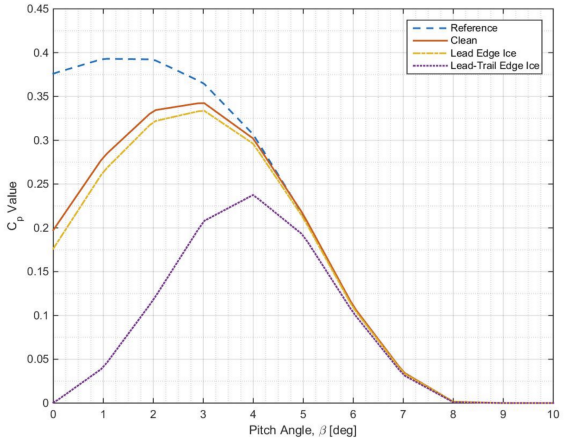

(b) C_P vs. λ at 16° pitch. All iced configurations show a decrease in C_P, and the peak shifts to higher tip-speed ratios, requiring potential control retuning. [45]

Figure A.1: Effect of ice accumulation on power coefficient curves at different pitch angles. The plots compare clean blades (measured), blades iced on the leading edge only, and both leading and trailing edges (measured), against the reference data. At low pitch, the shape of the curve is significantly distorted by severe icing; at higher pitch, the peak shifts but the shape remains mostly intact. [45]

Section-Based Aerodynamic Impact of Ice Accretion on Wind Turbine Blades

Hu et al. [46] investigated how ice accumulates along different spanwise sections of a wind turbine blade and how this affects the aerodynamic performance, focusing on lift and drag coefficients. Using simulations of the NREL 5-MW reference turbine, the authors analyzed different blade regions. Each region was evaluated under different icing scenarios using a combination of CFD analysis and ice accretion modeling. These simulations considered both the geometric distribution of ice and its effects on the aerodynamic coefficients.

The aerodynamic impact of ice accretion varied significantly along the blade. In the root section, where aerodynamic loading is relatively low, the effect of icing was minimal. However, the mid-span and tip regions, which contribute most to power generation, exhibited substantial degradation. The lift coefficient (C_L) decreased, and the drag coefficient (C_D) increased significantly in these regions, leading to a reduced lift-to-drag ratio. To analyze this in detail, Hu et al. examined six specific sections near the tip—labeled F through K—spanning from 80% to 98.5% of the blade length. These sections were chosen because they experience the highest ice accumulation and aerodynamic loads. The simulations showed that the lift coefficient (C_L) decreased and the drag coefficient (C_D) increased significantly in these regions, leading to a sharp reduction in lift-to-drag ratio and aerodynamic efficiency.

Figure A.2 shows the changes in lift and drag coefficients across the blade span due to ice accretion. The plots indicate that the degradation becomes more severe towards the blade tip, where both C_L reduction and C_D increase are more pronounced. The left subplot illustrates this by comparing sections F through K: as the analysis moves from F (less ice) to J (more ice), the lift coefficient consistently decreases over the full range of angle of attack, showing that increasing ice buildup leads to a substantial aerodynamic change. However, the data for section K deviates slightly from this trend. This may result from a significant geometric alteration in that part of the blade due to heavy ice accretion, which could have led to complex flow separation effects not present in other regions [46]. The right subplot shows the corresponding drag behavior; as the severity of icing increases from F to K, the drag coefficient increases significantly. This combination of reduced lift and increased drag confirms that ice accretion has a negative impact on aerodynamic efficiency toward the blade tip.

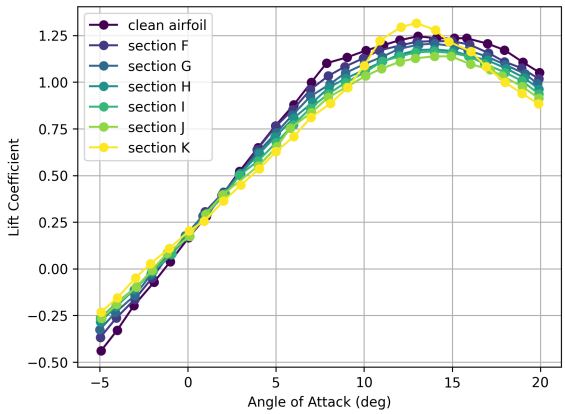

(a) Spanwise distribution of the lift coefficient (C_L) under clean and iced blade conditions, based on data from Hu et al. [46]. Ice accumulation significantly reduces C_L toward the blade tip, where aerodynamic loading is highest.

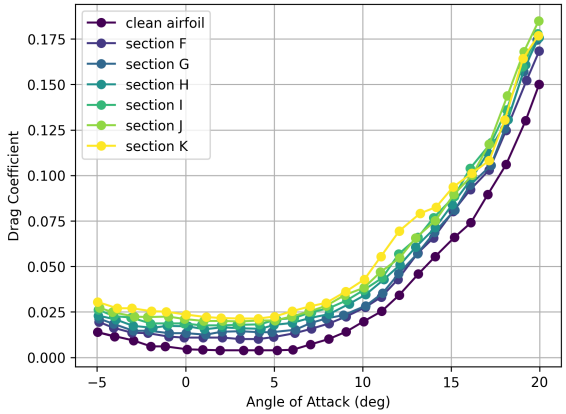

(b) Spanwise distribution of the drag coefficient (C_D) for clean and iced blades, based on data from Hu et al. [46]. Ice causes a sharp increase in C_D, especially near the tip, indicating higher aerodynamic losses.

Figure A.2: Effect of ice accretion on aerodynamic performance along the blade span. Ice buildup leads to reduced lift and increased drag, particularly in mid-span and tip regions that contribute most to power generation.

A.1.2. Blade Crack Scenarios

No studies were found on aerodynamic property changes due to blade cracks. It is possible that cracks do not significantly alter aerodynamic properties but instead increase the risk of structural failure. A possible approach is to introduce a crack and monitor its effect on the $C_P - \lambda$ curve over time. In this case, no change may be detected until the blade (partially) breaks.

This section has shown that LEE and icing consistently reduce aerodynamic performance and, consequently, turbine energy output. A key insight observed in both erosion and icing studies (Kumar and Rotea [37], Imran et al. [45]) is that increasing degradation not only lowers the maximum power coefficient ($C_P^{\max}$) but also shifts the C_P–λ curve to the right, indicating a higher optimal tip-speed ratio. This shift reflects the need for operational retuning to maintain maximum performance. Another important finding, demonstrated by Malik and Bak [38], is that the aerodynamic impact of erosion can be efficiently modeled by simply scaling the clean airfoil's lift and drag coefficients, offering a practical and computationally efficient method for simulating degraded blades. These insights provide both theoretical understanding and strategies for integrating degradation into a degradation framework.

A.2. Simulation of Blade Degradation Scenarios in BEMT

Wind turbine blades are subject to various degradation mechanisms that alter their aerodynamic and structural properties. These damage scenarios are modeled in OpenFAST by modifying the relevant module inputs. The most common damage cases and their associated OpenFAST modules are summarized below.

Ice Formation on Blades

Ice formation on turbine blades increases the mass and alters the aerodynamic properties by reducing lift and increasing drag. These changes require updates in the AeroDyn module to reflect the degraded aerodynamic performance accurately. Additionally, the blade mass distribution changes. The ElastoDyn module is updated to account for the increased blade mass distribution caused by the ice, which can affect the structural dynamics of the turbine. Ice formation also increases the demands on the turbine's pitch and yaw control systems as they work harder to maintain optimal performance, requiring changes in the ServoDyn module to simulate these control adjustments.

Material Loss

Besides LEE, other degradation is the physical loss of material at the blade surface, which can lead to a reduction in the local chord length of the airfoil. This kind of material loss alters the blade's

aerodynamic geometry, reducing its energy capture capability. In AeroDyn, this can be modeled by decreasing the chord length in the airfoil input definitions, allowing for simulation of the performance impact due to geometric erosion.

Blade Cracks
Blade cracks reduce the structural integrity of the blade, which is modeled by decreasing stiffness and damping in the ElastoDyn module. These changes are essential to capture the effects of the cracks on the blade's dynamic behavior, such as vibrations and increased flexibility. The aerodynamic performance near the crack is also altered due to localized drag increases. These effects are difficult to represent in AeroDyn. This is because AeroDyn utilizes spanwise sections with assigned airfoil properties. As a result, it is not currently possible to simulate one or more cracks directly within the AeroDyn framework.

Structural Deformation
Structural deformation, such as blade bending or twisting, alters the effective angle of attack and stress distribution along the blade. These effects are typically modeled in the ElastoDyn module using stiffness reduction and shape functions to capture dynamic response. In this work, ElastoDyn is not used. Instead, the aerodynamic effects of blade twisting are approximated in AeroDyn by manually increasing the blade twist angle, which allows the influence of twist-induced aerodynamic degradation to be analyzed.

Other Degradation Cases
In addition to the degradation scenarios discussed above, other forms of damage or operational faults can also affect turbine performance. For example, lightning strikes may cause localized surface damage or structural failure in blade sections, resulting in altered aerodynamic properties or even partial breakage. Another example is bug accumulation. This increases surface roughness and disrupts laminar flow. Though typically considered minor, long-term accumulation can reduce lift and increase drag, similar to mild erosion [47].

These types of degradation typically manifest as either changes in the blade's aerodynamic polar data or through yaw misalignment, where the rotor is not aligned with the incoming wind. Although they are not always physical changes to the blade geometry, their aerodynamic impact can still be represented within AeroDyn through appropriate adjustments to input parameters such as airfoil polars.

A.2.1. Implementation of Degradation in AeroDyn
All aerodynamic degradation scenarios are implemented in OpenFAST through modifications to the AeroDyn input files. AeroDyn models each blade as a sequence of spanwise sections, with each section linked to an airfoil file that defines aerodynamic coefficients as a function of angle of attack. Blade degradation could be simulated in AeroDyn by changing certain information:

- **Replacing airfoil polars:** Degraded regions can be assigned to airfoil files representing eroded, iced, or otherwise damaged aerodynamic surfaces.
- **Adjusting chord length:** In cases involving material loss, such as severe erosion, the local chord length can be reduced in the AeroDyn blade input to simulate changes in airfoil geometry.
- **Altering twist angle:** Structural twist or deformation may be modeled by increasing or decreasing the twist angle at specific blade sections, affecting the inflow angle and local aerodynamic force.
- **Altering yaw alignment:** Yaw misalignment, where the rotor is not perfectly aligned with the incoming wind, reduces aerodynamic efficiency and alters the distribution of inflow angles along the rotor plane. Although not a form of blade damage, it can be used to study the aerodynamic consequences of misalignment or control system degradation.

A.3. Methodology General Degradation Scenarios
A.3.1. General Degradation
In this section, three generic degradation mechanisms are simulated by modifying settings in the AeroDyn input files. These scenarios are modeled independently to isolate their individual effects on the power coefficient. The modifications are as follows:

- **Blade deformation:** Cracks in wind turbine blades are known to locally reduce stiffness and alter structural properties, particularly affecting the bending of the material [48]. This structural weakening can result in localized twist deformations under aerodynamic loading. Especially in outer blade regions where aerodynamic forces are highest [49].

  In this study, the effects are modeled in AeroDyn by applying a perturbation to the blade twist angle at approximately 70% of the blade span (i.e., the outer 30%). This location is selected based on findings from structural analyses and field reports, which indicate that outer-span regions are prone to damage while also being aerodynamically sensitive [48].

  Although no existing studies directly quantify the twist angle shift due to cracks, the deformation modes observed in collapse testing of a 47-meter composite blade under combined bending and torsion suggest angular perturbations consistent with twist changes between 1–3° [50]. Therefore, the selected twist perturbations of 1°, 2°, and 3° represent a plausible range of deformation severity. While simplified, this modeling strategy enables a controlled sensitivity analysis of perturbations in twist angle of the blade.
- **Material loss:** Wind turbine blades are subject to harsh environmental conditions, which can lead to surface wear and material loss, particularly toward the blade tip where impact velocities are highest. In AeroDyn, this is represented by reducing the chord length of the blade in the outer 20 m of the span. This region is commonly associated with tip degradation in field observations and experimental studies [49, 35].

  Chord reductions of 3%, 5%, and 10% are applied at the affected blade nodes. While these values are not taken directly from a specific dataset, they are representative of mild to severe surface degradation scenarios. In this work, they are chosen as part of a controlled sensitivity analysis to study the aerodynamic impact of geometric degradation alone. Associated changes in lift and drag coefficients are not included in this scenario, in order to isolate the effects of pure geometric alteration.
- **Shaft tilt:** Long-term mechanical wear or yaw misalignment can result in a change in the shaft tilt angle, causing the rotor axis to deviate from its intended alignment. Such deviations alter the inflow direction seen by the rotor plane and may reduce aerodynamic efficiency by changing the effective angle of attack along the blade span.

  In this study, the degradation is simulated by modifying the default shaft tilt angle from $-5°$ to $-6°$ and $-7°$. These changes are not derived from specific field measurements, but they are chosen to reflect deviations that could plausibly result from structural issues over time. This range of values is consistent with those reported in sensitivity analyses or misalignment cases in existing literature [51, 52]. The purpose of this scenario is to evaluate, through controlled sensitivity analysis, the aerodynamic consequences of small deviations in rotor alignment.

Each degradation scenario is simulated independently to assess its isolated effect on the power coefficient (C_P) curve. While these forms of degradation often occur concurrently in real-world turbines, separating them in analysis allows for clearer interpretation of their individual aerodynamic impacts.

B

Results General Degradation Cases

This appendix contains supplementary simulation results for aerodynamic degradation mechanisms beyond leading-edge erosion. These include:

- Blade Twist Deformation: Modeled by increasing the twist angle in the outer 30% of the blade span.
- Material Loss: Simulated through reduction in chord length near the blade tip.
- Shaft Tilt Variations: Simulated by altering the rotor tilt angle.
- Icing: Aerodynamic coefficients obtained from literature and approximated through scalar fits.

These results provide additional context for understanding how different types of degradation influence aerodynamic performance. While the primary modeling focus of this thesis remains on LEE, the simulations here illustrate the diversity and sensitivity of C_P–λ responses under other conditions.

B.1. General Degradation

In addition to the LEE degradation, three other degradation scenarios are simulated using AeroDyn: blade deformation, material loss, and shaft tilt.

The impact of blade deformation, modeled as increased twist at 70% of the blade span, is shown in Figure B.1a. Increasing the twist angle leads to a slight reduction in aerodynamic performance, although the maximum power coefficient remains relatively unaffected. Additionally, the tip-speed ratio at which this maximum occurs does not shift significantly.

Material loss effects, shown in Figure B.1b, are simulated as a reduction in chord length over the outer 20 meters of the blade. In contrast to twist deformation, the chord reduction leads to a noticeable rightward shift in the optimal tip-speed ratio. However, a significant drop in performance is only observed when the chord reduction exceeds 5%. A reduction that is unlikely to occur in practice through delamination or surface erosion alone.

The influence of shaft tilt is presented in Figure B.1c, where deviations from the nominal -5° to -6° and -7° show little to no impact on the power coefficient curve. This suggests that small variations in shaft tilt have limited aerodynamic effect in the partial-load region.

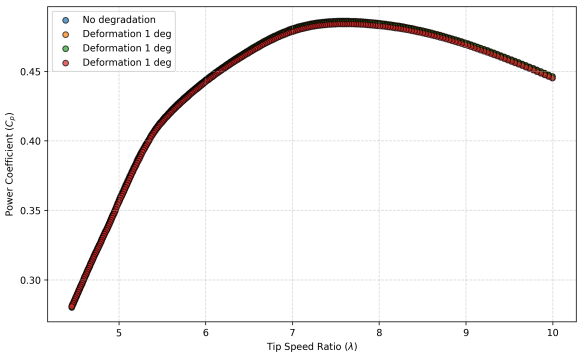

(a) Effect of blade twist deformation: 1°, 2°, and 3° increases at 70% blade span. Minor performance loss is observed with no significant change in optimal tip-speed ratio.

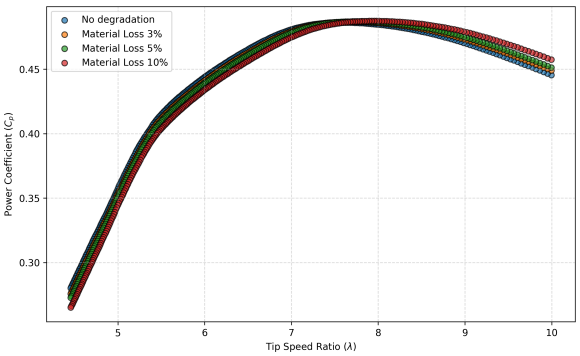

(b) Effect of material loss due to chord reduction by 3%, 5%, and 10% in the outer 20 m of the blade. Larger reductions cause a shift in the optimal tip-speed ratio.

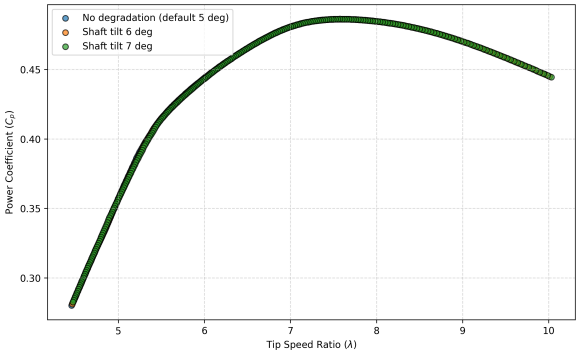

(c) Effect of shaft tilt variation from the baseline -5° to -6° and -7°. No significant impact on aerodynamic performance is observed.

Figure B.1: Power coefficient curves under general degradation scenarios. Each plot shows the aerodynamic impact of a specific degradation type on wind turbine performance.

B.2. Real Icing Case

While the primary focus of this thesis is on leading-edge erosion (LEE), icing presents a fundamentally different degradation mechanism with distinct aerodynamic signatures. This section investigates whether the proposed degradation modeling framework can approximate the effects of icing using the same scalar-based fitting approach developed for LEE. The aim is not to develop a full icing model, but rather to test the framework's flexibility and identify its limitations when applied to a structurally different type of surface degradation.

As discussed in Section A.1.1, to model icing, aerodynamic coefficient data was extracted from the paper by Hu et al. [46], which studied icing effects on the NREL Phase VI wind turbine. The paper selected six blade sections (F through K) for icing simulations. Figure B.2 shows the normalized cross-sectional geometries for these sections, demonstrating increasing ice buildup.

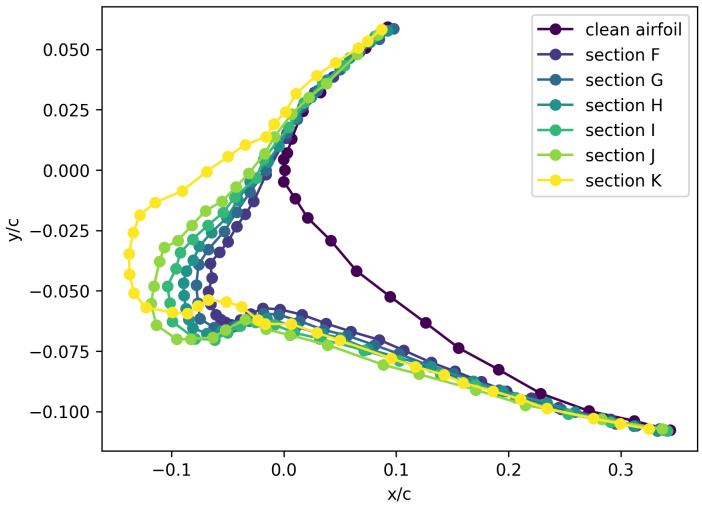

Figure B.2: Normalized airfoil geometries (x/c, y/c) for the clean airfoil and iced blade sections F through K, as digitized from Hu et al. [46]. Ice accretion visibly increases toward the tip.

To interpret the aerodynamic effect of icing within the proposed degradation framework, scalar multipliers were identified to best approximate the iced C_L and C_D curves. These scalars were applied to the clean reference coefficients $C_L^\diamond$ and $C_D^\diamond$.

The best-fit scalar multipliers for each blade section are reported in Table B.1. Figure B.3a shows the results of scaling $C_L^\diamond$ to fit the iced C_L curves. Although the overall trend is captured, mismatches appear near the peak lift values, especially at high angles of attack. This suggests that icing alters the C_L curve in a manner that cannot be fully captured by scalar scaling. Figure B.3b presents the corresponding C_D fit results. Here, the deviations are even more pronounced. The scalar-scaled curves fail to match the observed curves in the iced sections. This highlights that icing alters the C_D curve in manner that cannot be fully captured by scalar scaling.

Table B.1: Best-fit scalings for iced blade sections (applied to $C_L^\diamond$, $C_D^\diamond$)

Section	Scaling on $C_L^\diamond$	Scaling on $C_D^\diamond$
F	0.9680	1.1836
G	0.9470	1.2439
H	0.9223	1.2834
I	0.9044	1.3254
J	0.8818	1.3457
K	0.9151	1.3428

Despite these limitations, the fitted scalars provide a useful approximation that allows iced blade sections to be positioned in the (k_1, k_2) space. The results form the basis for the trajectory analysis presented in the next section.

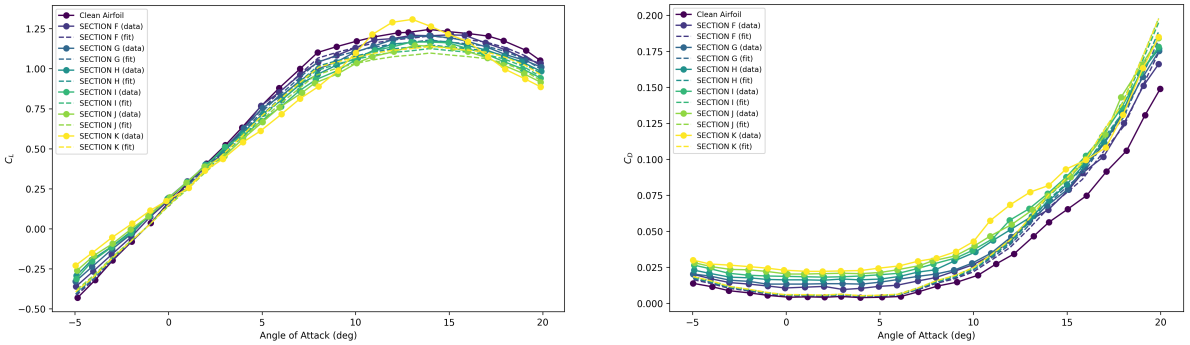

(a) Best-fit scaling of degraded C_L curves from Hu et al. [46].

(b) Best-fit scaling of degraded C_D curves from Hu et al. [46]

Figure B.3: Approximation of iced airfoil aerodynamics using scalar fits. (a) Lift coefficient curves are moderately well captured except near stall. (b) Drag coefficient fits show larger errors, indicating that scalar scaling does not fully represent icing effects.

One important observation is that the degradation caused by icing may not be well represented by a multiplicative scaling approach. In particular, the drag coefficient curves show a more uniform upward shift rather than scaling. This suggests that a model using additive offsets (i.e., representing the degraded coefficients as $C_L = C_L^{\diamond} + \Delta C_L$ and $C_D = C_D^{\diamond} + \Delta C_D$), might better capture the effect of ice accretion on airfoil aerodynamics.

B.2.1. Degradation Path for Icing

This section supplements the (k_1, k_2) degradation space analysis presented in Section 4.5.2 by including a projected degradation path for icing. Using the best-fit scalar approximations of iced C_L and C_D values from Table B.1, each iced blade section was mapped into the (k_1, k_2) space using the same decoupled degradation model framework introduced earlier. Although this scalar-based approach does not fully capture the complex aerodynamic changes caused by icing, it provides an estimate of its aerodynamic change over time in the parameter space. Figure B.4 shows the trajectory of the icing degradation states overlaid on the degradation severity maps introduced earlier.

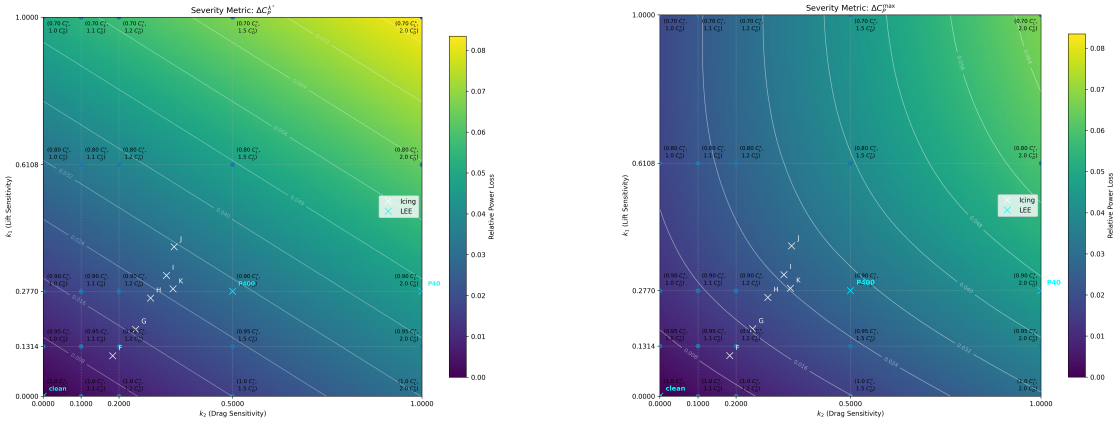

(a) Relative power loss at λ^* with the icing cases F until K overlaid.

(b) Relative maximum C_P loss map with iced blade cases. Degradation follows a distinct path in (k_1, k_2) space.

Figure B.4: Icing degradation path visualized in (k_1, k_2) space using scalar-fitted aerodynamic coefficients. (a) shows power loss at the optimal tip-speed ratio λ^*; (b) shows the relative loss of peak C_P.

The icing trajectory differs notably from the LEE path: it shows a stronger influence on both lift and drag degradation components, as reflected by the simultaneous growth in both k_1 and k_2. This distinction supports the idea that degradation types may be distinguishable in sensitivity space—potentially enabling classification and diagnostic applications. However, due to the limitations of using scalar fits for icing, the presented (k_1, k_2) trajectory should be treated as illustrative rather than definitive.