Surrogate-based optimization under uncertainty of wind farm control using combined control strategies

Department of Wind Energy Master Report

Charles Marie J Debusscher

DTU Wind Energy-M-0411

July 2020

DTU Wind Energy
Department of Wind Energy

Author: Charles Marie J Debusscher

Title: Surrogate-based optimization under uncertainty of wind farm control using combined control strategies

DTU Wind Energy-M-0411
July 2020

Project period:
November 2019 – July 2020

ECTS: 45

Education: Master of Science

Supervisors:
Tuhfe Göçmen
Søren Juhl Andersen
DTU Wind Energy

Carlos Simão Ferreira
Jan-Willem van Wingerden
TU Delft

Remarks:
This report is submitted as partial fulfillment of the requirements for graduation in the above education at the Technical University of Denmark.

DTU Wind Energy is a department of the Technical University of Denmark with a unique integration of research, education, innovation and public/private sector consulting in the field of wind energy. Our activities develop new opportunities and technology for the global and Danish exploitation of wind energy. Research focuses on key technical-scientific fields, which are central for the development, innovation and use of wind energy and provides the basis for advanced education at the education.

We have more than 240 staff members of which approximately 60 are PhD students. Research is conducted within nine research programmes organized into three main topics: Wind energy systems, Wind turbine technology and Basics for wind energy.

Technical University of Denmark
Department of Wind Energy
Frederiksborgvej 399
4000 Roskilde
Denmark
www.vindenergi.dtu.dk

Abstract

Due to the limited availability of sites on both land and sea, there is a need to maximize the power density of wind farms whilst limiting the adverse wake effect between the wind turbines that cause power losses and/or increased load cases. The study focuses on analyzing the potential of combining yaw-based wake-steering and constant blade pitch control down-regulation to mitigate the detrimental wind turbine wake interactions. Data-driven surrogate models based upon polynomial chaos theory have been used to model the statistical distributions of the wind farm power production and short-term Damage Equivalent Loads (DEL) in function of the control inputs. The model calibration data has been generated for a range of control settings for two V27 wind turbines aligned with the wind direction through large eddy simulations using the flowsolver EllipSys3D and the aeroelastic code Flex5. A power-based optimization with DEL constraints has been performed and median power gains ranging from +1% to +3% have been observed at close spacings depending on the severity of the imposed constraints. It was identified that larger power gains corresponded to an increase in DEL. At larger spacings, the wind farm control strategies shows limited performance increases and revert to baseline operations.

Preface

This thesis was prepared in fulfillment of the requirements for acquiring a MSc degree in Aerospace Engineering from the Delft University of Technology and a MSc degree in Wind Energy from the Technical University of Denmark. The project was carried out from Novermber 2019 to July 2020 under the supervision of Søren Juhl Andersen (DTU), Tuhfe Göçmen (DTU), Carlos Simão Ferreira (TU Delft) and Jan-Willem van Wingerden (TU Delft). This double degree joint education program is part of the European Wind Energy Energy Master (EWEM).

DTU, Kongens Lyngby, July 30, 2020

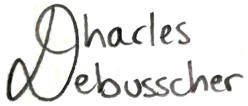

Charles Marie J Debusscher

Acknowledgements

I would like to express me deep and sincere gratitude to my supervisors, without whom the completion of this thesis would not have been possible. Our (regularly prolonged) weekly meetings have been an immense source of inspiration and motivation, however challenging, virtual or noisy they might have been at times.

In no particular order, I would like to thank Søren Juhl Andersen for always being available to answer my questions, whether it be over rubber duck debugging or the complex intricacies of wind farm aerodynamics and modelling. I am also thankful for Tuhfe Göçmen for all the clever insights and her patience when repeatedly and creatively having to explain those insights to me. I would like to express my gratitude to Carlos Simão Ferreira for sparking my interest for the field of wind farm aerodynamics by being an amazing teacher and for always putting things into perspective in such that I would not lose sight of the bigger picture.

I am very grateful for my friends, Casper, Mario and Kamran for the support and many (un)intelligent conversations we have had during this journey. A great thank you to everyone involved in the organisation of the EWEM programme, whether it be at the Delft University of Technology or the Technical University of Denmark, for having made this unforgettable experience possible.

Finally, I am eternally thankful to my family for all the named and unnamed sacrifices that they have made to allow me to pursue my education and prepare myself for the future.

Contents

List of Figures

List of Tables

Nomenclature

List of Abbreviations

(C)CW: *(Counter) ClockWise*

ABL: *Atmospheric Boundary Layer*

ADM: *Actuator Disk Method*

ALM: *Actuator Line Method*

ANN: *Artificial Neural Network*

CDF: *Cumulative distribution function*

CFD: *Computational Fluid Dynamics*

DEL: *short-term Damage-Equivalent Load*

DNS: *Direct Numerical Simulation*

FAST: *Fatigue, Aerodynamics, Structures, and Turbulence*

FLORIS: *FLOw Redirection In Steady-state*

KDE: Kernel Density Estimate

LES: *Large-Eddy simulation*

MLE: *Maximum Likelihood Estimate*

NREL: *National Renewable Energy Laboratory*

PCE: *Polynomial Chaos Expansion*

PDF: *Probability Density Function*

QRS: *Quadratic Response Surface*

RANS: *Reynolds-averaged Navier–Stokes*

SOWFA: *Simulator for Offshore Wind Farm Applications*

List of Symbols

α_{turb}	Kolmogorov constant
ϵ	Specific turbulent kinetic energy of the system
Γ	Coefficient reflecting the anisotropic character of the turbulence field
γ	Rotor rotational position
λ	Tip-speed-ratio
ν	kinematic
ν_{SGS}	Subgrid-scale eddy viscosity
ω	Rotor rotational speed
ψ	Nacelle yaw angle
ψ_1	Upstream wind turbine yaw input
ψ_2	Downstream wind turbine blade pitch input
ρ	Ambient air density
θ	Blade pitch angle
θ_1	Upstream wind turbine yaw input
θ_2	Downstream wind turbine blade pitch input
A	Weibull distribution scale parameter
a	Axial induction factor of the wind turbine rotor
A_d	Rotor swept area
C_P	Wind turbine rotor power coefficient
C_T	Wind turbine rotor thrust coefficient
f_i	Density-normalized body forces
H_{hub}	Hub height
k	Weibull distribution shape parameter
L_{turb}	Turbulent length scale
M_{BRE}	Blade Root Edgewise bending moment

M_{BRF} Blade Root Flapwise bending moment

M_{MS} Main Shaft torque

M_{TB} Tower Bottom bending moment

M_{TT} Tower Top bending moment

M_{yaw} Tower top torque/yaw moment

p Pressure

Q Generator torque

Q_X X^{th} distribution percentile

R Rotor radius

S_{ij} Strain rate tensor

S_X Wind turbine spacing

T Rotor thrust

t Temporal dimension

T_{ij} Subgrid scale tensor

u Wind velocity

U_∞ Free-stream wind speed

U_{wake} Axial wake velocity far downstream of wind turbine

w_0 Double Weibull distribution relative weight coefficient

x Position vector in Cartesian coordinates

CHAPTER 1

Introduction

The global installed capacity of wind energy has more than tripled during the last decade (IRENA, 2019b), as shown in Figure 1.1. It is unrealistic to assume that this rising trend can be maintained solely by allocating more and more land to wind energy plants. As a result, there is a rising need to maximize the power production of (new *and* existing) wind farm per unit of area, i.e. to increase their *power density*. This goal is partially achieved by the increasing size and efficiency of modern-day commercially available wind turbines. As of the day of writing, the world first 12MW wind turbine has been constructed[1], and due to the ongoing innovations and technological advancements, turbine rotor diameters are excepted to keep increasing (IRENA, 2019a).

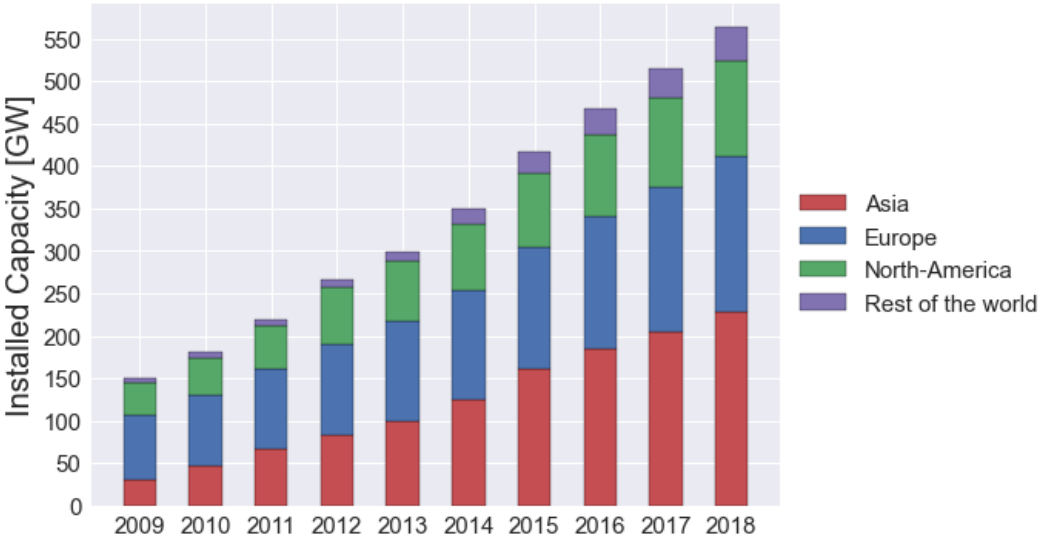

Figure 1.1: Global installed capacity trends within the wind energy sector from 2009 to 2018. Data has been retrieved from the International Renewable Energy Agency (IRENA, 2019b).

A first approach consists in maximizing the power production of each wind turbine individually, without taking the neighboring wind turbines into consideration. This approach is often referred to as the greedy-control strategy (Boersma et al., 2017; Kheirabadi & Nagamune, 2019). The downside of this approach is that the wind turbines may experience adverse aerodynamic effects on each other when the mean wind direction is aligned with the lay-out of a wind farm. The origins of these adverse aerodynamic effects lay with the creation of a wake by the upstream wind turbine. The viscous

[1] https://www.ge.com/renewableenergy/wind-energy/offshore-wind/haliade-x-offshore-turbine

interactions between the upstream wind turbine and the wind flowing through its rotor area create a region of air with an increased turbulence intensity and reduced velocity that propagates downstream (Lissaman, 1979). That region of air is referred to as the *wake* of the wind turbine. The wind turbines operating within the wakes of the upstream wind turbine will experience an increase in power losses and fatigue loads (Ainslie, 1988). For larger offshore wind farms, power losses due to the adverse aerodynamic interactions between the wind turbines have been estimated to run up to 10-20% of the total installed capacity (Barthelmie et al., 2007).

A more holistic approach can be taken by considering the control of wind farm as a collaborative group of wind turbines, rather than only focusing on the individual wind turbines it is made of. By limiting the adverse aerodynamic interactions between the wind turbines, the total generated power of a wind farm can be increased and fatigue loads minimized. Three main approaches are described by Kheirabadi and Nagamune (2019) to mitigate these adverse wake effects:

- **Spacing:** The inter-turbine distance is made sufficient such that the wake effects are limited. However, this approach violates the power density concept.
- **Lay-out optimization:** For a given site, historical wind data can be used to determine a wind farm lay-out that minimizes the aerodynamic interactions.
- **Wind farm flow control:** By means of the the control degrees of freedom of the individual wind turbines, the characteristics and propagation of the wake can be controlled to some extent.

The advantage of the latter of the three approaches is that it can be used for existing *and* future wind plants. For existing wind plants, the production could potentially be optimized by implementing wind farm control strategies within the existing control loop, limiting the power losses and increased fatigue loads occurring due to turbine-wake interactions. For future wind farms, their lay-out could be optimized as wind farm control strategies may allow for smaller inter-turbine distance, therefore increasing the power density. The two most prominent wind farm flow control strategies in the current literature are axial induction control and wake-redirection control (Boersma et al., 2017; Kheirabadi & Nagamune, 2019) and are discussed in the next section.

1.1 Introduction to Wind Farm Control Strategies

Axial induction control, introduced by Steinbuch et al. (1988), consists in reducing the moment extraction from the upstream wind turbine, therefore reducing its power production. As a result, the mean velocity deficit within its wake is reduced and the downstream wind turbine will operate in more favorable wind conditions (Johnson & Thomas, 2009). Ideally, the power production of the downstream turbine will increase such that the power losses of the upstream turbine are compensated for. One approach to axial induction control consists in reducing the thrust of the upstream turbine, as shown in Figure 1.2. The reduction in thrust is achieved by fine-tuning the blade pitch angle

and/or tip-speed ratio of the rotor away from their optimal design point, such that the thrust of the rotor is reduced and the amount of kinetic energy extracted from the free wind is reduced (Munters & Meyers, 2018a). A further benefit is the potential decreased turbulence within the wake, which could lead to a reducing the fatigue loads experiences by the downstream turbine (Dilip & Porté-Agel, 2017). A wind turbine operating in this state is referred to as *derated* or *down-regulated*. Similarly, wind turbines operating in above-rated wind conditions are down-regulated in order to limit the power production and avoid compromising the turbine components (S. Dhiman & Deb, 2020).

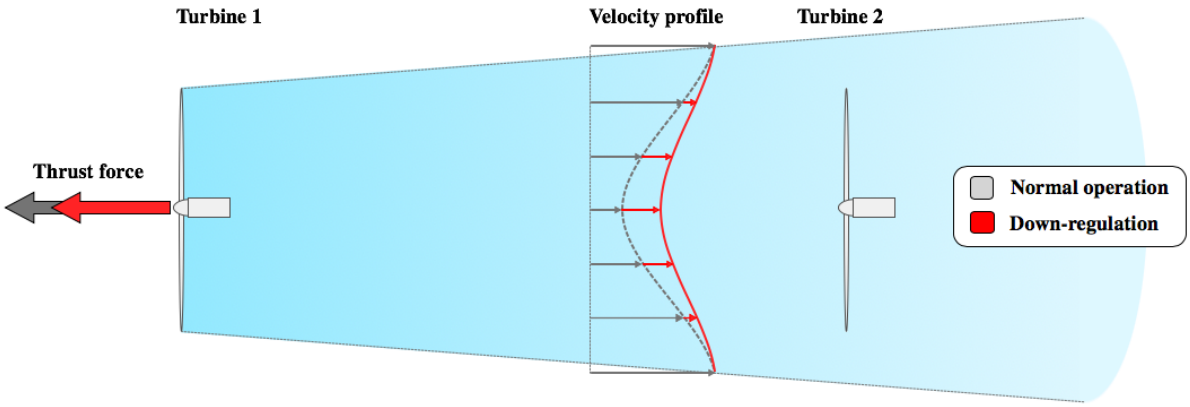

Figure 1.2: Schematic top-view of the mechanisms of axial induction control. The *red* elements indicate the operational changes of the turbines and their consequences on the wind flow. The thrust force of the upstream turbine is reduced, resulting in an increased mean wind speed in the wake.

The second approach consists in altering the direction of the thrust force of the upstream turbine with respect to the wind flow direction such that the wake can be redirected away from the downstream wind turbine. The mechanism of yaw-based wake-redirection is shown in Figure 1.3, where the nacelle yaw angle is misaligned with the incoming wind direction. As a result, a component of the thrust force acts in cross-wind direction and causes a deflection in the wake propagation direction (Bastankhah & Porté-Agel, 2016). As the wake is deflected away from the rotor of the downstream wind turbine, its power output will be optimized. A further benefit of the reduced wake effects its the reduction in fatigue loads, as the mean turbulence intensity encountered by the downstream wind turbine is decreased. However, the upstream wind turbine could experience an increase in fatigue loads due to the asymmetrical loading of its rotor (Zalkind & Pao, 2016).

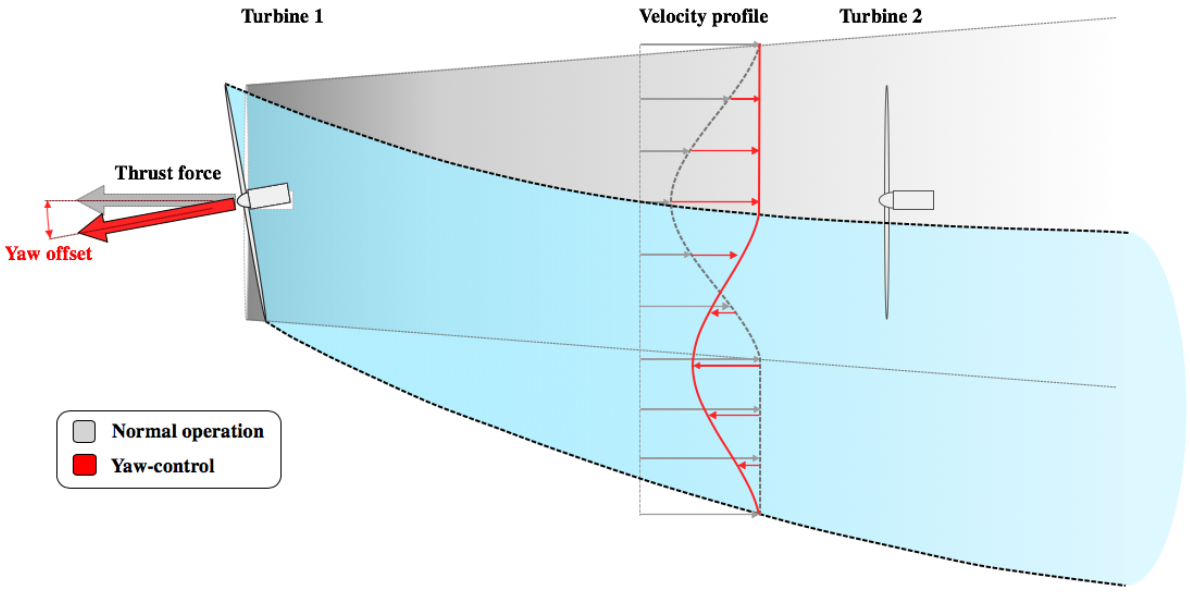

Figure 1.3: Schematic top-view of the mechanisms of yaw-based wake-redirection. The red elements indicate the operational changes of the turbines and their consequences on the wind flow. The thrust force of the upstream turbine is misaligned with the wind direction, steering the wake away from the downstream turbine.

1.2 Aim of Thesis & Contents

The objective of this thesis is to investigate the potential of wind farm control strategies to mitigate adverse wake effects in terms of power loss and fatigue loads by means of data-driven surrogate models. This work can be considered as a continuation of the work done in Hulsman et al. (2019). The novelty aspect of this study is twofold:

1. Both yaw-based wake-steering and down-regulation have proven to have the potential to control the power and loads within wind farms to some extent (Kheirabadi & Nagamune, 2019). However, both strategies come with their advantages and disadvantages. As a result, this study aims to combined both control strategies in order get the best of both worlds.
2. Due to the complex and the non-deterministic nature of the interactions between the wind turbines and the incoming wind, the potential of the different control strategies and their combinations are evaluated using statistical analyses. By means of fitting distributions to the power and loads outputs of the wind turbines for a given set of wind farm control inputs, the expected performance and the associated uncertainty can be determined. This allows to conducted an optimization of the individual and combined control strategies of the wind farm control inputs.

The study is conducted by means of large eddy simulations using the flowsolver EllipSys3D (J. A. Michelsen, 1992, 1994; N. Sørensen, 1995a) coupled with the aeroelastic code Flex5

(Oye, 1996) to simulate the flowfield and obtain the wind turbine power production and loads. The used wind turbine model is a V27. The simulations are run for a range of yawed or down-regulated cases, and their combinations, after which the flowfield is extracted at different downstream locations. By means of flow data extracted at each location, the potential power production and loads of a wind turbine placed at that location is determined using the Flex5 aeroelastic code.

Using the output obtained from Flex5, Polynomial Chaos Expansion (PCE) based surrogate models are constructed for the wind farm. The surrogate models are calibrated to predicte the distribution of the 10-minute statistics for a given set of wind farm inputs. By means of the constructed models, optimization under uncertainty of the input variables can be performed with the purposes of maximizing the wind farm power output and limiting the loads experienced by the wind turbines.

The following chapter, Chapter 2, aims to give the reader background on the topic of wind farm control and give a review of the current state-of-the-art. By means of the current literature, the different types of wind farm control models and strategies are analysed and compared, such that the placement of this study within the larger picture is made clear. Chapter 3 provides a description of the used methodology regarding the simulation set-up, the generation and processing of data, and the construction of the surrogates. The results are shown in Chapter 4 and are separated into two parts. In the first part, the performance of the surrogate model is evaluated in order to understand the limitations of the constructed models. In the second part, the surrogate models are used to perform optimization of the wind farm input parameters. In the last chapter, Chapter 5, the main conclusions are summarized and some recommendations are presented.

CHAPTER 2
Background & State-of-the-art

This chapter provides an overview of the field and state-of-the-art of wind farm flow control in order to provide a background and motivation for the thesis. In the first two sections, Section 2.1 & Section 2.2, basic wind turbine aerodynamics and control are discussed. The following section, Section 2.3, describes the wake of a wind turbine and is evolution as it propagates downstream. In Section 2.4, a review of wind farm flow models is presented and in the following section, Section 2.5, the state-of-the-art of the major control strategies is discussed. In the last section, Section 2.6, the main conclusions form the chapter are summarised and the motivation for the thesis is stated.

2.1 Wind Turbine Aerodynamics

The volume of air passing through the rotor of a wind turbine can be described by considering an actuator disk exerting a force field on the flow and the momentum theory within a streamtube. A streamtube, shown in Figure 2.1, is defined by the boundary between the air that will pass through the rotor plane and the air that will go around it (Burton et al., 2011). Steady, incompressible, inviscid flow is considered and it is assumed that no mass, momentum nor energy is exchange across streamtube's boundary, i.e. the conservation laws are respected within the streamtube.

Far upstream of the actuator disk, the wind flow is unaffected by the actuator disk and the flow parameters are undisturbed. Closer to the actuator disk, the wind slows down due to the presence of the actuator disk itself. This has two consequences. Firstly, due to conservation of energy, the pressure increases in order to compensate for the reduction

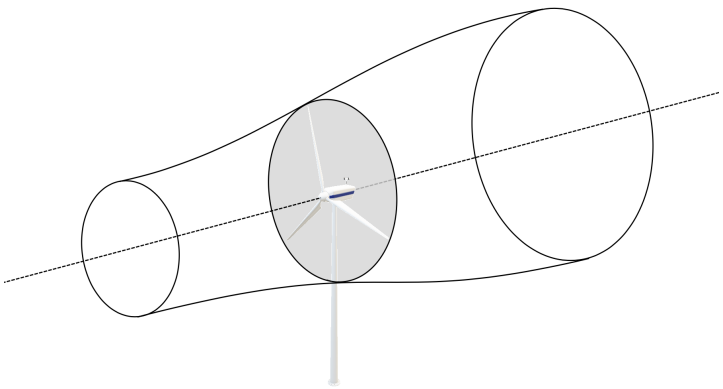

Figure 2.1: Visualisation of the streamtube passing through the rotor area.

in kinetic energy, as no other work has been performed on the airflow yet. Secondly, as the mass flow rate within the streamtube has to remain constant, its diameter increases as it nears the actuator disk. At the location of the actuator disk, work is performed by the actuator disk on the air, resulting in a sudden drop in pressure.

The air behind the actuator disk proceeds downstream with a decreased pressure and velocity. The pressure gradually recovers back the atmospheric pressure levels, which results in a further decrease in wind velocity and increase in streamtube diameter. Far downstream of the actuator disk, the pressure levels are restored to atmospheric levels. The flow with decreased pressure and velocity behind the actuator disk is referred to as the *wake*, and is discussed in further details in the next sections. For a more extensive and quantitative analysis of momentum method, the reader is referred to Snel (1998) and Burton et al. (2011).

Common measures for the performance of a wind turbine are the *thrust & power coefficients* , defined by,

$$C_P = \frac{TU_\infty}{\frac{1}{2}\rho A_d U_\infty^3} = 4a(1-a)^2, \tag{2.1}$$

$$C_T = \frac{T}{\frac{1}{2}\rho A_d U_\infty^2} = 4a(1-a), \tag{2.2}$$

where T is the thrust of the wind turbine, U_∞ is the free-stream wind speed, ρ is the ambient air density, A_d is the rotor swept area and a is the axial induction factor of the rotor. The axial induction factor is defined by,

$$a = \frac{U_\infty - U_{wake}}{2U_\infty}, \tag{2.3}$$

where U_{wake} is the axial velocity of the wake far downstream of the wind turbine. By combining Equation (2.1) and (2.2) with Equation (2.3), it can be observed that both the thrust and power coefficient are fully dependent on the upstream and downstream velocity within the streamtube. The performance of the rotor is thus directly related to the amount of kinetic energy extracted from the free-stream wind flow to the wind turbine rotor. The axial induction parameter has shown to be a crucial parameter to model the induction zone in front of the rotor (Troldborg & Meyer Forsting, 2017), the wake characteristics (Bastankhah & Porté-Agel, 2014) and wind farm power estimation (Johnson & Thomas, 2009).

In order for momentum theory the hold however, the axial induction factor a has to remain within $0 \leq a < \frac{1}{2}$, as for larger values of a the velocity in the wake would theoretically reache zero or negative values. For large values of a, empirical relations such as Glauert's empirical relationship are required (Burton et al., 2011).

2.2 Wind Turbine Control

Most commercially available wind turbines feature three degrees of freedom to optimize their performance during day-to-day operations (Boersma et al., 2017):

- **Blade pitch angle θ** : The pitch angle is defined as the angle between the tip of the blade chord line and the rotor rotational plane. Its magnitude can be adjusted by rotating the blade around its longitudinal axis, altering the force distribution along the blade span and directly influencing the rotor rotational velocity and rotor torque. Blade pitch control is used to limit the power generation in above rated conditions by reducing the aerodynamic efficiency of the rotor.
- **Generator torque Q** : The generator torque is adjusted to control or optimize the power capture in below and above rated conditions respectively. The discrepancy between the aerodynamic and generator torque will cause a deceleration or acceleration of the rotor.
- **Nacelle yaw angle ψ** : The nacelle yaw angle is defined as the angle between the mean flow direction and the nacelle axial direction in the horizontal plane. It is adjusted by rotating the nacelle around the wind turbine tower in order to ensure that the rotor is perpendicular to the flow, maximizing the power capture.

Pitch regulated variable speed wind turbines regulate their power production by adjusting the rotor rotational speed ω by means of the pitch controller and torque controller (Bossanyi, 2000). A simplified control diagram is shown in Figure 2.2, where both pitch and torque controller can be observed. In below-rated conditions, the power production is mainly regulated by the torque controller. In those conditions, the pitch is fixed at the

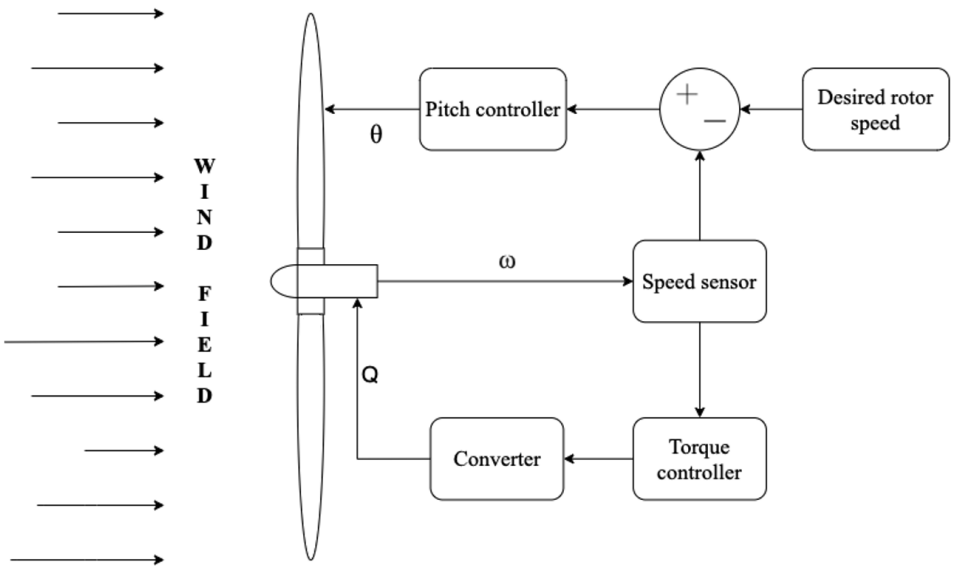

Figure 2.2: Conceptual control diagram of a pitch regulated variable speed wind turbine.

optimal pitch angle such that the design tip-speed-ratio is maintained, maximizing the power production. The tip-speed-ratio λ is defined by $\lambda = \frac{\omega R}{U_\infty}$, which corresponds to the ratio of the rotor tip speed and the free-stream wind speed.

Once above-rated conditions are reached, the power regulation is mainly taken over by the pitch controller. The pitch is adjusted away from its optimal point such that the rated torque is maintained. The action of deliberately limiting the power production of a wind turbine by means of a pitch controller is referred to as wind turbine *derating* or *down-regulation*, and will be discussed more in depth in Section 2.5. For further reading about wind turbine controllers, the reader is referred to Bossanyi (2000) and Pao and Johnson (2009).

Wind turbines and wind farms are connected to a larger network referred to as the grid. As a result, at times the power produced by the wind plant has to be controlled to ensure that the grid remains balanced. Active power control functions are used to limit or control the power output of given power plant. *Absolute power limitation* consists in applying a a hard upper boundary for the power production over a given time interval. Other approaches, such as *delta control*, apply a down-regulation based on the instantaneous power output of the wind plant (Elorza et al., 2019). This can be done so by maintaining a constant amount of reserve power (i.e. constant delta control) or having the amount of reserve power size proportionally with the possible power. The principle of the three active power control functions are shown in Figure 2.3. These three control strategies require to continuously modify the above described wind turbine degrees of freedom to produce to requested down-regulated power. In this study, wind turbines are down-regulated by means of *constant pitch control*, which to some extent is a form of proportional delta control. The pitch angle is increased away from its design set-point and is kept constant during the simulation time. Other power control function exist, such as *Power Balance control* and *power rate limitation*, and are discussed in more elaborately in Kristoffersen (2005).

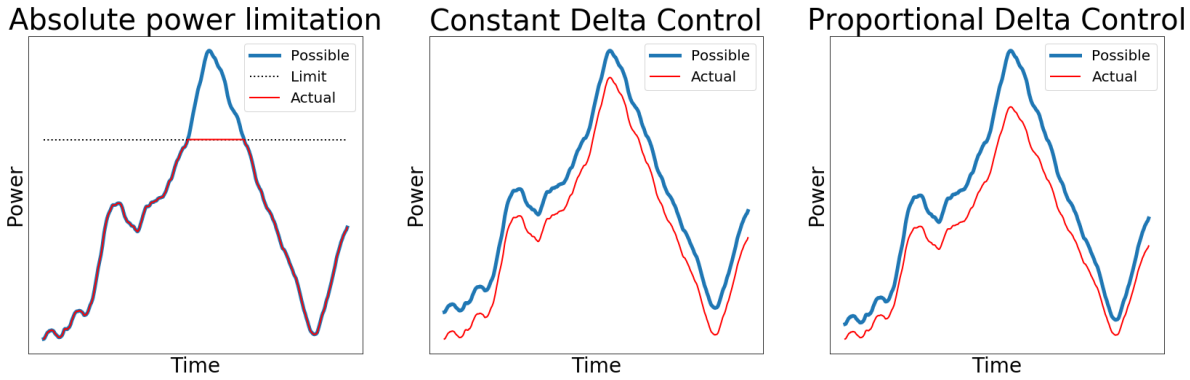

Figure 2.3: Schematic power curves of active power control functions. The blue line represents the possible achievable power and the red line represent the actual produced power when an active power control function is applied.

2.3 Wind Turbine Wake

Due to the viscous interactions between the wind turbine and the wind flowing through its rotor area, a volume of air with a reduced mean velocity and increased turbulence intensity is created that propagates downstream with the mean flow direction (Lissaman, 1979). This volume of air, referred to as the *wake* of the wind turbine, can be divided into two distinct regions with different characteristics depending on the downstream distance. Vermeer et al. (2003) defines these two regions as the *near*- and *far* wake regions, as depicted in Figure 2.4.

The near wake, located directly behind the rotor, is a region of reduced velocity of which the diameter slightly exceeds the rotor diameter. The shape and characteristics of the near wake are mainly determined by the rotor geometry, such as its diameter, rotational velocity and blade specifications (Vermeer et al., 2003). The distribution of the velocity deficits within the near wake can be modelled using a top-hat (Jensen, 1983) or double Gaussian distribution (Bastankhah & Porté-Agel, 2014). A maximal velocity deficit is reached around 1 or 2 diameters downstream (Ainslie, 1988). Thickening shear layers are formed between the wake and surrounding air due to the velocity and pressure gradient, creating turbulent eddies, or *turbulence* (Sanderse, 2009). The turbulent eddies are convected downstream and the thickness of shear layers increases due to turbulent dissipation, until the center of the wake is reached at around 2 to 5 diameters (Crespo et al., 1988).

Around the points where the shear layers reaches the core of the wake, the far wake starts and the wake characteristics are dictated by the ambient turbulence intensity and convection (Vermeer et al., 2003). The wake deficit within the far wake follows a

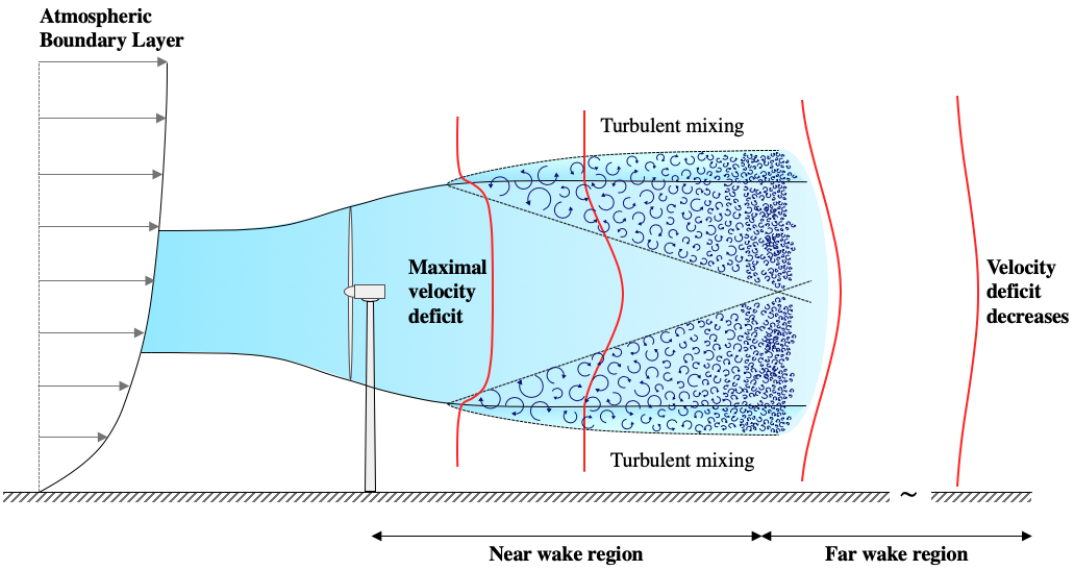

Figure 2.4: Schematic side-view of the wake created by a wind turbine. Inspired by Sanderse (2009).

self-similar axisymmetric profile with a Gaussian distribution, of which the magnitude decreases with downstream distance as the wake recovers (Burton et al., 2011). The control setting of the rotor, e.g. tip-speed-ratio and blade pitch angle, have a direct influence on the velocity deficit, as its magnitude is determined by the thrust coefficient of the rotor (Ainslie, 1988). As the downstream distance increases, the velocity deficit reduces due to turbulent dissipation and the wake homogenizes with the ambient air.

2.4 Wind Farm Flow Modelling

Conducting in-field experiments within wind farms can be a complex, expensive and time-demanding task. Wind farm flow models offer a less costly approach to modelling the operations and efficiency of the wind farm, as preliminary results can be obtained with less infrastructure and on smaller time scales (Boersma et al., 2017). A downside of using models is that they need to be verified and validated, and results are only a representation of the true physical system. The extent to which a model represents the states of real-world physical system accurately is defined as the *fidelity* of the model. The fidelity-spectrum can be roughly divided into *low*-, *medium*- and *high*-fidelity bands (Kheirabadi & Nagamune, 2019). High-fidelity simulations have considerable higher computational cost than low-fidelity models. While most low-fidelity models can be run in a handful of seconds on a present-day desktop PC, high fidelity simulations can require up to several days on cluster computers (Boersma et al., 2017). The different wind farm flow models are discussed in the following sections based upon their fidelity level.

2.4.1 Low-fidelity

Low fidelity models, also referred to as *parametric* or *engineering* models, are based on the integral relations in fluid dynamics in which the change in mass and momentum are conserved within the considered volume. The main advantage of these models are the low computational cost (Annoni et al., 2014). However, this comes at the price of incomplete system representations, as the velocity and pressure gradients can not be determined straight from the model integral equations (Emanuel, 2015). As a result, they have to be included using parametric assumptions of the wake characteristics. These state variables are fundamental to describe complex phenomena including wake expansion, turbulent mixing and ground effect (Kheirabadi & Nagamune, 2019). In order to obtain realistic results, the low-fidelity models require to be calibrated and/or tuned using measurements (Knudsen & Bak, 2012) or high-fidelity simulations (Gebraad et al., 2014). The Jensen wake model is an example of one of the earliest parametric models (Jensen, 1983).

The Jensen model assumes ideal axially symmetric flow with no rotation nor turbulence. A constant value is used to model the wake expansion in function of the downstream distance (Göçmen et al., 2016). The Park wake model expansion, based on Katic et al. (1986), allows for the determination of the total kinetic energy deficit of multiple wakes by setting the total energy deficit equal to the root-sum-square of individual wake energy deficits (Peña & Rathmann, 2014). The effect of the nacelle yaw was added to the

model in Gebraad et al. (2016) by using the analytical derivation from Jiménez et al. (2010). Concentric zones within the wakes were created where the wake characteristics could vary in order to increase the accuracy of the velocity deficit distributing. The tuning parameters were calibrated using high fidelity simulations to create the FLOw Redirection In Steady-state (FLORIS) model which is made freely available by NREL (NREL, 2019). This model has been further extended to create the quasi-dynamic wake model FLORIDyn developed in Gebraad and Van Wingerden (2014), where the effects of a change in operational setting of a turbine is propagated downstream in a step-wise manner. As a result, a rough boundary between pre-change and aft-change is advancing downstream with the wake propagation velocity with each time step.

2.4.2 Medium- & High-fidelity

High-fidelity models are based upon differential form of the Navies-Stokes equations, wherein mass and momentum are conserved (Kheirabadi & Nagamune, 2019). These methods are base upon Computational Fluid Dynamics (CFD). Three approaches, ordered by their fidelity level, are discussed in this section. The approaches are Reynolds-averaged Navier–Stokes (RANS), Large-Eddy simulations (LES) and Direct Numerical Simulations (DNS).

2.4.2.1 Reynolds-averaged Navier–Stokes

The RANS models are based upon the Reynolds decomposition of the Navier-Stokes equations. In these equations, the flow variables are separated into their time-averaged and fluctuating components. Three approaches are considered for the RANS models, namely linearized, parabolized and full original elliptic modelling (Sanderse et al., 2011). The linearized modeling is the fastest out of the three approaches and can be considered a mid-fidelity model.

2.4.2.2 Large Eddy Simulations

In LES, the filtered Navies-Stokes are solved on a grid level, while on a sub-grid level a subgrid-scale model is used, e.g. Smagorinsky model, Dynamic model, amongst others (Vreman, 2004). The advantage of high-fidelity models over the low-fidelity models is that the velocity and pressure gradients are evaluated, meaning that no assumption describing the main wake characteristics are required (Emanuel, 2015) and flows governed by turbulent mixing and large-scale eddies can readily be modelled (Sanderse, 2009). The downside is that the increase in fidelity comes at the cost of an increase in computational cost. Popular LES flow-solvers include amongst others SOWFA from NREL (Churchfield et al., 2012), SP-Wind (Allaerts & Meyers, 2015), UTD-WF from the University of Texas at Dallas (Martínez-Tossas et al., 2015) and EllipSys3D from the Risø National Laboratory (N. Sørensen, 1995b). The latter, EllipSys3D, includes both RANS and LES models (Sanderse et al., 2011).

If no turbulence model is used and all spatial and temporal ranges of turbulence are solved directly from the Navies-Stokes equations, the CFD model is referred to a DNS. These models have the highest fidelity level, but require a massive amount of computational power (Sanderse et al., 2011). According to Churchfield et al. (2015), modelling a wind farm within a $1km \times 1km \times 1km$ domain would require 10^{18} grid cells, exceeding the current available computational power.

2.4.3 Data-driven Models

An alternative approach to the physical models described above are the data-driven models, which are calibrated using data generated by high-fidelity simulations (Gebraad et al., 2014) or measurements (Japar et al., 2014). Once calibrated, data-driven models allow for a fast and reliable determination of the physical quantities of interest, whilst also ideally maintaining (to some extent) the accuracy of the high-fidelity models. Hence, data-driven models have the potential to take the best of both worlds considering the advantages of the low- and high-fidelity models. However, the quality of the output of the data-driven surrogate models relies heavily on the quality and quantity of input data used to calibrate the model (Gebraad et al., 2014; Hulsman et al., 2019).

Dimitrov et al. (2018) used different surrogate methods including Polynomial Chaos Expansion (PCE), Quadratic response surfaces (QRS), universal Kriging, importance sampling, and nearest-neighbor interpolation to evaluate the fatigue loads of wind turbines. Although the Kriging method had a slight advantage in accuracy over PCE method, it required additional computational time compared to the other methods. The PCE method appeared to be the optimal method when considering both evaluation time and accuracy. In Schröder et al. (2018), site-specific fatigue loads were determined using data-driven surrogate models based on Artificial Neural Networks (ANN), QRS and PCE. The QRS approach appears requires the less evaluation time, but lacked in accuracy compared to the two other approaches. The ANN model slightly outperformed the PCE model in terms of performance error and evaluation time. A similar conclusion was drawn in Dimitrov (2019) for the PCE and ANN models. It was observed that the ANN model surpassed the performance of the PCE model in terms of accuracy and computational cost. However, the author emphasised that the ANN models are more sensitive to overfitting and require a larger data set than for PCE-based models. In (Hulsman et al., 2019) PCE-based data-driven models were used to optimize the fatigue loads and power production of two wins turbines. Due to the computational speed of the calibrated models, they can be used as optimization tools for given wind turbine parameters.

2.5 Wind Farm Flow Control Strategies

In order to regulate the power production of the wind farm, facilitate the integration of wind energy within the electrical grid and/or minimize the structural loads of the turbines, the wind farm is equipped with a wind farm controller (Knudsen et al., 2015).

A simplified schematic of the closed-loop control system of a wind farm is shown in Figure 2.5. A certain power demand is communicated from the grid operator to the wind farm controller. The wind farm controller sets the operational set-points of the wind turbines with the purpose to achieve a certain reference power. The wind farm will interact with is environment and the resulting power, loads and operational settings are measured. These signals are sent back to the wind farm controller, which resets the turbines set-points and communicates the active power with the grid operator.

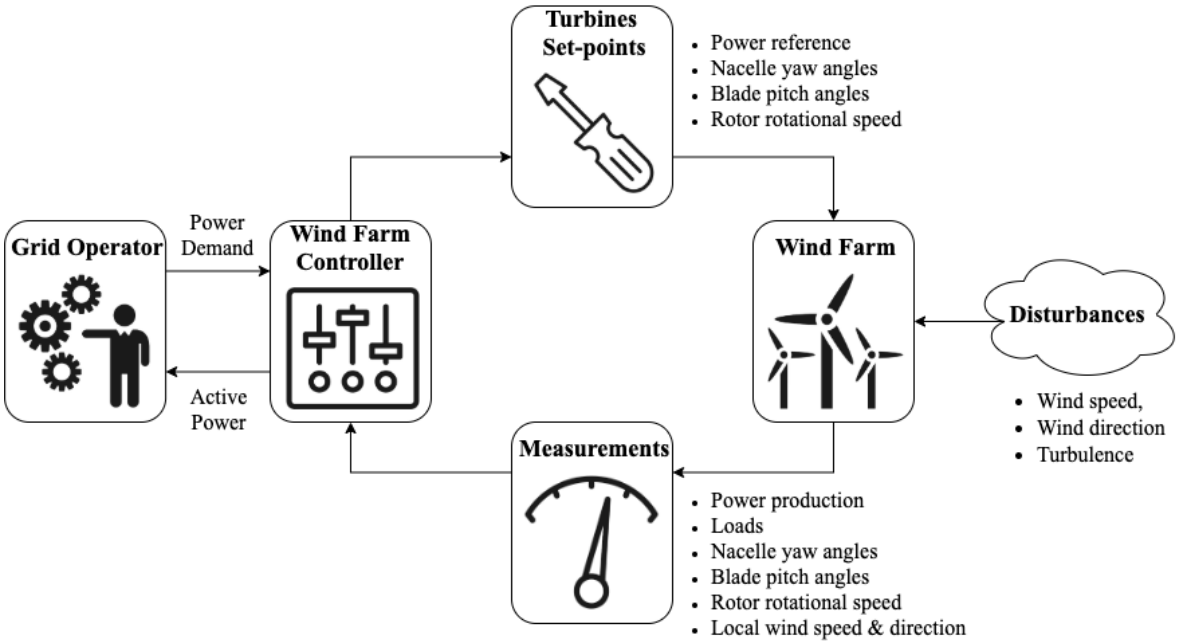

Figure 2.5: Simplified schematic of the closed-loop control of wind farms.

Each individual component of the schematic shown in Figure 2.5 is its own field of study, however the focus of this work lays with the determination of the optimal turbines set-points and their effects on the power and loads of the wind farm. The state-of-the-art of the different approaches are discussed within this section, with a focus on static axial induction control and yaw-based wake-steering. For other wake-steering approaches, such as tilt-based wake-redirection and individual pitch control, the reader is referred to (P. A. Fleming et al., 2014, 2015).

2.5.1 Down-regulation

As previously introduced in Section 1.1 and depicted in Figure 1.2, down-regulation consists in purposefully reducing the thrust of the upstream turbine, such that the downstream turbine can operate in more favorable conditions. The expectation is that the increased power production of the downstream wind turbine makes up for the power loss of the upstream turbine (Johnson & Thomas, 2009). The down-regulated upstream wind turbine is run into sub-optimal conditions by altering the collective blade pitch

angle and tip-speed-ratio away from its design point. The majority of the studies focus on static set-points, as reviewed in Boersma et al. (2017), Knudsen et al. (2015), however dynamic approaches have been the subjects of studies too. Munters and Meyers (2018c) observed that by introducing sinusoidally varying thrust in the first row of turbines a substantial increase in total wind farm power was achieved. The quantification of the benefits of the wind farm flow control strategies are case and method dependant, as a discrepancies exist between the results of the models and physical experiments. The following sections review the existing literature that investigates axial-induction control as a control strategy using different types of models.

2.5.1.1 Low- to Mid-fidelity Models

The majority of studies evaluating the steady-state potential of axial induction control have been conducted using low to medium fidelity models with varying optimization schemes and wind turbine control inputs (Kheirabadi & Nagamune, 2019). In Tian et al. (2014), power gains up to +26% were observed using the steady Park model by optimizing the blade pitch angle and tip-speed ratio of ten aligned wind turbines under constant speed conditions. In contrast, Herp et al. (2015) observed a potential power gain between +1% and +5% using the steady Park model to optimize the axial induction factors within the Nysted offshore wind farm. Although both studies did not measure any efficiency loss, the large spread between the two results indicate a high variability of the predictions of low-fidelity models for axial induction control.

2.5.1.2 Mid- to High-fidelity Models

Annoni et al. (2016) compared low-fidelity FLORIS model with the high-fidelity SOWFA model in terms of the summed power gains of two aligned wind turbines. While the low-fidelity FLORIS model predicted gains up to nearly +25%, the high-fidelity SOWFA model predicted a gain loss of more than −9%. According to (Annoni et al., 2016), this large discrepancy can be partly attributed to the following two points. Firstly, as the thrust of the upstream wind turbine is reduced, the momentum recovery and wake rate of expansion are reduced, which the FLORIS model fails to model. Secondly, the additional momentum within the wake due to the down-regulated upstream turbine lays mostly within the wake boundary, which does not pass through the swept area of the downstream wind turbine. The FLORIS models assumes a uniform distribution within its wake zones, incorrectly increasing the momentum capture of the downstream turbine. Using the finite volume model developed in Soleimanzadeh et al. (2010), a reduction in fatigue loads of the upstream down-regulated turbines was observed (Soleimanzadeh et al., 2012). The downstream wind turbine experienced an increased in fatigue loads, as it is subjected to a higher momentum wake.

2.5.1.3 Wind Tunnel & Field Experiments

Power gains ranging ranging from +3 to +4% were observed by Adaramola and Krogstad (2011) using two aligned wind turbines within a wind tunnel by derating the upstream

turbine by means of the blade pitch angles and the tip-speed-ratio. Due to the lower scales of the experiment compared to field experiments, the wind turbines are operating under lower Reynolds numbers and the wake rate of expansion is reduced (McTavish et al., 2013). Similarly to the conclusion made in the previous paragraph, the reduction of wake rate of expansion leads to an increased power capture of the upstream turbine. Full scaled field experiments conducted by Schepers and Van Der Pijl (2007) and Boorsma (2012) measured no noticeable power gains when implementing the derating strategy within the ECN wind turbine test station Wieringermeer.

2.5.2 Yaw-based Wake-steering

The yaw-based wake-steering control strategy, as introduced in Section 1.1 and depicted in Figure 1.3, consists in altering the direction of the resulting thrust force of the rotor by misaligning the nacelle direction with the incoming wind speed, such that the wake is deflected away from the downstream turbine(s). As a results, the downstream operates in more favorable conditions and can ideally make up for the power losses of the upstream turbine. In the following paragraphs, results from studies that made use of different types of experiments are investigated and compared.

2.5.2.1 Low- to Mid-fidelity Models

Quick et al. (2017) optimized the power production of the Princess Amalia Wind Farm by means of yaw-based wake-steering with the FLORIS model, taking into account the yaw angle uncertainty, and predicted power gains ranging from +0% to +1.4% depending on the wind direction. Annoni et al. (2018) optimized the same wind farm by means of yaw-based wake-steering and measured power gains reaching +17.8%. Similarly to the observation made for down-regulation, the predictions of the low-fidelity models have a high variability between studies.

2.5.2.2 Mid- to High-fidelity Models

P. A. Fleming et al. (2015) used high-fidelity model SOWFA to model two aligned wind turbines under a constant wind speed. It was observed that yawing the upstream wind turbine counterclockwise and clockwise by 25 degrees resulted in power gains of +4.6% and +4.0% respectfully. According to (Gebraad et al., 2016), the asymmetry originates from the interaction between the wake rotation and the vertical shear. A wind turbine with no yaw angle rotating clockwise will create a counterclockwise rotating wake. Looking downstream, the lower velocities from the atmospheric boundary layer will be transported upward to the right and the higher velocities from the the atmospheric boundary layer will be transported downwards to the left. This will lead to a wake deflection to the right (Gebraad et al., 2016). As a result, yawing counter-clockwise is more effective than yawing clockwise, as the wake deflection due to the yaw-control input matches the natural deflection of the wake, as concluded in P. Fleming et al. (2018). In Bay et al. (2019), the SOFWA model was used to evaluate the power gains of three aligned turbines under constant wind speed and under different turbulent intensities. I

was observed that the difference in turbulence intensities had a major impact on the power production predictions. For low turbulence intensity (5.6%) and high turbulence intensity (10.0%) conditions, power gains of +19.4% and +7.9% were predicted respectively. These results indicate the high sensitivity of the results with respect to the turbulence intensity. Gebraad et al. (2016) used the FLORIS model to optimize the yaw control strategy in terms of power and simulated the power production and loads using SOWFA and the FAST (Fatigue, Aerodynamics, Structures, and Turbulence) aeroelastic code. Although the optimization was performed such to optimize the power production, a general reduction in fatigue loads was observed for the upstream turbine. These results match the findings in Kragh and Hansen (2014), where it was observed that yaw misalignment can reduce the loads of the wind turbines in both deterministic and turbulent flows. Furthermore, Gebraad et al. (2016) observed that yaw-based wake-steering can cause both reduction and increase in fatigue loads, depending on the imbalance caused by the wake overlap. In P. A. Fleming et al. (2015), it is stated that these increased fatigue loads could potentially be mitigated by means of individual pitch control.

2.5.2.3 Wind Tunnel & Field Experiments

A wind tunnel study consisting of five aligned wind turbines with a rotor diameter of $0.15m$ was conducted in Bastankhah and Porté-Agel (2019). A power gain of ranging from 14% to 17% was achieved by introducing a yaw angle of descending magnitude for the first four wind turbines (i.e $\psi_1 > \psi_2 > \psi_3 > \psi_4$, $\psi_5 = 0°$). In a field experiment consisting of five aligned wind turbines in Alberta, Canada, Howland et al. (2019) measured gains up ranging from -13% to $+47\%$ for varying wind directions and wind speeds. The model used to tune the controller was constructed using five years of historical data, creating a spike of interest for the use of historical data to calibrate models (Kheirabadi & Nagamune, 2019).

2.6 Summary & Thesis Motivation

Due to the finite amount of sites available to construct wind farms, it is essential to optimize the production of the wind plant in terms of power production per unit area. In order to maximize the power density, the adverse aerodynamic interactions between the wind turbines should be minimized to reduce power losses and fatigue loads. Both the derating and yawed wake-steering control techniques have shown to be potential candidates to gain more control over the flow within wind farms and limit the adverse wake effects.

The majority of studies investigating the potential of static down-regulation control and static yaw-based wake-steering have been conducted using low-fidelity models under a constant wind speed and direction. Additionally, the results of these studies are spread over too large of a range to make satisfactory conclusions on the potential of these control strategies. The addition of time-varying wind speed, direction and turbulence has shown to have a unfavorable effect on the performance of the controllers (Kheirabadi

& Nagamune, 2019). Furthermore, besides the studies conducted in Frederik et al. (2020) and Munters and Meyers (2018a, 2018b, 2018c), little work as been conducted so far on the potential of the implementation of dynamic axial induction control and dynamic yaw-based wake-steering within the wind farm controllers.

Both down-regulation and yaw-based wake-steering have been reviewed in the literature review. The majority of existing literature is based upon low-fidelity models. Considering down-regulation, the predictions of low-fidelity models are spread over a larger range and are more optimistic than the predictions of the high-fidelity models and physical experiments. Contrarily, for yaw-based wake-redirection, models and experiments both appear to predict power gains spread over a similar range. Considering the loads, down-regulation causes a reduction in loads for the upstream down-regulated wind turbines, but a increase of loads for the downstream wind turbine. Yaw-based wake-steering has the potential to reduce the fatigue loads of the upstream wind turbine, but may cause an increase in loads for the downstream wind turbine. Furthermore, the direction in which the wind turbine is yawed appears to have a considerable impact on the performance in terms of power and loads. On a general note, the performance of the control strategies both showed a high sensitivity to the level of turbulence intensity.

The literature review of down-regulation and yaw-based wake-steering is performed in separate sections, as little work has been done on their combinations at the time of writing to the author's knowledge. As both strategies have their inherent advantages and disadvantages, this study makes an attempt to get the best of both worlds by combining the down-regulation and yaw-control strategies.

Data-driven surrogate models are used for this study, as they have shown to have potential to bridge the gap between low and high-fidelity models by offering reliable, quick and robust predictions (Dimitrov et al., 2018; Hulsman et al., 2019; Schröder et al., 2018). High-fidelity models are reliable and accurate, however they can not be practically used for optimization purposes due to their computational cost. Low-fidelity models have a low computational cost, however their predictions are associated with a high variability (Boersma et al., 2017). Due to the rapid nature of the surrogate models, they could be incorporated into the fast-paced wind farm controllers in order to optimize the operation of wind plants.

The vast majority of studies presented within the past sections are of deterministic nature. However, the interactions between the flow and the wind turbines are non-deterministic by nature. A further goal of this study is to capture this non-deterministic nature by taking an uncertainty-based approach. Instead of modelling the mean wind farm output for a given set of input parameters, this study aims to model the expected output and its associated uncertainty.

CHAPTER 3
Methodology

The methodology is divided into three distinct sections. The first section, Section 3.1, gives an overview of the simulation set-up by reviewing the required background for LES, describing the EllipSys3D flow solver and aeroelastic tool Flex5, and giving an overview of the different scenarios that are simulated. The following section, Section 3.2, reviews how the data of the selected scenarios is generated and processed such that it can be used for the surrogate models. An overview of the structure of the different surrogate models is given at the end of the section. Finally, in Section 3.3, polynomial chaos expansion theory is review and it application to the surrogate models is explained.

3.1 Overview of Simulation Set-up

This section gives an overview of the set-up used to generate the data for the surrogate model calibration. As discussed in the background and state-of-the-art in Chapter 2, the data can originate from low-, mid-, and high-fidelity models, or wind tunnel and field experiments. This study relies on high-fidelity LES, generated with the flowsolver EllipSys3D. Even though high-fidelity simulations have a high accuracy with respect to other lower cost solvers, no model can be assumed to be a perfect representation of reality. However, in order to be able to conduct this study, the assumption has to be made that the results obtained from the LES are correct.

This chapters reviews the required theoretical background on LES in Section 3.1.1. Afterwards, the description and the set-up of EllipSys3D is given in Section 3.1.2 and Section 3.1.3 respectfully. The set-up of the aeroelastic tool Flex5 is described in Section 3.1.4. Lastly, the simulated scenarios are discussed in Section 3.1.5.

3.1.1 EllipSys3D: LES Theory

The computational power required to fully solve the turbulent Navies-Stokes equations for real-world applications using DNS currently exceeds the available computational power (Churchfield et al., 2015). In order to reduce the computational requirements, a combined method consisting of mixing direct numerical simulations and turbulence modelling is commonly used. In LES, large scale eddies are solved directly using the filtered Navies-Stokes equations, while eddies smaller than a certain given grid size Δx are solved using a sub grid-scale model (Lesieur, 1996). In order to separate the sub grid-scale structures from the actual flow, the actual flow is convoluted with a predefined filter function, resulting decomposition of the actual flow u in a filtered flow component $\tilde{u}$ and a subgrid-scale flow component u' (Göçmen et al., 2016),

$$u = \tilde{u} + u'. \tag{3.1}$$

The continuity equation and filtered incompressible Navier-Stokes equations are shown in Equations (3.2) and (3.3), where u is the velocity, t is the temporal dimension, x is the position vector in Cartesian coordinates, ρ is the density, p is the pressure, ν is the kinematic viscosity, S_{ij} is the strain rate tensor, T_{ij} is the subgrid scale tensor, $f_i = \frac{F_i}{\rho}$ represents the external density-normalized body forces, i, j are the directional components and the tilde $\sim$ indicates the filtered variable.

$$\frac{\partial u_j}{\partial x_j} = 0 \tag{3.2}$$

$$\frac{\partial \tilde{u}_i}{\partial t} + \frac{\partial}{\partial x_j}\left(\tilde{u}_i \tilde{u}_j\right) = -\frac{1}{\rho}\frac{\partial \tilde{p}}{\partial x_i} + \frac{\partial}{\partial x_j}\left(2\nu\widetilde{S_{ij}} + T_{ij}\right) + \tilde{f}_i \tag{3.3}$$

The strain rate tensor S_{ij} defines the deformation rate of change within the fluid and is defined by Equation (3.4). The subgrid scale tensor T_{ij}, shown in Equation (3.5), determines the coupling between the small and large scale turbulence (Göçmen et al., 2016).

$$\widetilde{S_{ij}} = \frac{1}{2}\left(\frac{\partial \tilde{u}_i}{x_j} + \frac{\partial \tilde{u}_j}{x_i}\right) \tag{3.4}$$

$$T_{ij} = \tilde{u}_i \tilde{u}_j - \widetilde{u_i u_j} \tag{3.5}$$

Under the Boussinesq hypothesis, the subgrid scale tensor T_{ij} can be related to the filtered strain rate tensor $\widetilde{S_{ij}}$ by Equation (3.6), where ν_{SGS} represents the subgrid-scale eddy viscosity (Boussinesq, 1877; Schmitt, 2007).

$$T_{ij} = 2\nu_{SGS}\widetilde{S_{ij}} \tag{3.6}$$

An more extensive review of a number of subgrid-scale eddy viscosity models, including the commonly used Smagorinsky model (Smagorinsky, 1963), is presented in Lesieur (1996).

3.1.2 EllipSys3D Description

EllipSys3D is 3-dimensional CFD flow solver featuring both RANS and LES that discretizes the incompressible filtered Navier-Stokes equations shown in Equations (3.2) and (3.3) using multiblock finite volumes in general curvilinear coordinates (J. A. Michelsen, 1992, 1994; N. Sørensen, 1995a). For the purposes of this study, the LES approach is chosen with a hybrid scheme consisting of the third-order QUICK scheme (Leonard, 1979) and fourth-order central differencing schemes (Göçmen et al., 2016). By means of hybrid scheme, the undesired numerical wiggles associated with the fourth order numerical scheme are reduced and the numerical diffusion due to the upwind nature of the QUICK

scheme is limited (Troldborg, 2009). The external forces f_i, shown in Equation (3.3), are decomposed into three components,

$$\tilde{f}_i = f_{i,turb} + f_{i,ABL} + f_{i,WT}, \tag{3.7}$$

where $f_{i,turb}$ represents the external body forces introducing the atmospheric turbulence, $f_{i,ABL}$ the external body forces to model the prescribed Atmospheric Boundary Layer (ABL) and f_{WT} the external body forces representing the wind turbine aerodynamic loads (Hulsman et al., 2019). The modelling of the external body forces are discussed in the following subsections.

3.1.2.1 Atmospheric Turbulence $f_{i,turb}$

Atmospheric turbulence is created by means of Mann turbulence box (Mann, 1994, 1998), which is used to introduce a synthetic three-dimensional force field within an upstream plane of the simulation domain, as done in (Troldborg et al., 2014). The generated turbulence is based upon *Taylor's frozen wake hypothesis*, which allows to relate the temporal and spatial turbulent fluctuations of the flow (Taylor, 1938). The hypothesis holds under the assumption that if the turbulent components are small enough relative to the mean flow speed, the turbulent structures remain unchanged as they are advected downstream by the mean flow. As a result, the turbulence experienced by a fixed point over the simulation time will be, to a certain extent, similar to the instantaneous turbulence over the axis of the mean flow direction.

The three-dimensional turbulent velocity field is generated by means of the spectral velocity tensor, which reflects the spectrum of frequencies contained within the atmospheric turbulence. The resulting turbulent velocity field is incompressible, homogeneous, stationary, Gaussian and anisotropic (Sanderse et al., 2011). *Rapid Distortion Theory*, which implies linearization of the Navier-Stokes equations, is used to model the stretching of the turbulent structure due to shear (Mann, 1994). Additionally, the lifetime of the eddies is related to the assumed rate of viscous dissipation of specific turbulent kinetic energy of the system ϵ and the turbulent length scale L_{turb}, which are both input parameters defined by the user (Mann, 1994). The input parameters required for the model are: $\alpha_{turb} \cdot \epsilon^{2/3}$, L_{turb} and Γ, where α_{turb} is the Kolmogorov constant and Γ is a coefficient reflecting the anisotropic character of the turbulence field.

3.1.2.2 Atmospheric Boundary Layer $f_{i,ABL}$

Similarly to the works of Hulsman et al. (2019) and Troldborg et al. (2014), the atmospheric boundary layer is modelled independently from the atmospheric turbulence by imposing body forces over the whole computational domain. The body forces are determined by means of the discretized momentum equation at each time step over the computational domain for a desired mean wind shear profile. An alternative method could consists in letting the flow develop a shear layer over time by tuning the bottom boundary condition and turbulence fluctuations.

3.1.2.3 Wind turbine Aerodynamic Loads $f_{i,WT}$

The Actuator Line Method (ALM), introduced in J. Sørensen and Shen (2002), allows to model the loading along actuator lines rotating within a three-dimensional flow field. Each actuator line, representing an individual rotor blade, is subjected to a radially varying aerodynamic load derived from the velocity, angle of attack and airfoil type at the corresponding span wise location. At each span wise location, the local velocity and angle of attack is used as an input to determine the local lift and drag forces from the tabulated data of the corresponding airfoil, which are then projected into the CFD mesh of the flow field.

Due to numerical discontinuities between the actuator line nodes and CFD mesh, the forces are distributed along the actuator line nodes using a Gaussian distribution (Ivanell et al., 2007). The advantage of the the ALM over the Actuator Disk Method (ADM) is that since wind turbine blades are modelled individually, opposed to ADM where the resulting force field is derived from the circumferential integration of the aerodynamics of the blades. As a result, ALM allows to capture the tip and root vortices, as shown in Ivanell et al. (2007).

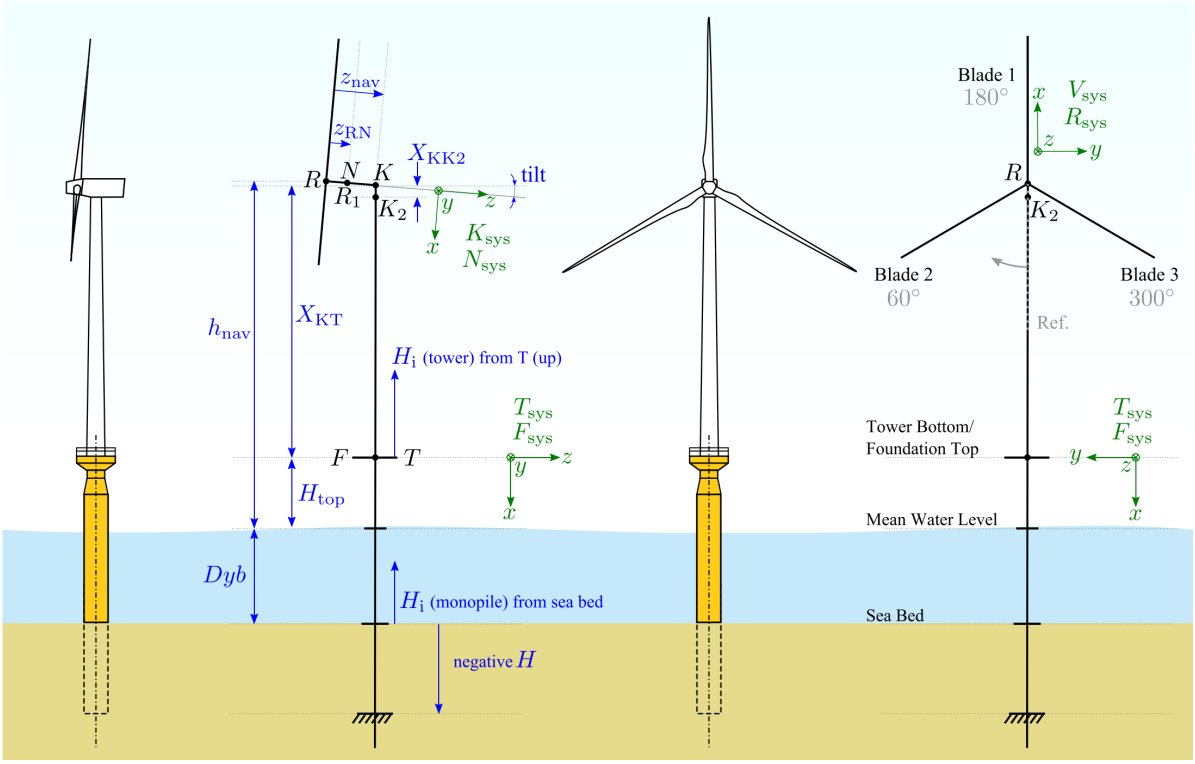

Figure 3.1: Flex5 geometry and coordinate systems. Reprinted from (Sessarego et al., 2017).

The actuator lines can be connected with the structural properties of the blades in order to capture the structural dynamics using the aeroelastic code Flex5 developed by Oye

(1996). The structural dynamics are modelled in Flex5 by means of the virtual work principle and up to 28 degrees of freedom (Sessarego et al., 2017). The Flex5 geometry and coordinate systems is shown in Figure 3.1. The actuator line model is coupled to Flex5 by transferring the velocity field at the actuator line from EllipSys3D to Flex5 at each time step, with which Flex5 determines and returns the corresponding deflections and loads of the entire wind turbine to EllipSys3D. The coupling allows to capture the full aeroelastic response of the wind turbine, including the blade deflections modelled by moving the actuator lines accordingly. For more information about the actuator line model composed of the EllipSys3D and Flex5 coupling, the reader is referred to Sorensen et al. (2015).

3.1.3 EllipSys3D Set-up

The following sections give an overview of the EllipSys3D set-up used to generate the data for the surrogate models. Firstly, a description of the numerical domain is given in Section 3.1.3.1. In the next section, Section 3.1.3.2, the generation of atmospheric turbulence by means of the Mann turbulence box is reviewed. The following section, Section 3.1.3.3, gives a description of the imposed boundary layer. Finally, the last section Section 3.1.5, gives an overview of the simulated cases in terms of yaw-control and down-regulation.

3.1.3.1 Numerical Domain

The LES simulation run over a period of 1619s (= 26min59s) and the numerical domain spans over a domain of $58.6420\frac{x}{R} \times 36.9248\frac{y}{R} \times 13.6861\frac{z}{R}$ represented by $920 \times 192 \times 192$ grid points. The wind turbine tower dynamics are included within the wind turbine model, however it aerodynamic effects are not simulated. The boundary conditions are set to *no slip* for the bottom, *farfield* for the top, and *cyclic* for the edges. The xy-plane and yz-plane of the numerical domain is shown in Figures 3.2 and 3.3 respectfully. The evolution of the cell sizes over the domain are shown in the three directions in Figure 3.4.

3.1.3.2 Atmospheric Turbulence

The input parameters used for the Mann turbulence box, described in Section 3.1.2, are shown in Table 3.1. A resulting turbulence intensity of 7.7% is calculated across the simulation domain. In Hulsman (2018), it was observed that the turbulence length L_{Turb} increases with height due to the higher wind speeds at higher altitudes, whilst the terms related to the energy dissipation ($\alpha_{turb}\epsilon^{2/3}$ and Γ) decrease at larger altitudes. The value of these parameters are thus depending on the height at which the calibration data was measured. For the sake of this analysis, it is assumed that the current set of input parameters Table 3.1 are valid to model the turbulence over the complete simulation domain.

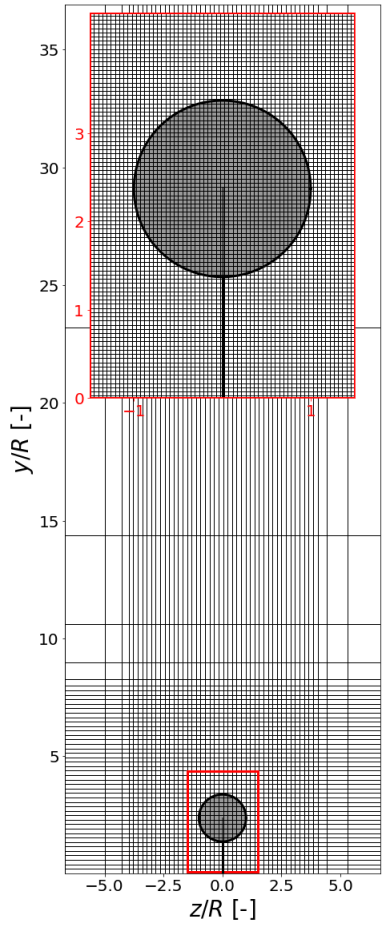

Figure 3.2: Vertical xy-plane of the mesh used in for the numerical domain in EllipSys3D. Figure axes have a 1:1 scale ratio w.r.t on another. Mesh shown within red frame is shown in full resolution.

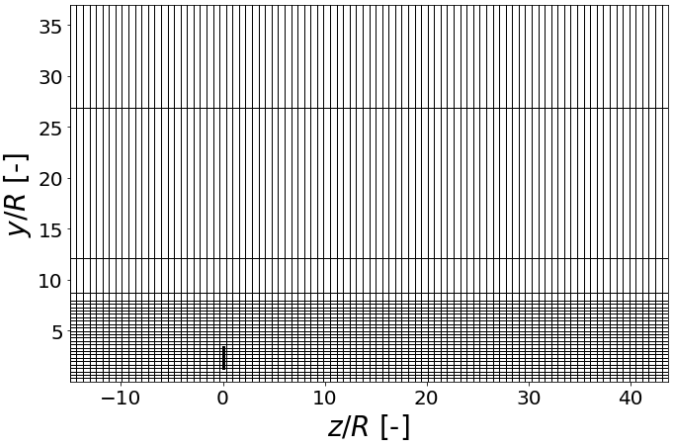

Figure 3.3: Horizontal yz-plane of the mesh used in for the numerical domain in EllipSys3D. Axes have a 1:1 scale ratio w.r.t on another. The mesh is not shown in full resolution.

Table 3.1: Input parameters for Mann turbulence box.

Mann box input	Value	Unit
$\alpha_{turb}\epsilon^{2/3}$	0.03379	$[\frac{m^{4/3}}{s^2}]$
L_{Turb}	156.69	$[m]$
Γ	3.0516	$[-]$

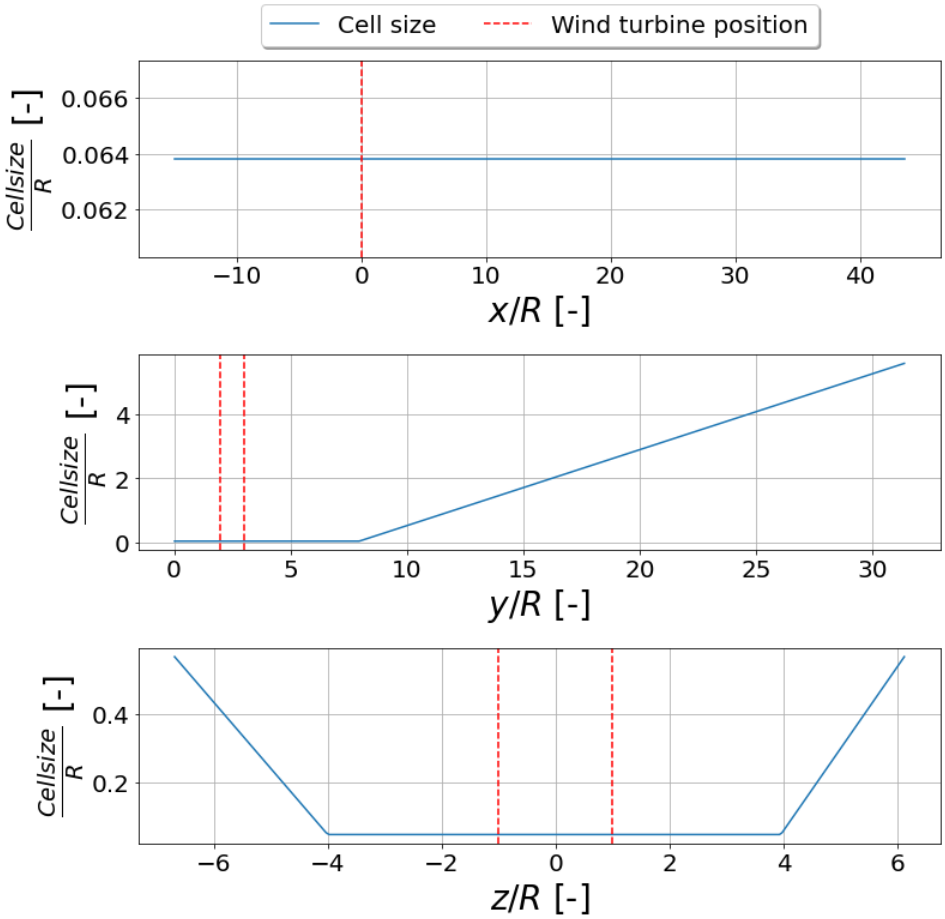

Figure 3.4: Radially normalized sizes of the cell in x-,y- and z-direction. The dotted red lines show the location of the wind turbine within the domain.

3.1.3.3 Atmospheric Boundary Layer

The atmospheric boundary layer is modelled independently from the atmospheric turbulence by means of the introduction of body forces over the whole computational domain. The settings of the body forces are based upon the study performed in Hulsman (2018). The inflow velocity is set to $U_\infty = 8.3283 m/s$ and it is assumed that the flow has no veer. The resulting prescribed boundary layer and its standard deviation is shown in Figure 3.5.

3.1.4 Flex5 Set-up

As described in Section 3.1.2, the actuator line model in the EllipSys3D LES is coupled with Flex5. The dimensions and characteristics of the Flex5 model are based upon the Vestas V27 and the model parameters are to some extent based upon the V27 described in Resor and LeBlanc (2014). The relevant geometric and operational properties, retrieved from Petersen (1990) & Resor and LeBlanc (2014) respectfully, are summarised in

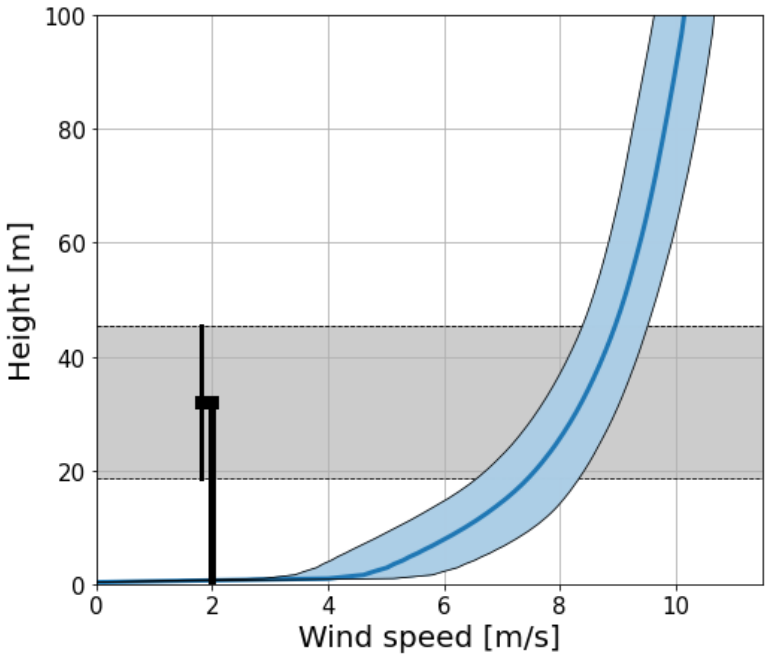

Figure 3.5: Prescribed atmospheric boundary layer imposed in the LES. The mean and standard deviation are represented by the blue solid line and shaded area respectfully.

Table 3.2. The free stream wind speed is set to 8.3283 m/s within the LES. As the V27 has a rated speed of 11.74 m/s, the wind turbine is said to operate within the below-rated regime in these simulations. The wind turbine is equipped with a PI-controller for the blade pitch angle.

Table 3.2: Main geometry and performance parameters of the V27 (Petersen, 1990; Resor & LeBlanc, 2014).

Parameter	**Symbol**	**Value**	**Unit**
Rotor radius	R	13.5	[m]
Hub height	H_{hub}	31.5	[m]
Tilt angle	θ_{tilt}	4	[deg]
Coning angle	θ_{coning}	0	[deg]
Blade length (profiled)	/	13.0 (11.5)	[m]
Blade profiles	/	NACA 62-200 series	/
Tower height	T_{height}	31	[m]
Rated power	P_{rated}	225	[kW]
Rated wind speed	U_{rated}	11.74	[m/s]
Rated rotational speed	ω_{rated}	44	[rpm]
Design tip-speed-ratio	λ_{design}	7.61	[-]

The modal frequencies of the model of the turbine contained within Flex5 are based upon the experimental values from Resor and LeBlanc (2014). A comparison of the main experimental modal frequencies obtained from the experiment and the Flex5 modal are shown in Table 3.3. A sweep analysis of the effects of the wind speed on the power, operational settings and loads of the V27 is conducted and discussed in Appendix B.

Table 3.3: Comparison of the experimental modal frequencies obtained in Resor and LeBlanc (2014) and the modal frequencies from the Flex5 model.

Modal Frequencies		
Name	**Experiment [Hz] (Resor & LeBlanc, 2014)**	**Flex5 [Hz]**
1st & 2nd Tower FA	1.00 & 7.83	0.98 & 8.06
1st & 2nd Tower SS	1.01 & 7.97	0.98 & 8.06
1st Shaft torsion	2.65*	2.65
1st & 2nd flapwise	2.40 & 6.67	2.36 & 7.11
1st edgewise	3.65	3.64

* Value computed with the polymax algorithm.

3.1.5 Simulation Cases

Two wind farm flow control strategies are analysed within this study, namely axial-induction control (under the form of constant pitch control down-regulation) and yaw-based wake-steering (or yaw-control). For both strategies, the thrust force plays an essential role in the development of the wake. Firstly, the thrust force determines the wake deficit within the wake (Ainslie, 1988). Secondly, for the yaw-based wake-steering strategy, the cross-wind thrust force component causes the wake deflection (Bastankhah & Porté-Agel, 2016). Due to the essential role of the thrust force, the LES scenario were selected based on cases of equivalent thrust.

The cases were identified using the following procedure. Firstly, an LES with no wind turbine is run to obtain the undisturbed flow across the domain. A cross section of the undisturbed dynamic velocity field is extracted from the LES at the location of where the wind turbine will be located. The undisturbed dynamic velocity-field cross section is then used as flow input for the Flex5 aeroelastic code. Note that this approach differs from the coupling described in Section 3.1.2.3, where the actuator line model used within the LES is *directly* coupled with Flex5. As there is no wind turbine present within the used LES, Flex5 is *indirectly* coupled with the LES by extracting a cross section dynamic velocity field. This approach, referred to as the *ghost* wind turbine approach, allows to retrieve wind turbine power and loads for the desired wind turbine settings without having to run a different LES each time. The thrust of the wind turbine over the simulation time for a selected number of yaw angles and the time-averaged thrust force is stored for each case. Afterwards, a range of down-regulated cases are run in

Flex5 by introducing different constant pitch control angles. The five blade pitch angles resulting in a time-averaged thrust force which is equivalent to their counterclockwise (and clockwise) yaw counterparts are saved and used for the LES. The procedure used to determine the combined cases is described in Appendix A. An overview of the LES scenarios inputs is given in Table 3.4[1].

The resulting thrust equivalent scenarios obtained through Flex5 are shown in Figure 3.6, where the normalized time-averaged thrust force of the cases is plotted against the yaw input setting used for the simulation. Each down-regulated case (*green circle*) lays less than 1% away of the same equivalent thrust line (*horizontal line*) as the its corresponding yaw-controlled case(s) (*red/brown squares*). The difference in effectiveness between yawing CW and CCW can already be noted, as for a similar yaw angle it appears that the thrust of the CCW yawed turbine is reduced slightly more than the thrust of the CW yawed turbine. This discrepancy between the CW and CWW yaw-control has been observed in different studies discusses in Section 2.5.2.2 and is further explained in (Gebraad et al., 2016).

Using the control inputs determined by the method described above, five down-regulated cases, eight yaw-controlled cases and four combined cases are defined and run in Ellip-Sys3D. Another LES is run without a down-regulation or yaw control input to serve as a

[1]Due to a partial loss of data as a result of a glitch on the DTU university cluster, the CW yaw cases have not been included in the calibration data set of the surrogate models.

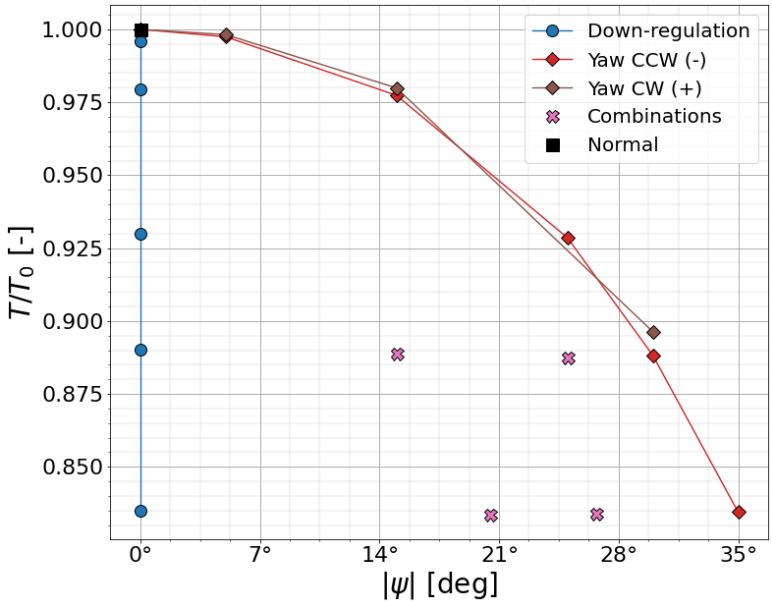

Figure 3.6: Time-averaged thrust for the down-regulated and yawed cases in function of the input yaw angle as determined by Flex5 with the undisturbed dynamic inflow from the turbine-less LES. The thrust is normalized with respect to the case with no down-regulation or yaw input, referred to as *normal* operation

.

Table 3.4: Overview of blade pitch angle and yaw angle inputs used for LES simulations.

Case Name	**Blade pitch angle input:** θ_1 **[deg]**	**Yaw input:** ψ_1 **[deg]**
No turbine	/	/
Normal	0	0
Down-regulation (1)	0.1	0
Down-regulation (2)	0.3	0
Down-regulation (3)	1.0	0
Down-regulation (4)	1.6	0
Down-regulation (5)	2.4	0
Yaw-control CCW (1)	0	-5
Yaw-control CCW (2)	0	-15
Yaw-control CCW (3)	0	-25
Yaw-control CCW (4)	0	-30
Yaw-control CCW (5)	0	-35
Yaw-control CW (1)*	0	5
Yaw-control CW (2)*	0	15
Yaw-control CW (3)*	0	30
Combination (1)	1.3	-15
Combination (2)	1.6	-20.5
Combination (3)	0.6	-25
Combination (4)	1.1	-26.7

*Not included in surrogate model calibration data.

reference. Finally, including the initial LES without the wind turbine, a total of nineteen simulations have been ran. A predefined turbulence seed is used for every simulation in order to guarantee similar initial conditions for each run.

The results of the LES containing a wind turbine model are compared in Figure 3.7 in terms of time-averaged power and thrust, normalized by the reference *normal* case. Note that the data is taken directly from the output of the LES/Actuator line coupling set-up. It can be observed that decreasing the thrust causes a larger decrease in power for yaw-control than for down-regulation. Furthermore, yawing counterclockwise appears to have a slight advantage in power compared to yawing clockwise for a similar thrust.

3.2 Data Generation, Processing and Fitting.

This section gives an overview of the process used to produce the calibration data for the surrogate models. Four main steps can be identified within this process, as shown in Figure 3.8. The first step consists in actually generating the data with the set-up described in the previous section. The methods used to generate the time series data for the upstream and downstream wind turbines are discussed in Section 3.2.1. Once

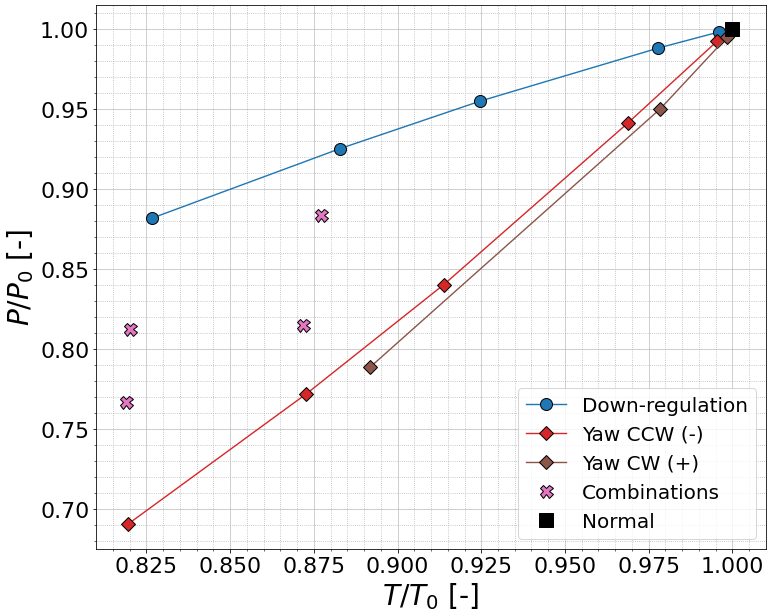

Figure 3.7: Comparison of the selected LES scenarios in terms of normalized thrust and power. The used data is taken straight from the LES/Actuator line model outputs. The down-regulation cases are shown by the green circles, the CCW yaw cases by the red squares and the CW yaw cases by the brown squares.

the time series are obtained, the second steps starts and the 10-minute statistics are determined and then normalized, as described in Section 3.2.2. The data has to be processed such that a statistical distribution can be fitted from the values obtained from the time series. The third step, which involves the fitting of the distributions and other plausible methods, are discussed in Section 3.2.3. Finally, during the fourth and final step, the surrogate models are constructed. A description of the set-up of the different surrogate models is given in Section 3.2.4.

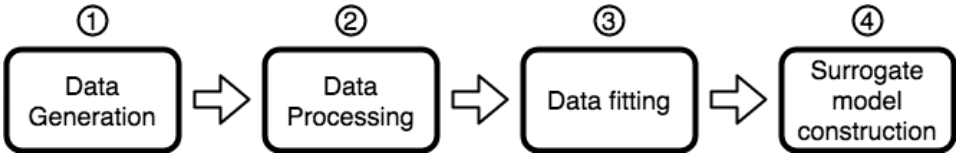

Figure 3.8: Diagram showing the required steps to prepare the data for the calibration of the surrogate models.

3.2.1 Data Generation

For the upstream turbine, the channel outputs can be retrieved from the direct EllipSyS3D/actuator line coupling for each LES. However, no downstream wind turbine is present within these LES scenarios. Therefore, an alternative approach has to be taken in order to obtain the channel outputs for the downstream turbine for different spacings

and inputs. This approach is discussed in Section 3.2.1.1. As a result, the implications of this alternative approach have to be considered for the upstream turbine, as done in Section 3.2.1.2. Finally, an overview of the simulations cases is given in Section 3.2.1.3.

3.2.1.1 Downstream Turbine

For each upstream wind turbine case shown in Table 3.4, an LES is run with the listed control inputs for the upstream turbine. There is no downstream wind turbine present during the LES. The power and loads of the downstream turbine are determined by means of the process shown in Figure 3.9. Given a LES scenario, a "*ghost*" turbine is placed at downstream distance sx of the upstream wind turbine. A cross section perpendicular to the axial direction of the velocity field at the downstream distance sx is selected. The dynamic inflow at cross section is extracted and used as input for the Flex5 aeroelastic code with the wind turbine model described in Section 3.1.4. This method is similar to the one described in Section 3.1.5, where the LES with no wind turbine has been used instead. The procedure is repeated for every LES simulation, every downstream distance of the second wind turbine and every combination of control inputs of the down-stream turbine listed in Table 3.5 in the second column.

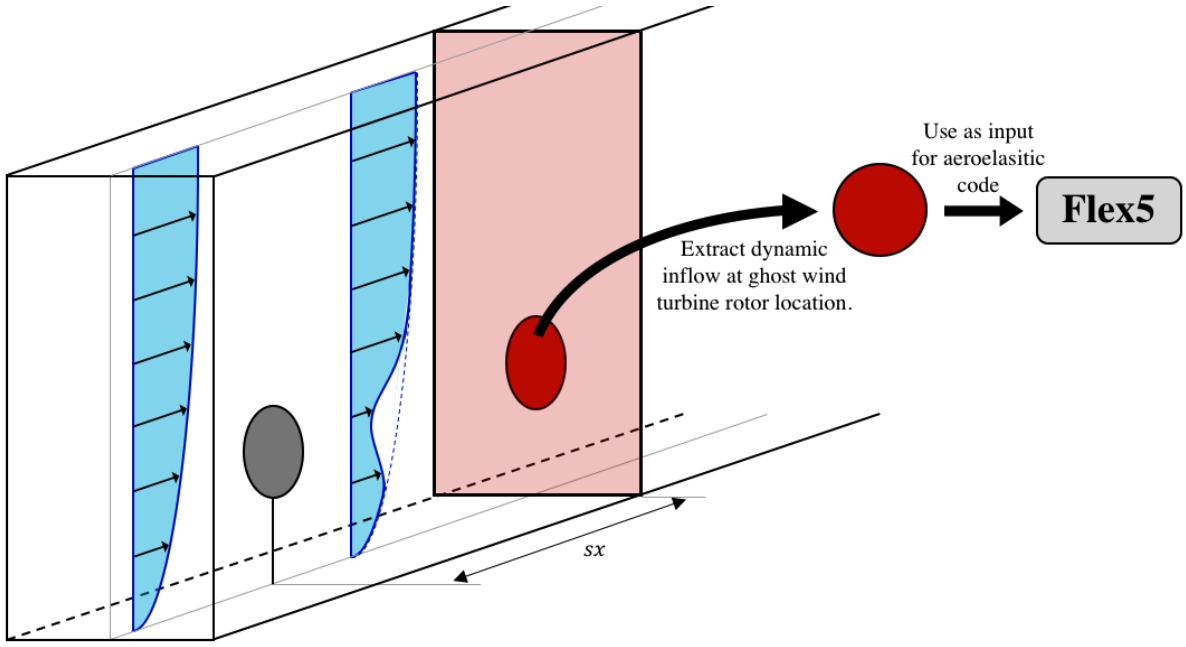

Figure 3.9: Sketch illustrating the ghost wind turbine approach. The flow field at the location of where the second wind turbine would have been located (red cross section) is extracted from the LES and is used as a dynamic inflow for the aeroelastic code Flex5.

3.2.1.2 Upstream Turbine

The power and loads of the upstream wind turbine can be computed in two different ways. The first approach is similar to the *ghost* wind turbine approach used for the downstream

turbine and the indirect LES/Flex5 coupling described in Section 3.1.5. Using the LES without a wind turbine, a cross-section perpendicular to the axial direction is extracted at the location of the upstream turbine. The extracted dynamic inflow is then used as input for the wind turbine model in Flex5 with the desired control inputs. The second approach consists in obtaining the data directly from the LES/actuator line model coupling.

The advantages of the ghost wind turbine approach over the direct EllipSys3D/Flex5 coupling are that the control inputs for the upstream wind turbine are not limited to the specific cases selected for the LES simulations. Furthermore, as the downstream wind turbine is using a similar approach, the data of the upstream and downstream wind turbine can be directly compared. The downsides of the ghost wind turbine approach are the larger temporal scales and the simplified physics due to the indirect coupling of the EllipSys3D and Flex5. However, the use of the ghost wind turbine approach introduces some uncertainty within the system. As an effort to quantify this uncertainty, the two approaches are compared quantitatively in Appendix C.

3.2.1.3 Data Generation Cases

An overview of the cases used for the data generation for the upstream turbine and downstream wind turbine are shown in Table 3.5. For the *upstream wind turbine*, data is generated using the 'ghost' wind turbine approach for combinations of yaw cases $\psi_1 = [-35° : 2.5° : 30°]$ and down-regulated cases $\theta_1 = [0.0° : 0.1° : 2.4°]$, as shown in the left column. A total of $27 \times 25 = 675$ cases are run for the upstream wind turbine using the ghost wind turbine approach.

Table 3.5: Overview of the inputs used for the Flex5 simulations using the ghost turbine approach for the upstream wind turbine (left) and downstream wind turbine (right).

	Upstream wind turbine cases	**Downstream wind turbine cases**
$\boldsymbol{\psi_1}$ [deg]	[-35:2.5:30]	LES scenarios*
$\boldsymbol{\theta_1}$ [deg]	[0:0.1:2.4]	LES scenarios*
$\boldsymbol{\psi_2}$ [deg]	/	[-30:5:30]
$\boldsymbol{\theta_2}$ [deg]	/	[0:0.3:2.4]
$\boldsymbol{S_x}$ [R]	/	[4:1:20]

*Excluding "*no turbine*" and "*CW Yaw-control*" scenarios listed in Table 3.4

For the *downstream wind turbine*, combinations of yaw cases $\psi_2 = [-30° : 5° : 30°]$ and down-regulated cases $\theta_2 = [0.0° : 0.3° : 2.4°]$ are run for all LES scenarios listed in Table 3.4, for the exception of the '*no turbine*' scenario and the CW yaw LES scenarios. The former is omitted as it is only used for the upstream ghost turbine. The latter has not been omitted by choice, but due to a glitch on the DTU Cluster resulting in the loss of the LES data for the scenarios. As a result, a total of $15 \times 13 \times 9 = 1755$ cases are run in Flex5 for the downstream wind turbine using the ghost wind turbine approach.

This lack of positive yaw data can possibly result in instability of the surrogate models when the upstream yaw angle ψ_1 approaches 0°. The effects of the lack of positive yaw data are further discussed in Section 4.2.1.

3.2.2 Data Processing

Within the next paragraphs, the data processing method is discussed. Firstly, the channels of interest are listed and described in Section 3.2.2.1. Once the channels have been selected, the time series have to be converted into 10-min statistics. This is done in different ways for the different channels, as explained in Section 3.2.2.2. Finally, the data scaling and normalization process is described in Section 3.2.2.3.

3.2.2.1 Channel Selection

Ten load channels are considered to compare the performance of the different control strategies. The selection of loads is similar to the studies conducted in Dimitrov et al. (2018) and Damiani et al. (2018). The selected loads are the main contributors to the fatigue limit state and blade extreme design loads for the load cases described in *DLC1.2* and *DLC1.3* in the *IEC 61400-1* (International Electrotechnical Commission, 2005). The selected load channels are the electrical power P, blade root flapwise bending moment M_{BRF}, blade root edgewise bending moment M_{BRE}, main shaft torque M_{MS}, tower top torsion (or yaw moment) M_{yaw}, tower top bending moment M_{TT}, and tower bottom bending moment M_{TB}, as shown in Table 3.6. Due to the symmetrical nature of the wind turbine tower, the two latter, M_{TT} & M_{TB}, are determined the resultant bending moment from the axial and lateral bending moments. A spectral analysis of the selected loads and their corresponding eigenfrequencies is conducted in Appendix D.

Table 3.6: Overview of Flex5 channels and their corresponding symbol and unit.

Load Nr.	Load channels	Symbol	Unit
1.	Electrical Power	P	[kW]
2.	Blade Root Flapwise (BRF) bending moment	M_{BRF}	[kNm]
3.	Blade Root Edgewise (BRF) bending moment	M_{BRE}	[kNm]
4.	Main Shaft (MS) torque	M_{MS}	[kNm]
5.	Tower top torque/yaw moment	M_{yaw}	[kNm]
6.	Tower Top (TT) bending moment*	M_{TT}	[kNm]
7.	Tower Bottom (TB) bending moment*	M_{TB}	[kNm]

* Combined bending moment: $M = \sqrt{M_x^2 + M_y^2}$

3.2.2.2 Short-term Equivalent Load

The short-term damage-equivalent loads (DEL) are determined by the expression shown in Equation (3.8), where S_i are the load cycle ranges and n_i are the corresponding number of cycles at a given range.

$$S_{eq} = \left[\sum \frac{n_i S_i^m}{T_{ref}} \right]^{1/m} \tag{3.8}$$

Both parameters are determined by the rainflow counting algorithm (Rychlik, 1987). The reference number of cycles T_{ref} is set to the chosen time interval such that $1Hz$ equivalent loads are obtained. The wöhler exponent m is a material dependent property derived from its S-N curve. The used Wöhler exponents m and their corresponding material and load channel are given in Table 3.7. Information about the composition of the V27 can be found in Petersen (1990).

Table 3.7: Overview of the relevant wind turbine materials, their Wöhler exponents and the corresponding load channel numbers.

Material	Wöhler exponent	Load channel Nr.
Welded steel	4	5-7
Cast iron	7	2
Glass fiber	10	3,4

3.2.2.3 Sample Selection, Scaling, and Normalization

The total simulation time of the LES is $1361.48s$. The first $\pm 100s$ are excluded from the analysis in order to allow the flow to settle. The resulting simulation time of $1260s$ is divided into 23 10-minute windows shifted by $30s$ each. Although the windows are not statistically independent, they do convey part of the inherent variability of the flow.

The full process for each $1260s$ time series corresponding to a specific set of inputs is shown in the diagram in Figure 3.10. Once 23 10-minute samples have been created, a distinction is made between the power and the load channels. For the power, the mean of each 10-minute window is computed. As the other loads are subject to high-frequency oscillations, the short-term equivalent load of each 10-minute window is determined. As a result, for each case described in Table 3.5, 23 samples are obtained.

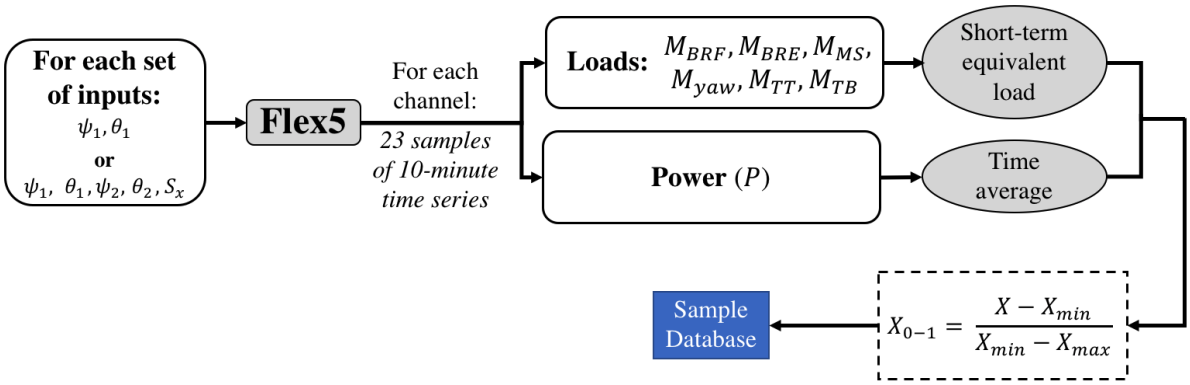

Figure 3.10: Diagram shown an overview of the steps involved in the processing of the data. For a selected set of inputs, 23 10-minute time series are obtained for each channel. Depending on the channel, the short term equivalent load or time average is determined. The data is then normalized and stored.

The following step in the diagram is the normalization of the data. All simulation inputs and channel outputs have to be normalized within the 0 to 1 value range. This is crucial for the construction of the surrogate models, as it is essential that all variables carry the same weight. The normalization is done according to Equation (3.9), where X_{min} and X_{max} represent that minimal and maximal input and output values obtained from the simulations.

$$X_{0-1} = \frac{X - X_{min}}{X_{max} - X_{max}} \tag{3.9}$$

3.2.3 Data Fitting

Once the data properly converted into 10-minute statistics and normalized accordingly, a discussion can be conducted regarding how the statistics of the 10-minute data points for each channel should be represented. Three approaches are considered, namely *percentile fitting*, *Kernel Density Estimate (KDE) fitting* and *empirical CDF fitting*. The advantages and disadvantages of each of these approaches are discussed within the following three subsections.

3.2.3.1 Percentile Fitting

The first approach consists in constructing a data set containing the median value of the 23 samples for each set of inputs listed in Table 3.5. The surrogate models calibrated with this data set would allow to approximate the expected median value (Q_{50}) for a given set of inputs. As it is the purpose to perform optimization under uncertainty, the procedure is repeated for different percentiles as well, e.g. 25^{th} (Q_{25}) and 75^{th} (Q_{75}) percentiles. A sketch of the approach is shown in Figure 3.11 Separate data sets containing the Q_{25} percentile and Q_{75} percentile in function of the set of inputs are constructed. The two

surrogate models calibrated using these two data sets allow to approximate the interval in which 50% of the samples would potentially be.

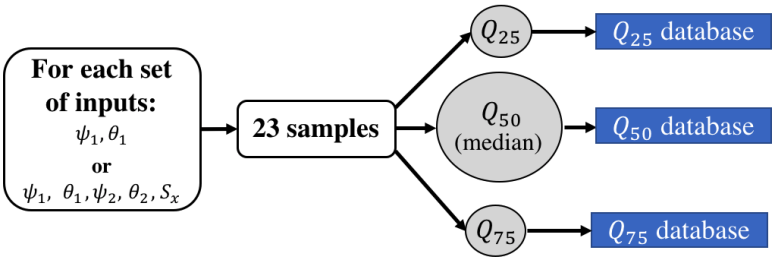

Figure 3.11: Diagram showing the percentile fitting approach. For each set of inputs for each channel, 23 samples are obtained. The 25^{th}, 50^{th} (median), and 75^{th} percentiles are stored.

The main advantage of using percentiles to estimate the median values and expected intervals is that no assumption has to be made regarding the distribution of the 23 samples across the input domain. However, the size of the intervals is limited to the percentiles for which the surrogates are constructed. If a larger interval is desired, e.g. an interval containing 90% of the data, new surrogate calibrated using Q_{05} and Q_{95} have to be constructed. Another issue arises due to the fact that there are only 23 samples for each input set. When determining the percentiles, only the median value will be an actual data point. The other percentiles, i.e. Q_{25} and Q_{75}, would have to be interpolated between two samples. Furthermore, due to the limited size of the data set with which these percentiles are determined, the percentile estimations might be noncontinuous across the input domain.

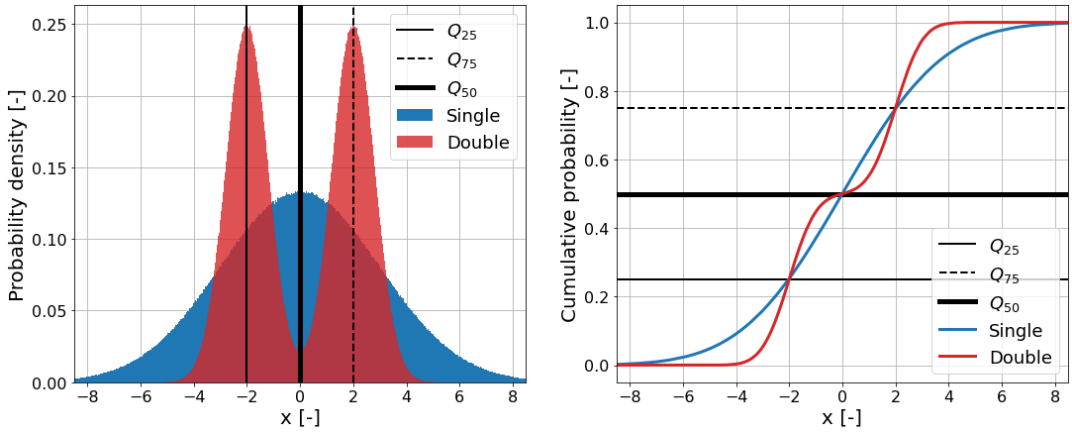

Figure 3.12: Probability distribution (left) and cumulative probability (right) of a single (blue) and double (orange) Gaussian distribution. Both distribution have similar 25^{th}, 50^{th}, and 75^{th} percentiles.

Finally, the main downside of using percentiles is shown in Figure 3.12. The left plot

shows the Cumulative Distribution Function (CDF) of both a *single Gaussian* and a *double Gaussian*. The main percentiles, namely Q_{25}, Q_{50} and Q_{75}, are indicated by the vertical lines. On the right plot, the corresponding CDF are plotted along the with the main percentiles represented by the horizontal lines. Although the distributions are totally different, they have exactly the same Q_{25}, Q_{50} and Q_{75} values. This means that no conclusions can be drawn about the distribution of the parameters when using these percentiles. For the reasons listed above, this approach has been discarded. The following approach, the Kernel Density Estimate (KDE) fitting, makes an attempt at modelling the actual samples distribution without making assumptions of the distribution itself.

3.2.3.2 Kernel Density Estimate Fitting

The second approach makes use of the KDE. KDE assigns a user-defined distribution (or *kernel*) with a user-defined width (or *bandwidth*) to each single data point. The distributions are summed up and then divided by the number of sample elements in order to keep the probability within the 0-to-1 interval. KDE is a very powerful tool to quickly estimate the distribution of a sample set, however there are some drawbacks. Firstly, two user-defined parameters, namely the kernel and bandwidth, are required. For large data sets, the kernel selection does not matter, as the resulting estimated distribution will tend to the same shape as the number of samples increases to infinity. However, in the case of 23 samples the chosen Kernel has a considerable impact. As for the bandwidth, its optimal value can be determined by using k-fold cross-validation in combination with a grid-search. However, the issue is that the optimal bandwidth should be determined for each set of inputs, which could be a source of discontinuity between different sample sets. The second drawback of KDE it does not return parameters or straightforward equations that describe the resulting distribution. This issue can be circumvented using the method shown in Figure 3.13. Given a certain set of inputs, the KDE generated Probability Density Function (PDF) can be computed once a bandwidth and kernel have been chosen. In this example, the output values of the KDE are then fitted using two convoluted Gaussians (or a Double Gaussian), after which the distribution parameters are stored.

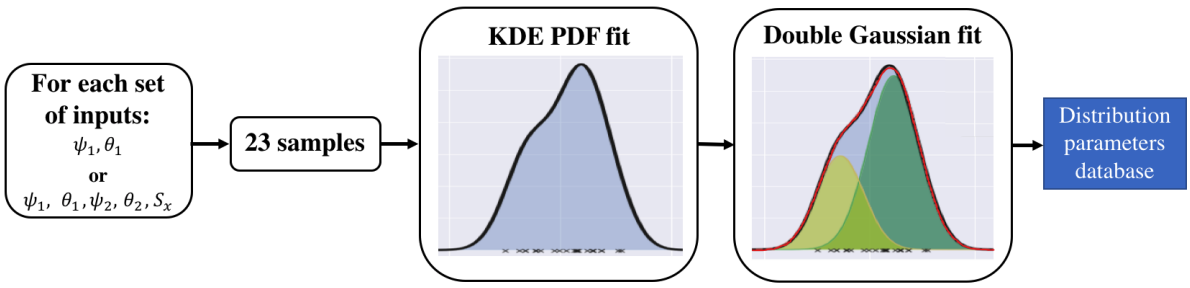

Figure 3.13: Diagram showing the steps involved in the KDE fitting approach. For each set of 23 samples, the KDE is estimated and fitted using double Gaussians. The parameter distributions are then stored.

Ideally, the fitted distribution should be directly related to the data. However, with this approach, an extra step is included with the KDE. This extra step within the data fitting process is not desired, as it adds a layer of complexity to the process and distances the final fit from the actual data. In order to simplify and shorten the fitting process, a selected distribution could be fitted to the empirical Cumulative Distribution Function (CDF) of each sample s set.

3.2.3.3 Emperical CDF Fitting

The final approach consists in fitting the empirical CDF of the samples with a chosen (set of) distributions. This allows to obtain distribution parameters that are directly related to the sample distribution, rather than to an *estimation* of the sample distribution, as described in the previous method. However, this comes at the cost of having to make an assumption regarding the distribution of the parameters.

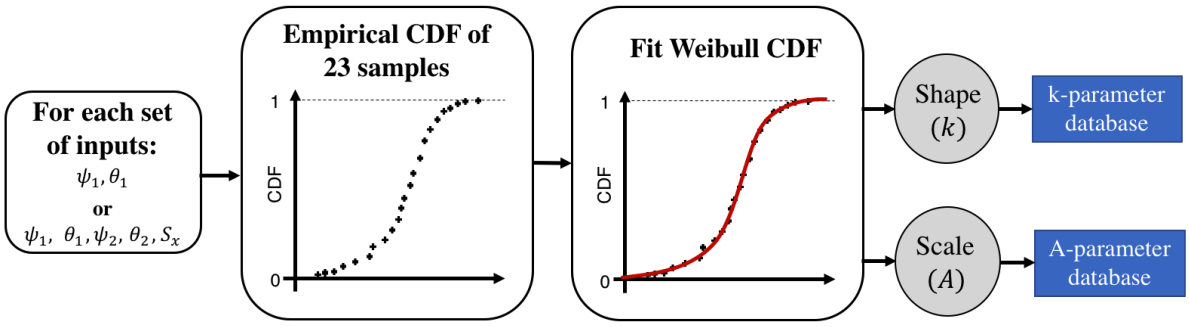

Figure 3.14: Diagram showing the steps involved in empirical CDF fitting. For each set of 23 samples, the empirical fit is determined, after which a single Weibull distribution is fitted to it. The distribution parameters parameters are then stored.

The approach is shown schematically in Figure 3.14. For a given set of inputs, 23 samples are obtained. The empirical CDF of these samples can readily be determined through the two following steps:

1. For a given sample size of n, the data is sorted in ascending order. Each sample is given a rank number k from 1 to n associated with its position in the list of sorted samples.
2. The empirical CDF for samples with rank $k = 1, ..., n$ is determined using Equation (3.10).

$$F_{empirical}(k) = \frac{k}{n+1} \tag{3.10}$$

The next step consists in constructing a fit for the empirical CDF. A decision has to be made regarding the assumed distribution of the samples over the input domain. Weibull distributions are selected, as they offer more flexibility than Gaussian distributions. The equation for the single and double Weibull CDFs are given in Equations (3.11) and (3.12)

respectively, where k_x are the shape parameters, A_x the scale parameters and w_0 the relative weighting of the fist and second Weibull. For each set of inputs, the k_x and A_x parameters have to be fitted to the empirical CDF and stored.

$$F_{Single,Weibull} = 1 - e^{-\left(\frac{x}{A}\right)^k} \tag{3.11}$$

$$F_{Double,Weibull} = w_0 \cdot \left[1 - e^{-\left(\frac{x}{A_1}\right)^{k_1}}\right] + (1 - w_0) \cdot \left[1 - e^{-\left(\frac{x}{A_2}\right)^{k_2}}\right] \tag{3.12}$$

However, in some cases, using a single Weibull is not enough to capture the intricacies of the samples distributions. Figure 3.15 shows the empirical CDF for the blade root flapwise moment $M_{BRF,2}$ given a spacing of $S_x = 10R$ and $\psi_1 = \psi_2 = \theta_2 = 0°$. A down-regulation input of $\theta_1 = 0°$ and $\theta_1 = 1.1°$ are used as input the left and right plot respectively. The *single Weibull* and *double Weibull* fits are shown in blue and orange respectively. The Mean Absolute Error (MAE) as function of the cumulative probability is given in the legend for each fit. It can be seen that the single Weibull has a larger MAE than the double Weibull, certainly when the CDF is more complex as seen in the right plot for $\theta_1 = 1.1°$.

The fits are performed using 5-fold cross-validation with the Maximum Likelihood Estimation (MLE). MLE helps determine the distribution parameters (e.g $\boldsymbol{\alpha}$) for which the samples ($\boldsymbol{x}$) are most likely to belong to. The optimal parameters set $\boldsymbol{\alpha}$ of a chosen distribution distribution for a given sample set $\boldsymbol{x}$ can be determined through Equation (3.13).

$$\boldsymbol{\alpha}_{MLE} = \arg\min \left[-\log \mathcal{L}\left(\boldsymbol{\alpha}|\boldsymbol{x}\right)\right] \tag{3.13}$$

Through visual inspection, it can be seen that for the case with no blade pitch angle input ($\theta_1 = 0°$), both the *single* and *double Weibull* represent the sample distribution quite well. Quantitatively, The double Weibull outperforms the single Weibull with nearly halved value of $MAE = 0.02$. When the upstream blade pitch angle is increased to $\theta_1 = 1.1°$, as shown in the right plot, the benefits of using the *double Weibull* can be observed. It appears that the samples are no longer continuous and are concentrated within different intervals. For this case, the advantage of the flexibility of the double Weibull is clearly observed within the 0.50 to 0.54 interval on the x-axis. However, even the double Weibull does not manage to capture the distribution of the first few samples accurately, as can be seen in the 0.43 to 0.46 interval.

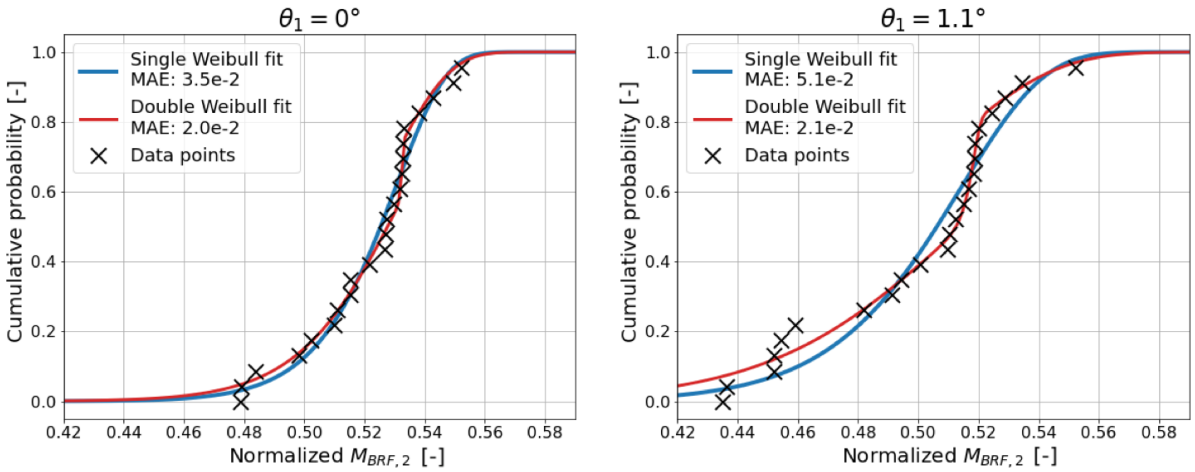

Figure 3.15: Comparison of the quality of the fit of the single (blue) and double (orange) Weibull distributions using the empirical CDF of the downstream wind turbine blade root flapwise bending moment with $S_x = 10R$, $\psi_1 = \psi_2 = \theta_2 = 0°$ and $\theta_1 = 0°$ (left) or $\theta_1 = 2.4°$ (right). The legend shows the corresponding mean absolute error of the fit.

Using a double Weibull over a single Weibull has the downside that a total of five parameters have to be fitted, namely the shape parameters (k_1,k_2), the scale parameters (A_1,A_2) and the relative weighting (w_0) of the single Weibulls of which the double Weibull is composed, as shown in Equation (3.12). A stability analysis of the single and double Weibull parameters is conducted in Equation Section 4.1.1.1. Within this analysis, it is concluded that the distribution parameters of the double Weibull are too unstable across the parameter domain to be used for the polynomial based surrogate models. Moreover, the analysis also goes over the fact that the double Weibull does not always result in a better fit. As a result, the surrogates are constructed using the single Weibull distribution due to their smooth distribution parameters across the input domain.

3.2.4 Surrogate Model Structure and Calibration

For the purposes of this study, two main surrogate models are constructed, namely one for the upstream wind turbine and one for the wind farm. Section 3.2.4.1 discusses the inputs used for the surrogate models and how they impact the meaning of the outputs of the surrogate models. Afterwards, the structure of the upstream turbine and wind farm surrogates are presented in Section 3.2.4.2 and Section 3.2.4.3 respectively.

3.2.4.1 Surrogate Model Inputs

The constructed surrogate models take the yaw and blade pitch angles used for the LES simulations as input. However, it is essential to consider that these simulation inputs, as listed in Table 3.5, are not the *effective* yaw angle or blade pitch angle that the wind turbine experiences. As the flow constantly varies due to gusts and turbulence, the instantaneous yaw and blade pitch angles will constantly vary. However, capturing

the instantaneous effective angles that the wind turbine makes with the flow is a very complex process.

For instance, the Rotor Equivalent Wind Speed (REWS) has to be determined in order to determine the effective yaw angle. The REWS can be computed by taking the spatial average of the wind direction across the rotor. However, performing this calculation for each time step is computationally expensive. Therefore, the REWS can be related to the hub wind speed by using a low-pass filter. This low-pass filter would filter out high frequency random local wind direction fluctuations that do not affect the rotor as a whole. The ways with which the low-pass filter can be constructed are determined in Appendix E. However, determining the effective blade pitch angles still poses challenge, as it dynamically varies over the rotor blade length, rotor rotational angle and instantaneous turbulence.

The constructed surrogate models are data-driven. The data-driven approach has the advantage that the individual steps of complex processes do not have to be quantified exactly, as the models only require data connecting a given input with a desired final output. Rather than determining each step within the complex processes that occur within the system, data-driven model make use of statistical approaches to determine the underlying relation. The surrogate models can be used to determine the relation between the initial control input and the final desired channel, without having to determine the effective flow angles.

By assuming that the simulation input angles correspond to the wind farm controller inputs, most of the heavy work is relayed to the data-driven model. The advantage is that no complex system of equations have to be solved to obtain the effective yaw and blade pitch angle, as this will automatically be taken into account by the data-driven models. However, it introduces a small nuance in the meaning of the model outputs. By optimizing the wind farm control inputs using surrogate models calibrated with the simulation input angles, the returned optimized inputs might not be the actual flow angles in which the wind turbine is operating. Rather, the returned quantities are the inputs that have to be given to these specific simulations in order to optimize the desired quantity. The consequences of this assumption are discussed more elaborately in Section 4.2.1.

3.2.4.2 Upstream Wind Turbine Surrogate Model

The first main surrogate model is calibrated only by the inputs and outputs of the upstream wind turbine. The upstream wind turbine surrogate model is referred to as *WT1* surrogate, and is shown schematically in Figure 3.16. Its two inputs, shown on the left side, are the yaw (ψ_1) and blade pitch (θ_1) inputs of the wind turbine. The main surrogate is divided into two sub-surrogates, used to generated the shape (k) and scale (A) parameters of the Weibull distribution. These surrogates are calibrated using the 675 cases listed in the left column of Table 3.5. Using the surrogate-generated distribution parameters, the CDF of the distribution can be constructed. The WT1 surrogate is constructed for each channel listed in Table 3.6 for the upstream turbine and are indicated

to belong to the upstream wind turbine output by the "1" subscript (e.g. P_1)

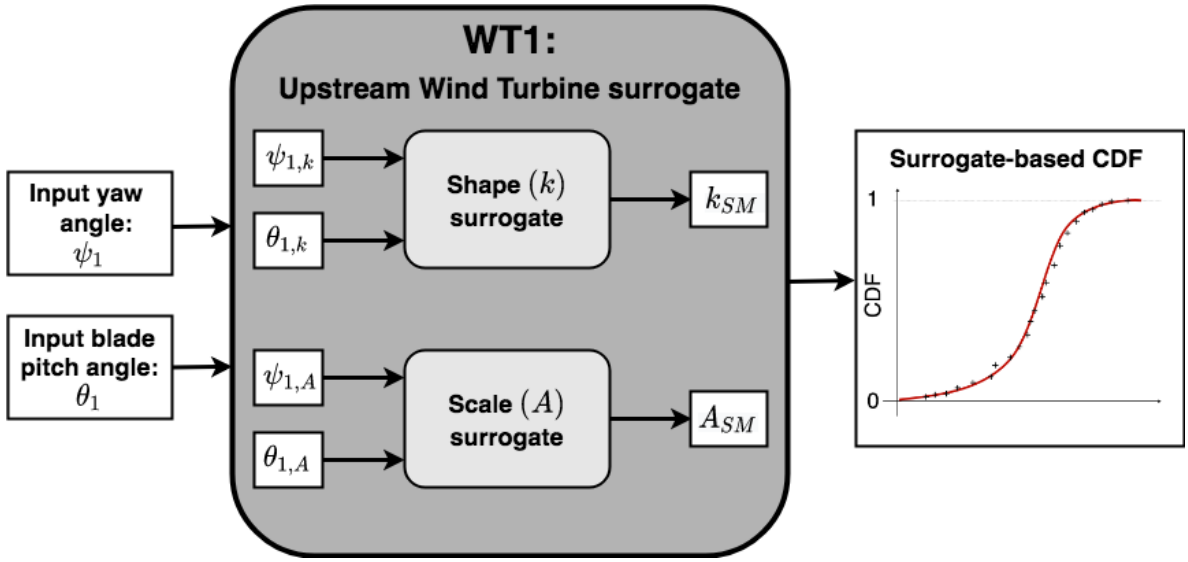

Figure 3.16: Structure of upstream wind turbine surrogate model.

3.2.4.3 Wind Farm Surrogate Model

The second main surrogate model encompasses both the upstream and downstream wind turbine inputs, as shown schematically in Figure 3.17. A separate surrogate is constructed for each wind turbine spacing, from $S_x = 4R$ to $S_x = 20R$ in steps of $1R$, as the spacing S_x is not considered an optimization parameter for this study. The drawback of constructing a separate surrogate model for each spacing rather than a larger surrogate model that includes the spacing as a parameter is that there is loss of continuity between the separate surrogate models in terms of the spacing input. However, the surrogates do no require spacing S_x as an input, which reduces their individual complexity. Furthermore, this study aims to analyse the effectiveness of the different control strategies for a range of given set inter-turbine distances. If the purpose is to reduce aerodynamic interactions between the wind turbines, a larger spacing is always beneficial. However, in reality, larger turbine spacing lead to higher costs. Thus, if it is desired to include spacing as an optimization parameter, a full financial model would have to be included.

The main wind farm surrogate, or *WF* surrogate, is constructed analogously to the WT1 surrogate model. It contains both a shape (k) and scale (A) surrogate model. However, there are four main differences. Firstly, a different WF surrogate model is constructed for each spacing S_x. Secondly, the shape and scale surrogates of the WF surrogate take two extra inputs, namely the downstream wind turbine yaw angle ψ_2 and blade pitch angle θ_2. Thirdly, these surrogates are calibrated using the 1755 cases listed in the right column of Table 3.5. Finally, the WF surrogate outputs the channels of the downstream wind turbine listed in Table 3.6, which is indicated by the subscript "2" (e.g. P_2). The WF

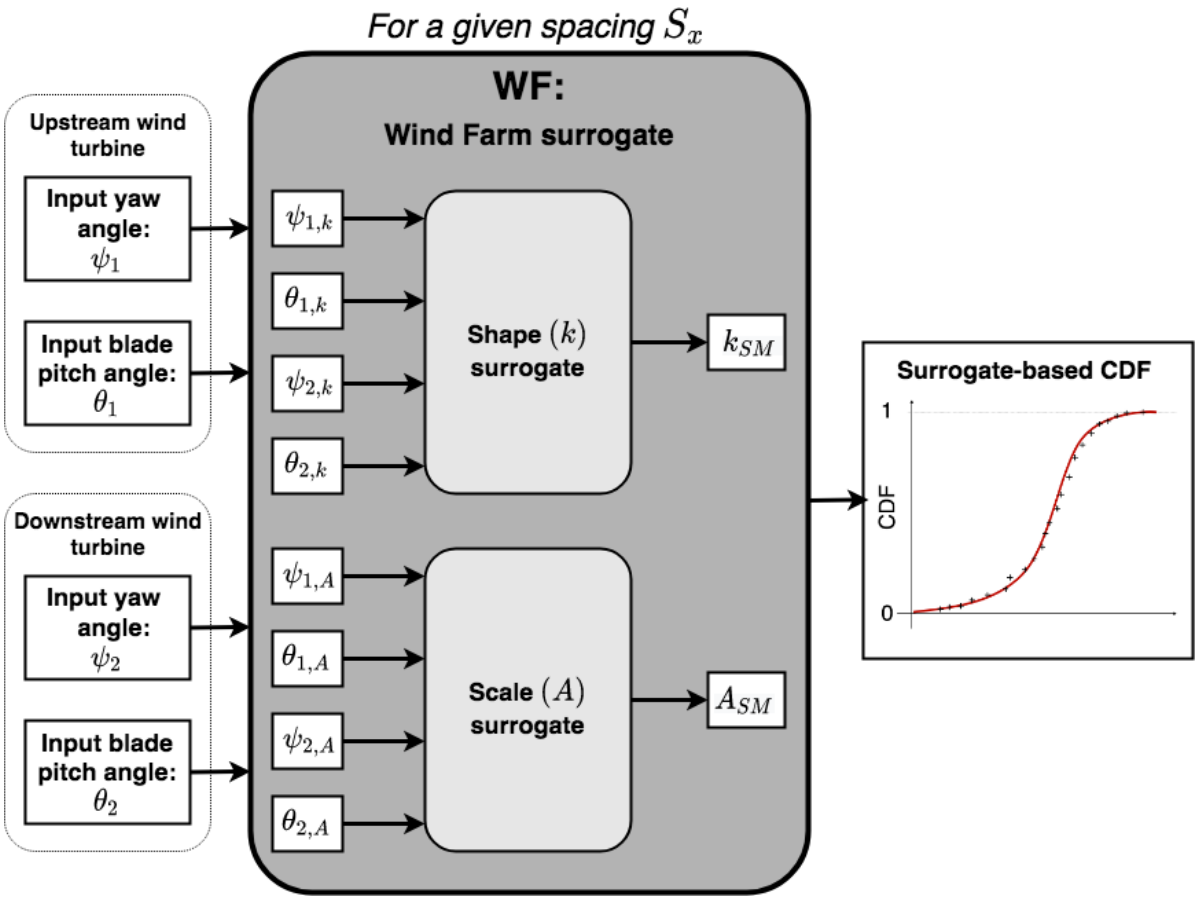

Figure 3.17: Structure of wind farm surrogate model.

surrogate features an additional channel, which is the wind farm total power, denoted by P_{tot}. This channel is obtained by taking the sum of the upstream and downstream power channels, i.e. $P_{tot} = P_1 + P_2$.

3.3 Polynomial Chaos Expansion

The Polynomial Chaos Expansion (PCE) method allows to construct an approximation of the stochastic response surface of a model with random input parameters by means of a orthogonal polynomial Hilbert basis. The PCE method, developed by Wiener (1938) and re-introduced into the field of engineering by Ghanem and Spanos (1991), serves as a popular alternative to Monte Carlos simulations for uncertainty quantification due to its simplicity and fast convergence rate (Murcia et al., 2018). As the number of variables to be computed increases drastically with the number of random input variable, sparse PCE surrogate models are created by means of adaptive training algorithms. These algorithms avoid over-fitting and detect the most relevant coefficient, minimizing the terms with a multivariate dependency and increasing the efficiency of the surrogate model (Blatman & Sudret, 2011).

In a study conducted by Dimitrov et al. (2018), the PCE model outperformed other surrogate models based Quadratic response surfaces (QRS), universal Kriging, importance sampling, and nearest-neighbor interpolation when considering the combination of computational time and accuracy. Murcia et al. (2018) PCE surrogate models to estimate the site-specific power production and lifetime equivalent fatigue loads of the 10MW DTU reference wind turbine without including the aeroelastic design. This study is a continuation of the work done in Hulsman et al. (2019), where PCE models were used to model the effect of yaw on a two aligned in terms of loading and power.

3.3.1 Theoretical Background

Adopting the notation from Sudret (2008), the stochastic models S defined as (non-linear) functions f with a finite number of random input variables $\mathbf{X}$ as input,

$$S = f(\mathbf{X}) = f(X_1, ..., X_M). \tag{3.14}$$

Assuming that the response surface S has a finite variance, it can be represented by the PCE method by Equation (3.15) (Soize & Ghanem, 2005),

$$S = \sum_{j=0}^{\infty} S_j \phi_j \left(X_1, ..., X_M\right), \tag{3.15}$$

where S_j are the unknown coefficient of ϕ_j and $\phi_j\{X_1, ..., X_M\}_{j=0}^{\infty}$ is basis within the corresponding Hilbert space of the response. This basis corresponds to a polynomial family, such as the Hermite or Legendre polynomial families which is orthogonal to a given distribution, as further described in Appendix F. The orthogonality of the polynomial families is defined by the inner product of the polynomials within the corresponding Hilbert space,

$$\langle \phi_m, \phi_n \rangle := \int_a^b \phi_m(x)\phi_n(x)w(x)dx = h_n \delta_{mn} \tag{3.16}$$

where ϕ_n is an element from the polynomial family with degree n, $w(x)$ is the corresponding probability density function or *weight function*, h_n is a value depending on the degree of the polynomial and δ_{mn} is the Kronecker delta, defined by,

$$\delta_{mn} := \begin{cases} 0 & \text{for m} \neq \text{n} \\ 1 & \text{for m} = \text{n} \end{cases}. \tag{3.17}$$

The weight function $w(x)$ has to be positive and continuous over the (in)finite interval [a,b], such that the moments of $w(x)$ exist. In order to limit the required computational power to construct the response surface, the maximal degree of the M-dimensional polynomials is limited to a degree p. After the truncation, Equation (3.15) can be rewritten as,

$$S \approx \sum_{j=0}^{P-1} S_j \phi(\mathbf{X}) = \sum_{j=0}^{P-1} S_j \phi(X_1, ..., X_M), \tag{3.18}$$

where P corresponds to the number of coefficients,

$$P = \binom{M+p}{p} = \frac{(M+p)!}{M!\,p!}. \tag{3.19}$$

From Equation (3.19), it can be observed that PCE suffers from the "curse of dimensionality". The adaptive training algorithms described in Blatman and Sudret (2011) thus play an essential part in increasing the sparsity of the coefficient matrices.

The surrogate constructed in this study are all based upon Hermite polynomials. However, the chosen polynomial family is not relevant for the these surrogate models, as it is not their purpose to directly reconstruct a given distribution. As explained in Section 3.2.4, the surrogate models are used to compute a response surface of of the individual distribution *parameters*, an not the distribution itself. The usefulness of the chosen polynomial family is limited to its ability to quickly generate a polynomial basis. This basis is then used to shape a response surface by means of a weighted sum of the polynomials. The method with which the coefficients are determined is discussed in the next section.

3.3.2 Determination of Coefficients

Two main approaches exist to compute the coefficient of a PCE model, namely *intrusive* and *non-intrusive* approaches. Intrusive approaches, e.g *Galerkin Projection* method, require a modification to the system of equations such that the PCE coefficients are the unknowns in the new reformatted system (Eldred, 2009). By means of the Galerkin Projection method, all coefficients are solved in a single run. However, this comes at the cost of an increased computational effort. The non-intrusive methods, e.g. *Spectral Projection* or *Regression* methods, determine each coefficient independently and do no require reformatting of the system of equations, as the existing code can be used without alterations (Kaintura et al., 2018). The non-intrusive regression method is used for the purposes of this study and is described in the following paragraphs.

The regression method consists in conducting a regression of the exact solution S with respect to the chosen PCE basis. The relation between the exact model output S and the model output based upon the PCE $\tilde{S}$ is given in Equation (3.20),

$$S = f(\mathbf{X}) = \tilde{S}(\boldsymbol{\xi}) + \epsilon, \tag{3.20}$$

where $\tilde{S}(\boldsymbol{\xi})$ is defined by,

$$\tilde{S}(\boldsymbol{\xi}) = \sum_{j=0}^{P=0} S_j \phi_j(\boldsymbol{\xi}). \tag{3.21}$$

where $\boldsymbol{\xi}$ is a vector composed of the standardized normalized variables of the random vector $\mathbf{X}$ and ϵ is the zero-mean variable (Sudret, 2008). The PCE model $\tilde{S}$ fits the exact model S in as much as the variance of the zero-mean variable ϵ is minimized by the choosing adequate set coefficients, denoted by $\boldsymbol{\ell} = \{S_j, j = 0, ..., P-1\}$. As a result, the regression can be formulated as,

$$\boldsymbol{\ell} = \text{Argmin } E\left[\left\{f\left(\mathbf{X}(\boldsymbol{\xi})\right) - \tilde{S}(\boldsymbol{\xi})\right\}^2\right]. \tag{3.22}$$

Equation (3.22) is solved by choosing a set of regression points N within the standard normal space, $\{\xi^1, ..., \xi^N\}$, such that a set of random input variables of the input vector $\boldsymbol{X}$ are selected, $\{x^1, ..., x^N\}$, which form the experimental design (Sudret, 2008).

$$\boldsymbol{\ell} = \text{Argmin } \frac{1}{N}\sum_{i=1}^{N}\left\{f((x)^i - \sum_{j=0}^{P-1} S_j\phi_j\left(\boldsymbol{\xi}^i\right)\right\}^2 \tag{3.23}$$

3.3.3 Application of Polynomial Chaos Expansion to Surrogate Models

The surrogates models are constructed by means of the *Chaospy* Python library introduced in Feinberg and Langtangen (2015). Chaospy is a numerical tool that allows to use polynomial expansions for uncertainty quantification. For the purposes of this study, the library is used a tool to quickly generate and fit polynomials in multi-dimensional space. Its application can be compared to the simple *polyfit* function used in Python with the *NumPy* library or in MATLAB, but with the great advantage the the fit can easily be done in the desired amount of dimensions. The procedure to create a PCE-based surrogate model consists in the following steps.

1. Select the desired polynomial families for each input (i.e each dimension).
2. Set the maximal desired polynomial order for each input.
3. Generate the polynomials using three-term recurrence of the corresponding polynomial families.
4. Fit the polynomial coefficients to the data using the regression method.

In the first step, polynomial families for each input are set. As discussed earlier in Section 3.3.1, the selection of polynomial family is not especially relevant. Furthermore, the inputs of the models, namely ψ_1, θ_1, ψ_2, and θ_2, correspond to the dimensions of the system. Thus the surrogates of the upstream wind turbine shown in Figure 3.16 are two dimensional, whilst the surrogate of the wind farm shown in Figure 3.17 are four dimensional.

In the second step, the maximal order polynomial for each input has to be set. Taking the shape-surrogate depicted within Figure 3.16, individual orders have to be determine for both the ψ_1 and θ_1 inputs. This is done in Section 4.1.2.

The actual polynomials are generated using the selected orders and the three-term recurrence relationship of the chosen surrogate family in the third step. For more information on polynomial families and three-term recurrence, the reader is referred to Appendix F. The procedure is illustrated by means of Equation (3.24) with Hermite

polynomials in a one dimensional case with the order set to two. The first matrix, the *coefficient matrix*, is own to the polynomial family and is generated through the three-term recurrence relationship. The second matrix is a vector consisting of polynomials ranging from zero to the selected maximal order. The polynomials are generated by the product of the coefficient matrix and polynomial vector.

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix} = \begin{bmatrix} 1 \\ x \\ x^2 - 1 \end{bmatrix} \tag{3.24}$$

A two dimensional example is shown as generated through ChaosPy for Hermite polynomials in Equation (3.25), where the selected maximal order for the first and second dimension are set to two and three respectively. As it can be seen in the fifth row of the matrices, not only are the polynomial for each individual dimension generated, but also the corresponding cross-terms (i.e $x_0 x_1$).

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & -3 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ x_1 \\ x_1^2 \\ x_1^3 \\ x_0 \\ x_0 x_1 \\ x_0^2 \end{bmatrix} = \begin{bmatrix} 1 \\ x_1 \\ x_0 \\ x_1^2 - 1 \\ x_0 x_1 \\ x_0^2 - 1 \\ x_1^3 - 3x_1 \end{bmatrix} \tag{3.25}$$

In the fourth and final step, the generated polynomials are fitted to a set of inputs and outputs. This is done by assigning a weight to each generated polynomial and taking the resulting sum, as shown for the one dimensional example in Equation (3.26). Using the regression method, the best weights for each polynomial are computed and the polynomial fit is obtained.

$$\begin{bmatrix} w_0 & w_1 & w_2 \end{bmatrix} \begin{bmatrix} 1 \\ x \\ x^2 - 1 \end{bmatrix} = w_2 x^2 + w_1 x + (w_0 - w_2) \tag{3.26}$$

This whole process is used for each channel for all the shape and scale surrogate models shown in Figures 3.16-3.17 by replacing the x parameters by the corresponding yaw ψ or blade pitch angle θ inputs. An essential parameters of this whole process if the order selection. Too low or high of an order will result in under- or overfitting the data. Therefore, it s crucial to conduct a sensitivity analysis to determine the optimal order. This sensitivity analysis is done for the upstream wind turbine and wind farm surrogates in Section 4.1.2.

CHAPTER 4
Results and Discussions

This chapter is divided into two distinct main sections. Firstly, in Section 4.1, the surrogate models are developed and their performance is evaluated. This is done by calibrating the surrogate models and tuning the model hyper parameters accordingly. Furthermore, the error associated with each step during the construction of the surrogate model is quantified. Once the surrogate model errors have been determined, the surrogate models can be used to optimize the wind farm power production, as discussed in the second part of this chapter in Section 4.2.

4.1 Model Development and Performance Evaluation

Three essential steps are defined within the process of the construction of the surrogate models. The process is shown in a diagram format in Figure 4.1, where the three essential steps and the relevant errors are indicated. The first step consists in determining the empirical CDF of the 23 samples for each set of input combinations. In the following step, the empirical CDF is fitted with a chosen distribution. The Maximum Likelihood Estimate (MLE) in combination with a 5-fold cross-validation is used to obtain the distribution parameters. The error between the empirical CDF and the MLE fitted CDF is defined as the *fitting error*. The fitting error and how it is determined is discussed in Section 4.1.1.

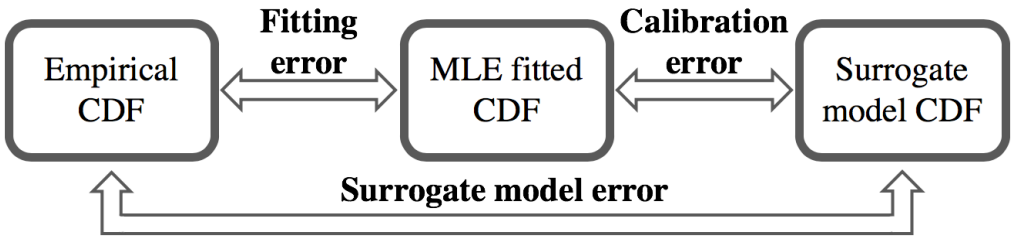

Figure 4.1: Diagram showing the different steps involved in the development of the model the corresponding errors.

The parameters determined by the MLE fit are used as calibration data for the surrogate models in the third step. As discussed in Section 3.2.4, separate surrogate models are constructed for the separate distribution parameters. The error between the MLE fitted CDF and the surrogate generated CDF is referred to as the *calibration error*. It is defined by the error between the calibration data MLE fitted parameters and the surrogate model

parameter outputs. The calibration error is used to determine appropriate orders for the surrogate model inputs in Section 4.1.2.

The error between the surrogate model CDF and empirical CDF is referred to as the *surrogate model error* and is quantified in a similar way as the fitting error. It quantifies the overall performance of surrogate model with respect to the empirical CDF and is discussed in Section 4.1.3.

4.1.1 Fitting Error

The fitting error is the metric used to estimate how well the MLE fitted CDF matches the empirical CDF. The order to illustrate the method with which the error is determined, an example of a not too successful fit where the error is clearly visible is used. For a given set of input parameters, the fit of the upstream turbine main shaft torque $M_{MS,1}$ is represented by Figure 4.2.

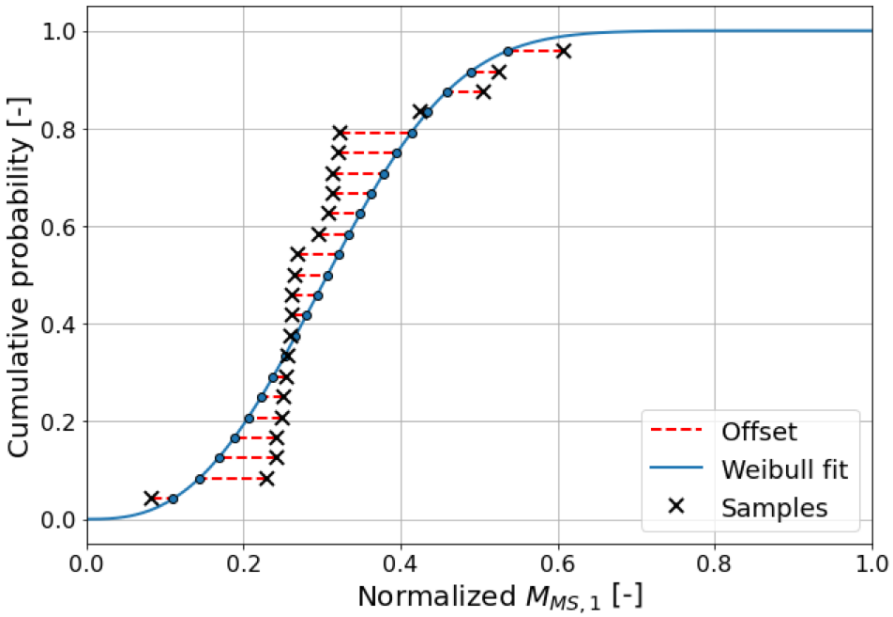

Figure 4.2: Illustration of the method used to quantify the fitting error using the $M_{MS,1}$ channel to illustrate. The fitting error (dotted red line) is defined by the difference between fitted single Weibull (solid blue line) and the samples (black crosses).

The black crosses represent the 23 samples corresponding to the given input and the blue curve is the MLE Weibull fit. For a given cumulative probability, the red dotted line is the difference between the normalized $M_{MS,1}$ according to the sample and the MLE fitted distribution. The error is the normalized format. In order to un-normalize the data, its value has to be scaled back using Equation (4.1), where the nomenclature is similar to the nomenclature used in Section 3.2.2.3.

$$X = X_{0-1}(X_{max} - X_{min}) \tag{4.1}$$

The error (red dotted line in Figure 4.2) can be computed for each sample for each CDF over the whole input domain. Thus, the number of errors computed is the product of the number of samples per CDF and the number of existing input combinations. The mean and standard deviation of the error data set can be computed. The mean error will give the bias of the sample over the input domain and the standard deviation will give an indication of how much it varies. Furthermore, a histogram can be computed with the created error data set to represent the distribution of the error of the CDFs over the whole input domain.

4.1.1.1 Distribution Selection

Due to the smooth nature of the polynomials of which the surrogate models are made off, it is essential that the general tendencies of the distribution parameters remain continuous throughout the input parameter domain. The magnitude of the random variations of the distribution parameters are not allowed to exceed the general tendencies of the distribution parameters within the input domain. In other words, the distribution parameters can be noisy, but large irregular discontinuities can not be captured effectively by the polynomials.

As discussed in Section 3.2.3.3, the double Weibull has the advantage over the single Weibull that it can capture more of the intricacies of the samples distributions for a given set of inputs. However, it comes at the cost of having to fit five distribution parameters, namely the scale an shape coefficients of both Weibull distribution and their relative weighting with respect to one another. The double Weibull equation is showed in Equation (3.12). An analysis of the smoothness of the double Weibull parameters is conducted within the next few paragraphs. For a given spacing of $S_x = 5R$ and upstream turbine running under normal operation with inputs $\psi_1 = 0°$ and $\theta_1 = 0°$, the tower top yaw moment $M_{yaw,2}$ is determined in function of the inputs ψ_2 and θ_2 of the downstream turbine. The corresponding value of the weight w_0 parameter of the double Weibull is shown in Figure 4.3. It can be observed that the majority of weights, mostly on the positive ψ_2 side of the plot, are equal to $w_0 = 0.5$. This means that an equal weight is attributed to both the first and second Weibull of the double Weibull.

The exact value of $w_0 = 0.5$ is no coincidence. This weight occurs when the shape and scale parameters of both Weibulls converge to the same values ($k_1 = k_2$ and $A_1 = A_2$). This is shown in Figure 4.4, where the difference between the scale ($\Delta k = k_2 - k_1$) and shape parameters ($\Delta A = A_2 - A_1$) are shown. It can be observed that the input combination that resulted in a weight of $w_0 = 0.5$ correspond to $\Delta k = \Delta A = 0$. This means that the double Weibull actually converged to a single Weibull distribution. For the other cases, the fit converged to a double Weibull. These sudden jumps between single and double Weibull distributions are source of discontinuity that can only be capture by the polynomial-based surrogate models if their order is set very high. However, this would result in overfitting the data, which is undesired. The purpose of the surrogate is to capture the general tendency of the data, not all its variations. The heat maps of the first and second shape and shape parameters (k_1, k_2, A_1 and A_2) are shown in

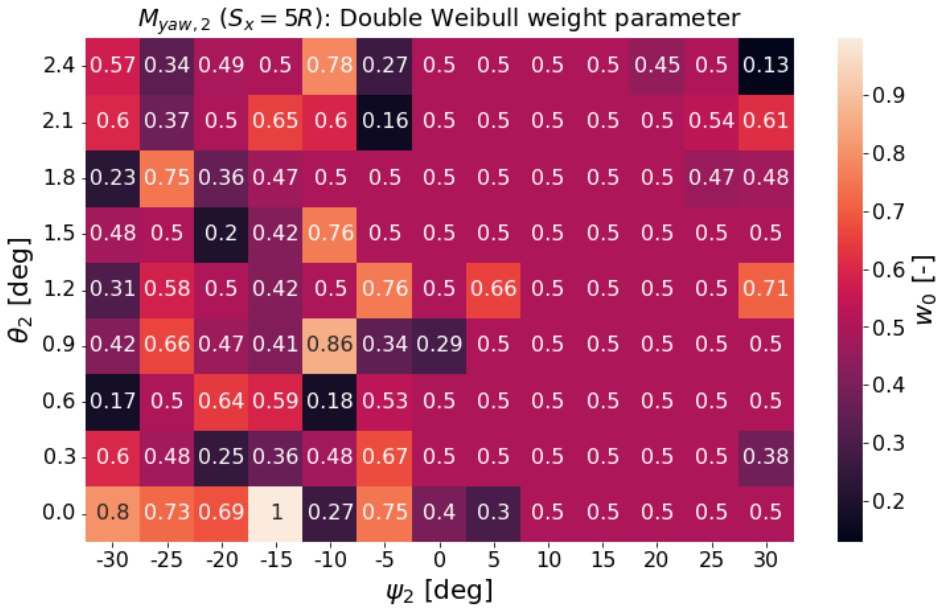

Figure 4.3: Weight parameter (w_0) of the double Weibull distribution for $S_x = 5R$ and $\psi_1 = \theta_1 = 0°$ for the downstream wind turbine tower top yaw moment channel.

Appendix G. In contrast, the distribution parameters of the single Weibull, shown in Figure 4.5, are less noisy and the general tendencies can be more easily discerned.

The difference in accuracy between the single and double Weibull approach are compared using the averaged Mean Absolute Error (MAE) metric over the input domain in Figure 4.6. Note that the values in the heat maps have been scaled by a factor of 10^2. Negative values a better performance of the double Weibull. Examining the left side of

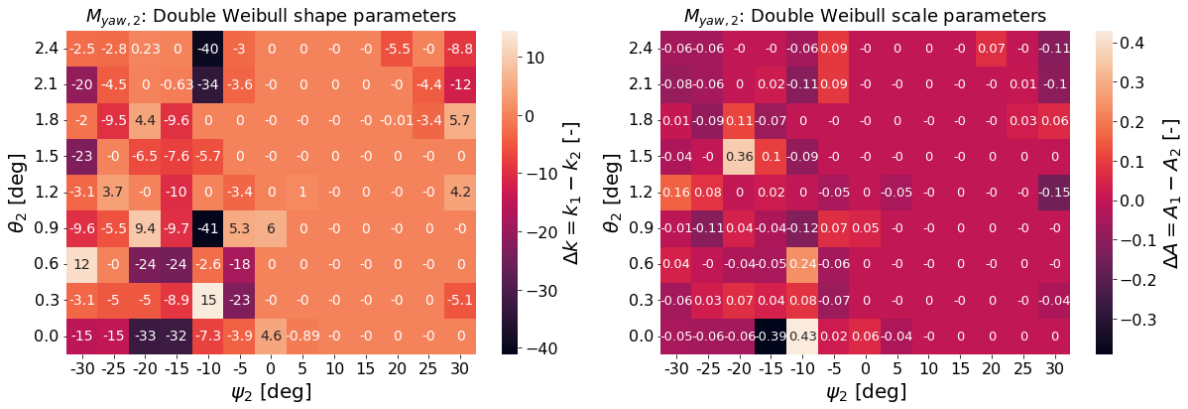

Figure 4.4: Difference between shape parameters (left) and difference between the scale parameters (right) of the two Weibull distributions of which the double Weibull is composed off for $S_x = 5R$ and $\psi_1 = \theta_1 = 0°$ for the downstream wind turbine tower top yaw moment channel.

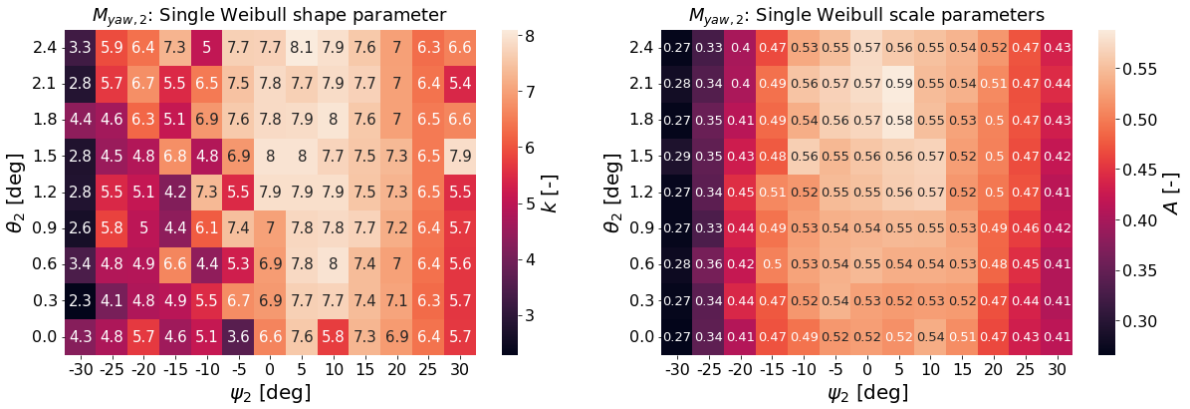

Figure 4.5: Shape parameter (left) and scale parameter (right) of the single Weibull distribution for $S_x = 5R$ and $\psi_1 = \theta_1 = 0°$ for the downstream wind turbine tower top yaw moment channel.

the plot, it appears that the double Weibull does not consistently outperform the single Weibull, as the error oscillates between positive and negative values.

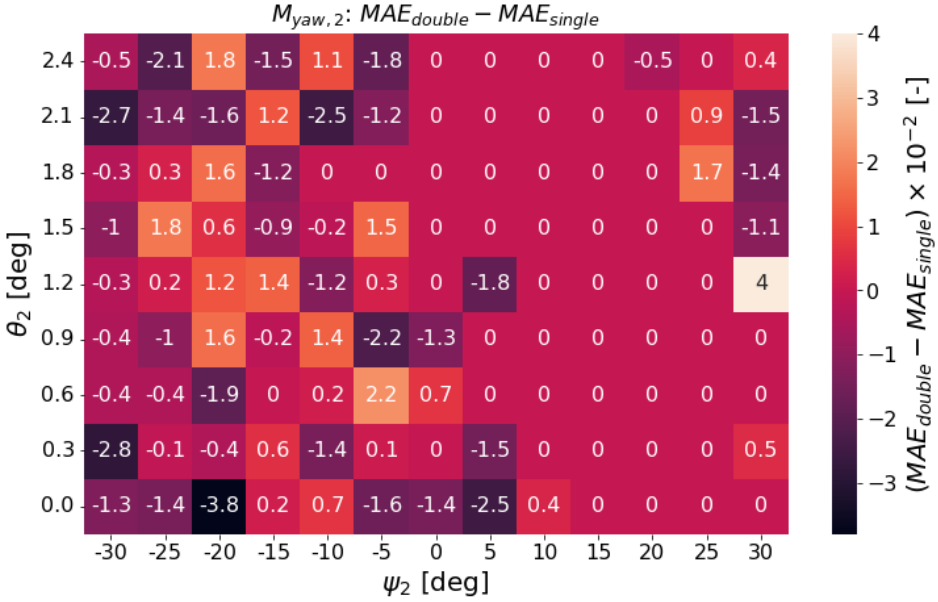

Figure 4.6: Difference in mean absolute error between the double Weibull and single Weibull fit for $S_x = 5R$ and $\psi_1 = \theta_1 = 0°$ for the downstream wind turbine tower top yaw moment channel.

The example above is focused on the downstream tower top yaw moment $M_{yaw,2}$ at given spacing of $S_x = 5R$. Similar observations have been made using different power and load channels at varying turbine downstream distances S_x. In the example above, the areas of instability are concentrated on the negative yaw ψ_2 side of the heat maps. However,

this is not the case for all channels and downstream distances S_x. Patches of parameter instability occur randomly throughout the whole input domain.

As the double Weibull has a unstable behaviour across the parameter domain, often times converges to a single Weibull, and does not always show an increased performance compared to the single Weibull approximation, the remainder of this study will be conducted using solely the single Weibull approximation. This decision will result in a loss of flexibility and accuracy of the model, however it guarantees a smooth parameter domain to create surrogate models with.

Further work could done regarding the distribution selection, as only the single Weibull and double Weibull were considered. Moreover, multiple other approaches could be taken. For example, a shift parameter could be introduced to the distributions. One could also fit a second (or third) distribution to the error of the fit of the first distribution. However, since the error can be negative, this could lead to negative distributions. Finally, one could also select multiple distributions with a fixed location.

4.1.1.2 Overview of Fitting Error

The fitting error is normalized with respect to the median output of each channel for the baseline case defined in Section 4.2.1. The baseline corresponds with inputs $\psi_1 = \theta_1 = \theta_2 = 0°$ and $\psi_2 \approx 0.8°$ (depending on the spacing S_x). The normalization is performed as shown in Equation (4.2), where $X_{Fit.CDF}$ and $X_{Emp.CDF}$ are the physical quantities as obtained from the fitted and empirical CDFs respectively, $X_{BL,median}$ is the median output of the baseline and X is the chosen channel.

$$\epsilon\ [\%] = \frac{X_{Fit.CDF} - X_{Emp.CDF}}{X_{BL,median}} \cdot 100 \tag{4.2}$$

The single Weibull fits are generated for each channel listed in Table 3.6 for both upstream (*WT1*) and downstream (*WT2*) wind turbines. The accuracy of the fits are evaluated by determining the fitting error, as defined defined in Section 4.1.1, throughout the whole input domain. The resulting fitting error is then normalized according to Equation (4.2). The fitting errors are presented in a $Mean \pm SD$ format, where the mean shows the bias of the error and the standard deviation (SD) gives an indication of how much this error varies. The results are summarized in Table 4.1 for every second spacing. As the $WT1$ case only consists of one wind turbine, no values for spacing S_x and P_{tot} are determined.

Comparing the errors of the $WT1$ case with the the mean fitting of the $WT2$ cases with varying spacings, it appears the $WT1$ case does not outperform the $WT2$ cases for all channels (e.g. M_{BRF} and M_{TT}). This is an unexpected result, as one would expect more complex empirical CDFs for the downstream turbine due to the aerodynamic interactions with the upstream turbine. As a result, it could be expected that the single Weibull fits would result in a larger error. Furthermore, all channels (excepted for the P_{tot} and P channels) have a decreasing standard deviation of the fitting error with increased spacing. A possible explanation for this is that the effect of the wake of the upstream turbine on

the downstream turbine is reduced, which causes the downstream turbine to operate in more steady and less turbulent conditions. However, the mean (*or bias*) of the fitting error behaves differently with spacing for different channels. For some channels, such as M_{MS}, M_{yaw}, and M_{TT}, a decreasing error is observed in function of distance. For other channels, such as P, an increasing error is observed. A final observation that can be made is that the WT2 surrogate errors are nearly all positive, excluding the P_{tot} channel. Visually, this means that on average the fitted CDF is fitted slightly to the right side of the empirical CDF. This means that for a given channel value, the fit will predict a lower cumulative probability. Essentially, the values predicted by the fitted CDF will be larger on average than the actual samples.

The M_{BRE} channel has a much smaller error compared to the other short-term equivalent load channels. The reason for this behaviour is that the load is dominated by gravity, resulting in very regular and consistent oscillations within its time signal. This is illustrated in Figure 4.7, where a segment of the time signal of the rotor rotational position (γ), blade root edgewise bending moment (M_{BRE}) and blade root flapwise bending moment (M_{BRF}) is shown. It can be seen that the oscillations in the M_{BRE} time signal nearly completely correspond to the rotor rotational frequency. This can be further visualised in Appendix D, where a very strong peak is observed at the frequency of the rotor rotation. Since the oscillations in the M_{BRE} time signal are so stable and dominated by the effects of gravity on the rotor, the outputs of the rainflow algorithm on the 10-minute time signals vary very little.

In contrast, the time series of the M_{BRF} channel is much more complex. The oscillations due to the rotor rotation can be seen, but other larger and irregular disturbances are present as well. A smaller peak can be observed in the spectral analysis of the M_{BRF} at the frequency of the rotor rotation in Appendix D. These larger disturbances occur due to the gust and turbulence present in the flow. As the oscillations are less smooth, the output of the rainflow algorithm varies strongly for every 10-minute time signal. The other loads for which the short-term equivalent load was determined vary similarly to the M_{BRF} channel.

Table 4.1: Overview of the normalized fitting error for upstream wind turbine (WT1) and downstream wind turbine (WT2) for different spacings S_x. The errors are given in a $Mean \pm SD$ format, such that the mean represents the bias of the error and the standard deviation gives an indication of how much the error varies throughout the input domain.

Model	S_x [R]	P_{tot} [%]	P [%]	M_{BRF} [%]	M_{BRE} [%]	M_{MS} [%]	M_{yaw} [%]	M_{TT} [%]	M_{TB} [%]
WT1	/	/	-0.06±0.34	0.78±4.52	-0.0±0.04	0.3±2.19	0.2±3.61	0.33±2.8	0.02±1.06
WT2	4	0.02±0.69	0.02±0.31	0.1±0.7	0.14±1.44	0.07±0.36	0.19±1.35	0.03±0.47	0.3±2.61
WT2	6	-0.01±0.76	0.01±0.41	0.12±0.76	0.12±1.07	0.1±0.51	0.22±1.68	0.03±0.45	0.11±2.28
WT2	8	0.01±0.87	0.02±0.54	0.05±0.63	0.05±0.81	0.15±0.58	0.32±1.78	0.03±0.42	0.22±2.7
WT2	10	-0.02±0.9	0.03±0.63	0.03±0.43	0.03±0.59	0.11±0.46	0.17±1.0	0.0±0.34	0.04±2.33
WT2	12	-0.01±0.96	0.05±0.7	0.04±0.41	0.03±0.54	0.06±0.36	0.07±0.68	-0.01±0.26	-0.01±2.41
WT2	14	-0.01±1.0	0.06±0.71	0.07±0.48	0.07±0.66	0.04±0.3	0.06±0.51	0.01±0.16	0.1±1.79
WT2	16	-0.01±0.97	0.06±0.69	0.05±0.42	0.06±0.61	0.05±0.34	0.06±0.52	0.02±0.12	0.13±1.62
WT2	18	0.05±0.97	0.08±0.7	0.06±0.4	0.06±0.57	0.05±0.36	0.05±0.51	0.01±0.14	0.07±1.46
WT2	20	0.07±0.99	0.11±0.73	0.07±0.4	0.08±0.62	0.04±0.33	0.04±0.46	0.01±0.16	0.05±1.35

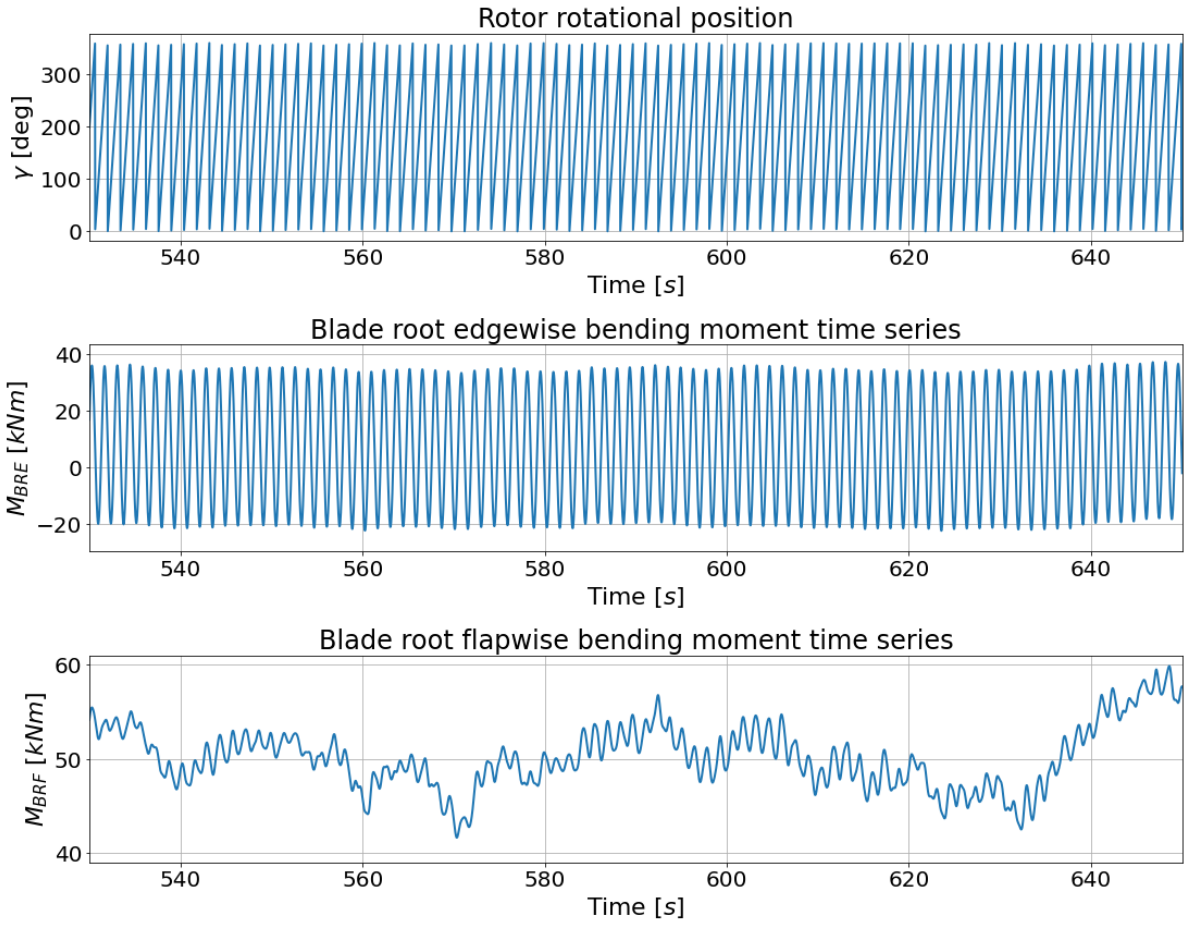

Figure 4.7: Partial time series of the rotor rotational position γ (top), the blade root edgewise bending moment $M_{BRE,1}$ (middle) and blade root flapwise bending moment $M_{BRF,1}$ (bottom) for the upstream wind turbine under normal operations ($\psi_1 = \theta_1 = 0°$). The $M_{BRE,1}$ channel can be seen to follow the cyclic behaviour of the rotor rotational position, whilst the $M_{BRF,1}$ shows a less regular behaviour.

4.1.2 Calibration Error and Surrogate Model Order Selection

The previous section discussed the *fitting error* of the MLE-based CDF fit with a single Weibull to the empirical CDF. The next error that can be found in Figure 4.1 is the *calibration error*. The calibration error quantifies the error between the MLE fitted distribution parameters and the distribution parameters generated through the surrogate models. The calibration error is used to fine-tune the input orders of the surrogate models in the following sections for the upstream and downstream wind turbine respectively.

4.1.2.1 Upstream Wind Turbine Surrogate

The upstream wind turbine surrogate, showed in Figure 3.16, consists of two sub-models, namely the shape (k) and scale (A) surrogate models. Each of these models take both the yaw angle ψ_1 and blade pitch angle θ_1 as inputs and output the shape or scale parameter of the distribution. These surrogate models have to be calibrated using the MLE fitted CDF parameters.

As described in Section 3.3.3, the order of each input of the surrogate model has to be set. It is essential to select the right order for each parameter. The order has to be high enough such that the tendencies of the data can be captured, but the order can not be so high that the data is overfitted. The shape surrogate model is used as an example to illustrate the method used to determined the surrogate model input orders. The shape surrogate model's inputs are denoted as $\psi_{1,k}$ for the yaw angle and $\theta_{1,k}$ for the blade pitch angle input. A grid search is conducted with orders ranging from 1 to 8 for each input. This upper limit is selected as orders higher than 8 are undesired and would increase the complexity of the model too much. For each input order combination in the grid, a surrogate model is constructed using 80% of the calibration data (i.e train data). The remaining 20% of the data (i.e. test data), is used to determine the Mean Absolute Error (MAE) of the surrogate model outputs with the MLE fitted CDF parameters.

Taking the shape parameters as an example, the MAE is determined for every combination of k_{fit} and k_{SM} that is produced with the test data. This procedure is repeated 5 times with alternating slices, such that each slice is used once as the test data, thus performing a 5-fold cross-validation. The mean and standard deviation of the 5-fold cross-validation of the MAE are computed for each set of input order combinations. Finally, the logarithm of the mean and standard deviation of the 5-fold cross-validation MSE is taken. The results of the grid search are shown for the power P_1 channel for the shape and scale parameters in Figure 4.8. The test MAE for the shape and scale parameters are shown in the upper and lower two heat maps respectfully. The left and right heat maps show the mean and Standard Deviation (SD) of the 5-fold cross-validation test MAE with the ψ_1 and θ_1 on the x- and y-axis respectively. A values between -1 and -2 in the heat maps indicate an error in the range of 0.1 to 0.01. These errors are quite small, but it is important to consider that the error is determined on *scaled* data, meaning that the range of the data is from 0 to 1, as described in Section 3.2.2.3.

Using these heat maps, a fitting order for the surrogate model inputs can be selected. The mean test MAE value has to be more or less converged to its minimal value and the standard deviation of the MAE can not be abnormally large compared to the other order combinations standard deviations. For the shape parameter sensitivity analysis for the P channel shown in Figure 4.8, an converged test mean MAE of -2.3 is observed for orders of 4 and 2 for the $\psi_{1,k}$ and $\theta_{1,k}$ inputs respectively. The standard deviation of the test MAE has also reached its near minimal value of -3.6 at for these orders. Using a similar reasoning for the scale parameters, orders of 6 and 3 are selected for the $\psi_{1,A}$ and $\theta_{1,A}$ inputs respectively. This example is used to illustrate the order selection approach as its features the largest convergence order of all the other channels, as can be seen in

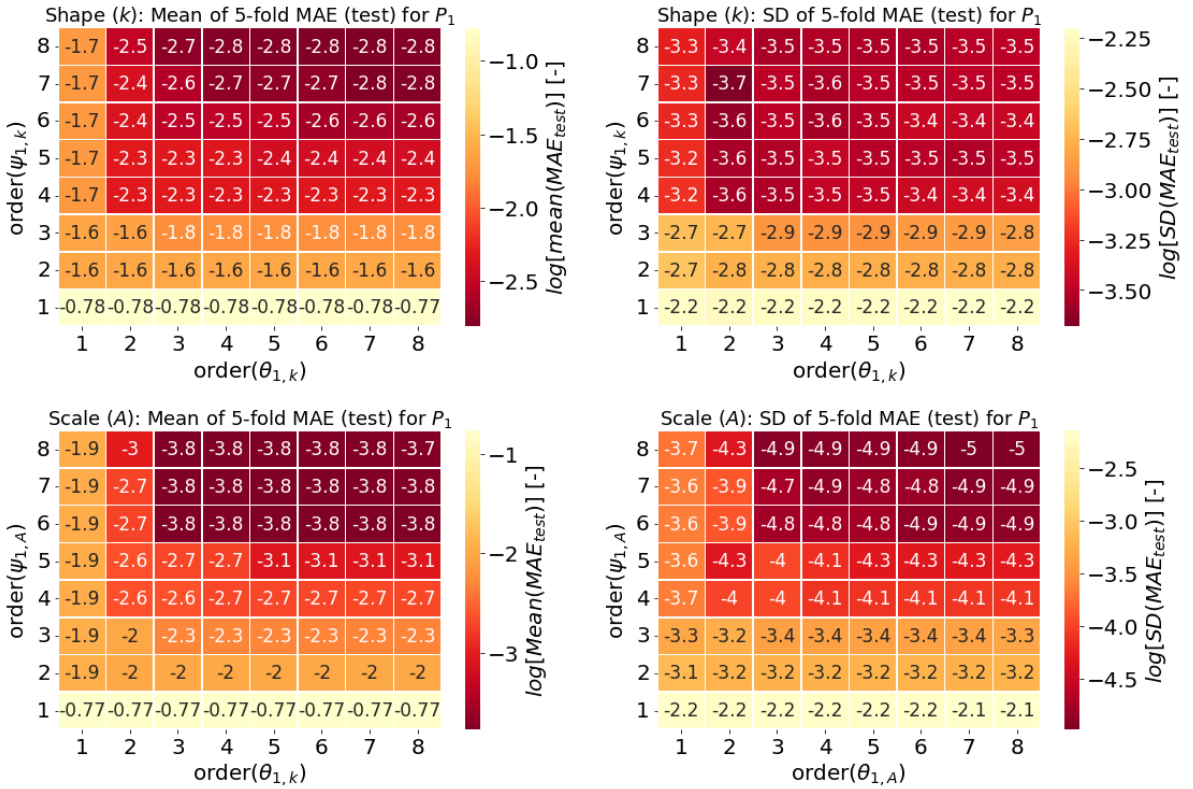

Figure 4.8: Heat maps of the test data MAE of shape k (top) and scale A parameters (bottom) for power channel P_1 in function of input order. The left column shows the means of the MAE and the right column the standard deviation of the MAE. These heat maps are used to determine the converge of the error with increasing input order.

Table 4.2, where the resulting orders for the different channels of surrogate model of the upstream wind turbine are shown.

4.1.2.2 Downstream Wind Turbine Surrogate

Referring back to the structure of the complete wind farm surrogate model in Figure 3.17, it can be observed that it consists of both a shape and scale surrogate models which both take four parameters, namely ψ_1, θ_1, ψ_2, θ_2. The order of each of these inputs for each surrogate model have to be determined. The approach is very similar to the one used for the upstream wind turbine surrogate model. However, rather than repeating the grid search in a four dimensional space, the orders of the downstream wind turbine inputs (i.e. ψ_2 and θ_2) are set the the same inputs shown in Table 4.2 for the upstream wind turbine parameters (i.e. ψ_1 and θ_1), as these orders were deemed sufficient to model the effects of the wind turbine's own inputs on its own power and loads in Section 4.1.2.1. This assumes that the effects of the upstream turbine wake on the downstream turbine channels can be fully captured by additional upstream wind turbine inputs of the surrogate models. This assumption greatly reduces the size of the grid search that has to be performed to

Table 4.2: Overview of upstream wind turbine surrogate model input order determined from sensitivity analysis. Both the input parameter orders for the shape surrogate (left columns) and scale surrogate (right columns) are shown.

	Orders			
	$\psi_{1,k}$	$\theta_{1,k}$	$\psi_{1,A}$	$\theta_{1,A}$
P_1	4	2	6	3
$M_{BRF,1}$	2	2	4	2
$M_{BRE,1}$	6	3	6	3
$M_{MS,1}$	2	1	2	2
$M_{yaw,1}$	2	2	4	2
$M_{TT,1}$	2	1	4	2
$M_{TB,1}$	2	1	2	2

tune the input orders for a minimized surrogate model error. As a result, only the orders of the $\psi_{1,k}$, $\theta_{1,k}$, $\psi_{1,A}$, and $\theta_{1,A}$ have to be tuned, where the second subscript indicates whether the inputs is for the scale (k) or shape (A) surrogate model.

A similar approach to the one used for the upstream wind turbine surrogate is used. Heat maps are constructed with the orders of inputs of both surrogate models, with the orders of ψ_1 and θ_1 and the x- and y-axis respectively. These heat maps are constructed for both the mean and standard value of the train and test data of MAE between the parameters obtained from the fitted CDF and the surrogate model outputs. Thus, separate heat maps are constructed for MAE of the shape and scale parameters.

Contrarily to the case of the upstream wind turbine surrogate, the wind farm surrogate only has a very limited amount of combined cases for the upstream wind turbine. As listed in Table 3.4, only four combined cases are available. As a result, there is a major concern that the surrogate will be overfitted to these four data points, whilst being unstable in other parts of the parameter domain. In order to avoid this, a 4-fold cross validation is used. For each fold, a specific combined LES scenario is omitted from the calibration data of the the surrogate models. The performance of the surrogate models constructed with the given set of orders is then evaluated with the omitted combined LES case. For a given channel and set of order inputs, the MAE error between the fitted distribution parameters and the surrogate model output parameters is determined for all the spacings S_x. This error is referred to as the *train* MAE, as the outputs of the surrogate are compared to the data it was calibrated with. The MAE is determined with all the different spacings S_x at once, as for simplicity's sake, it is desired to have surrogate models with similar input orders for every spacing. The *test* MAE is determined by determining the error between the fitted distribution parameters of the omitted LES case with the surrogate model output distribution parameters. Determining the test MAE is essential to avoid overfitting the surrogate models.

The resulting heat maps are shown for the total power P_{tot} of the wind farm for the

shape surrogate in Figure 4.9. The upper two plots show the logarithmic mean (left) and standard deviation (right) of the 4-fold cross-validation *test* MAE. The lower two plots show the logarithmic mean (left) and standard deviation (right) of the 4-fold cross-validation *train* MAE. Note that all the four heat maps are for the shape parameter. The heat maps of the train MAE are added to this analysis, as it gives an indication on whether the surrogates are overfitting the data. If the train error is disproportionally smaller than the test error, overfitting is very likely. Divergent errors (> -1), occurring when large input orders are used, are masked from the MAE heat maps, as these indicate drastic overfitting. These masks do no occur for the train MAE heat maps, as their errors decreases with increasing inputs order due to the overfitting of the data. The mean test MAE error appears to be minimal for orders of 2 for the $\psi_{1,k}$, $\theta_{1,k}$ inputs. Its associated standard deviations is also relatively low compared with the other input order combinations. As a result, these orders will be used for the $\psi_{1,k}$, $\theta_{1,k}$ inputs of the wind farm total power shape surrogate model. A similar approach is used for the scale surrogate models and the other channels. The resulting orders are shown in Table 4.3. Note that the orders for $\psi_{2,k}$, $\theta_{2,k}$, $\psi_{2,A}$, and $\theta_{2,A}$ in Table 4.3 are the same as the orders found in Table 4.2 for $\psi_{1,k}$, $\theta_{1,k}$, $\psi_{1,A}$, and $\theta_{1,A}$.

Table 4.3: Overview of wind farm surrogate model input order determined from sensitivity analysis. Both the input parameter orders for the shape surrogate (left columns) and scale surrogate (right columns) are shown.

	WF surrogate model input orders							
	Shape (k)				Scale (A)			
	$\psi_{1,k}$	$\theta_{1,k}$	$\psi_{2,k}$	$\theta_{2,k}$	$\psi_{1,A}$	$\theta_{1,A}$	$\psi_{2,A}$	$\theta_{2,A}$
P_{tot}	2	2	4	2	3	3	6	3
P_2	1	3	4	2	3	3	6	3
$M_{BRF,2}$	3	1	2	2	1	3	4	2
$M_{BRE,2}$	4	1	6	3	4	2	6	3
$M_{MS,2}$	2	1	2	1	2	1	2	2
$M_{yaw,2}$	1	3	2	2	4	1	4	2
$M_{TT,2}$	1	4	2	1	4	1	4	2
$M_{TB,2}$	1	4	2	1	2	2	2	2

4.1.3 Surrogate Model Error

Once the shape and scale surrogates have been calibrated and their input order set, the overall performance of the surrogate models can be assessed by means of the surrogate model error. This error is a direct consequence of the fitting and calibration errors measured along the way. The surrogate model error is defined between the empirical CDF and surrogate model generated CDF in a similar fashion as the fitting error, as described in Section 3.2.3. The errors are normalized with the baseline case described in Section 4.2.1. The normalization is performed as shown in Equation (4.3), where $X_{SM.CDF}$

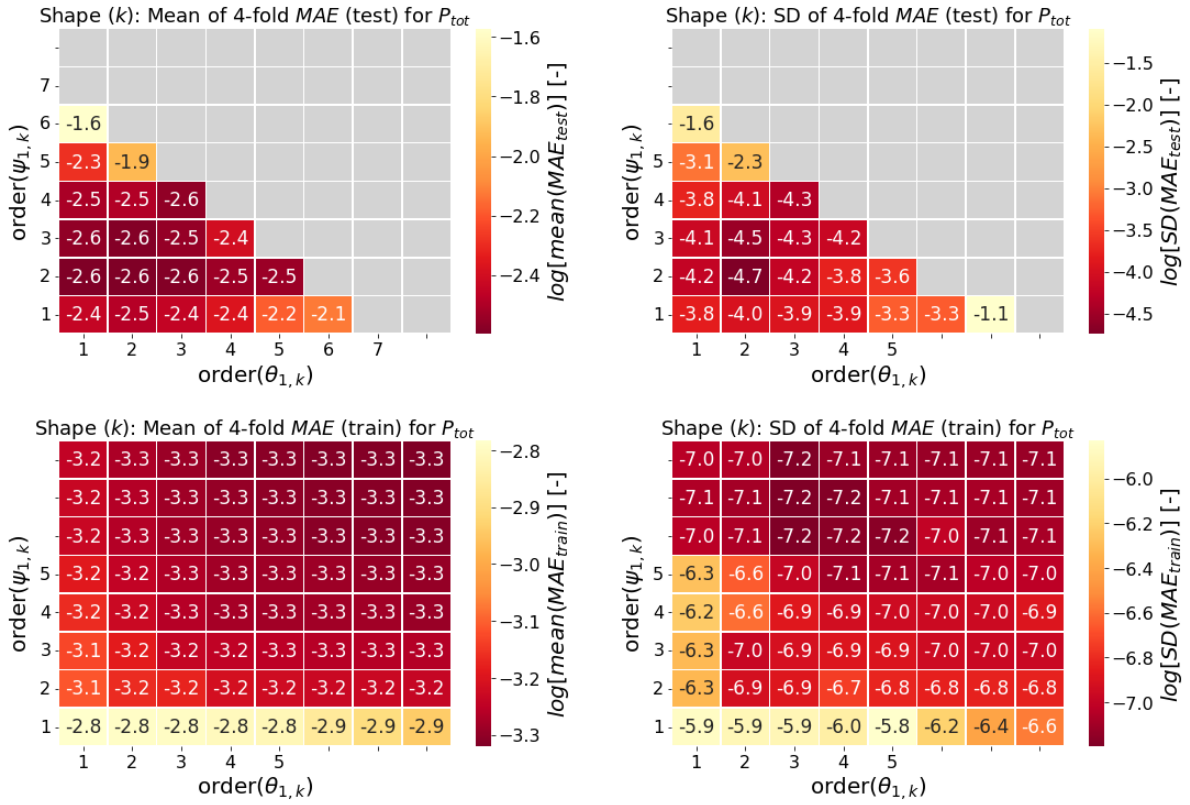

Figure 4.9: Heat maps of the test data (top) and train data (bottom) MAE of shape k (top) parameter for total power channel P_{total} in function of input orders. The left column shows the means of the MAE and the right column the standard deviation of the MAE. These heat maps are used to determine the converge of the error with increasing input order.

and $X_{Emp.CDF}$ are the physical quantities as obtained from the surrogate-generated and empirical CDFs respectively, $X_{BL,median}$ is the median output of the baseline and X is the chosen channel.

$$\epsilon = \frac{X_{SM.CDF} - X_{Emp.CDF}}{X_{BL,median}} \tag{4.3}$$

An overview of the surrogate model errors over the whole input domain is shown in Table 4.4 for the upstream (WT1) and downstream wind turbine (WT2) channels. The error is defined as described in Section 4.1.1 and is given in a $mean \pm SD$ format. The error of the downstream turbine surrogate model is shown for every second spacing S_x.

Examining Table 4.4, the upstream surrogate model appears to have a better overall performance, excepted for the M_{yaw} and M_{TT} channels. For both these channels, a larger mean error and standard deviation is observed. Very small biases are generally observed for the WT1 channels, certainly for the M_{BRE} where the absolute value of the bias is

smaller than 0.005%. The reason for this small error compared to other load channels is discussed in Section 4.1.1.2. The M_{BRF} and M_{MS} channels perform considerable worse than the other channels at small spacings, as the standard deviations of the error reaches values larger than 10%. These large standard deviations appear drop as as the spacing S_x is increased. A similar phenomenon is observed for the standard deviations of the other channels, excepted for the total power channel P_{tot}. A possible explanation for this is that the effect of the wake of the upstream turbine is reduced at larger spacings, resulting in a more steady and less turbulent flow for the downstream turbine, as discussed in Section 4.1.1.2 for the fitting error. Unsurprisingly, as the surrogate CDF are essentially calibrated to match the MLE fitted CDFs, the same positive bias tendency is observed. This means that the CDF generated through the surrogate models will on average be slightly to the right of the empirical CDF. Thus, it appears that the surrogate models have a tendency to overpredict the outputs associated with a given cumulative probability of the channels.

The table shown in Table 4.4 show the error of the surrogate models that have been calibrated using all the available data. The error has been normalized as defined in Section 4.1.1 These surrogate models and their errors are referred to as *full* surrogate models and *full* errors. The table gives very little indication on whether the data is being overfitted. Particularly the WT2 channels, as there are only four LES scenarios with combined upstream turbine inputs. Thus, it is a reasonable concern that the surrogate models may be overfitted to these four scenarios. In order to investigate whether this might be the case, a 4-fold cross-validation approach is used. The approach is similar to the one used in Section 4.1.2.2, where one of the combined cases is excluded from the calibration data for each fold. As a result, for each fold a *train* and *test* error can be determined. The distributions of the agglomerated train and test data of the 4-fold cross-validation can be visualised. If there is a large difference in the mean and spread of the train and test error, it is very probable that the surrogate models are overfitting the data.

The error distributions of the different channels are shown in Figure 4.10 with their associated error box plots. The outliers of the box plots are not shown. The train, test and full surrogate model errors are shown in blue, orange, and black respectively. Only a very small difference is observed between that train and full errors, as the difference in calibration data comes down to a single combined case. The full error is added to the figure nevertheless, as the training error only represents a conservative error of the full surrogate model error. Furthermore, it should be noted that the train and error distributions are a result of 4 different surrogate model, as this is the agglomerated error of the 4-fold cross-validation.

Although the errors vary in scale for each individual channel, the box plots appears to indicate that there are no signs of overfitting. However, the test error distributions of the total power P_{tot} and downstream turbine power P_2 channels stand out. It appears that the test distribution shows signs of bimodality and do not coincide fully with the train error distribution. This is further investigated in Appendix H, where the train and test

error are plotted separately for each fold within the 4-fold cross-validation. Within this analysis, it appears that the bimodel behaviour is associated with larger yaw angels. It is worth noting that this bimodel behaviour is only observed for the channels P_{tot} and P_2, which have a relatively small error scale compared to the other channels. Other channels, such as M_{BRF}, M_{MS}, and M_{TB} show distribution ranging to more than 10%.

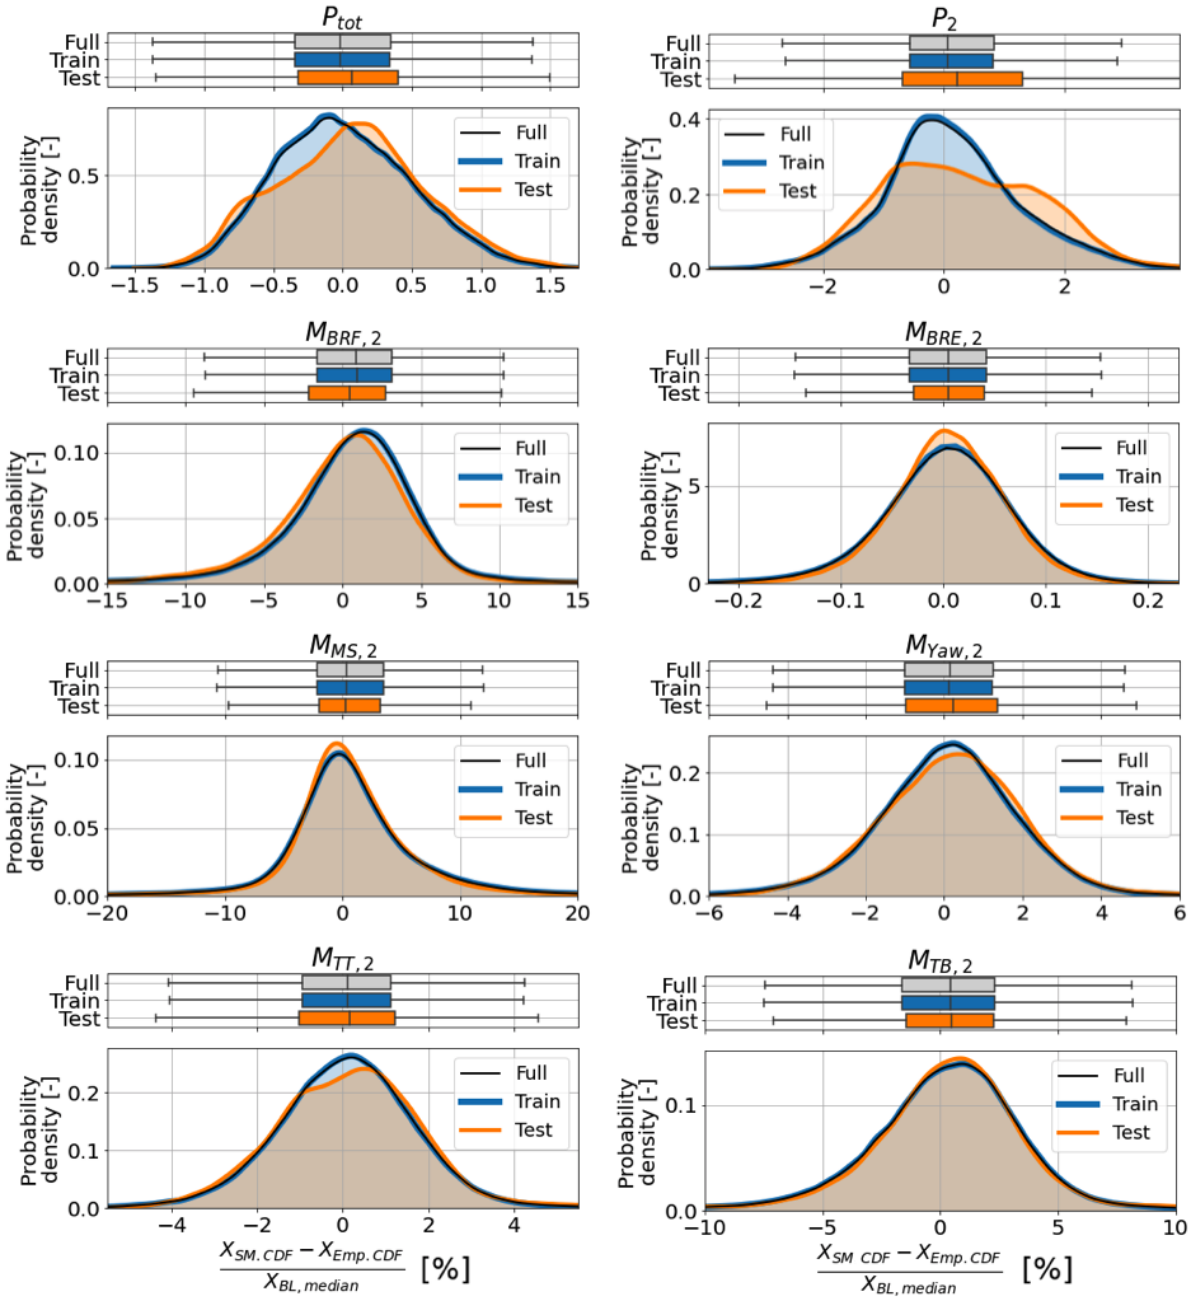

Figure 4.10: The error distribution of the full surrogate model (gray) and partial surrogate with the train (blue) and test (orange) data for P_{tot}, P_2, M_{BRF_2}, M_{BRE_2}, M_{MS_2}, M_{yaw_2}, M_{TT_2}, and M_{TB_2} channel.

Table 4.4: Overview of normalized surrogate model error for upstream wind turbine (WT1) and downstream wind turbine (WT2) for different spacings S_x. The errors are given in a $Mean \pm SD$ format, such that the mean represents the bias of the error and the standard deviation gives an indication of how much the error varies throughout the input domain.

Model	S_x [R]	P_{tot} [%]	P [%]	M_{BRF} [%]	M_{BRE} [%]	M_{MS} [%]	M_{yaw} [%]	M_{TT} [%]	M_{TB} [%]
WT1	/	/	-0.06±0.34	0.19±1.24	-0.0±0.02	0.02±0.17	0.04±0.58	0.05±0.45	0.1±1.91
WT2	4	0.03±0.45	0.13±1.49	0.9±10.62	0.01±0.28	1.53±10.05	0.29±2.99	0.2±2.75	0.6±7.32
WT2	6	-0.0±0.48	0.1±1.31	0.75±5.12	0.02±0.35	2.91±16.58	0.16±2.88	0.24±3.14	0.15±3.52
WT2	8	0.01±0.51	0.11±1.22	0.38±3.86	0.01±0.21	2.09±10.02	0.11±1.49	0.12±1.43	0.28±3.2
WT2	10	-0.01±0.51	0.14±1.15	0.35±3.67	0.02±0.35	1.23±5.16	0.07±1.41	0.05±1.33	0.18±2.73
WT2	12	-0.0±0.52	0.18±1.11	0.34±3.67	0.01±0.18	0.82±5.82	0.06±1.46	0.01±1.38	0.23±3.04
WT2	14	0.0±0.51	0.19±1.07	0.46±4.14	0.0±0.06	0.63±5.53	0.13±1.43	0.08±1.18	0.2±2.73
WT2	16	-0.0±0.49	0.18±1.03	0.38±3.8	0.0±0.06	0.63±4.1	0.15±1.59	0.14±1.62	0.19±2.77
WT2	18	0.03±0.48	0.2±1.01	0.39±3.48	0.0±0.05	0.56±4.0	0.09±1.83	0.1±1.89	0.15±2.7
WT2	20	0.04±0.48	0.24±1.02	0.43±3.6	0.0±0.13	0.42±3.39	0.07±2.0	0.1±2.03	0.13±2.83

4.2 Wind Farm Power Based Optimization

Two control approaches are considered for the purposes of the power based optimization of the wind farm. The first approach, shown on the left side of Figure 4.11, aims to optimize the power output of each wind turbine individually. In other words, the upstream wind turbine will optimize its yaw offset and down-regulation level such that its own performance is optimized without considering its adverse wake effects on the downstream turbine. The downstream wind turbine will try to optimize its performance within the wake of the upstream turbine. The total power is then the sum of the individually optimized upstream and downstream power channels. This control approach is referred as *greedy* control and will serve as the baseline, as discussed in Section 4.2.1.

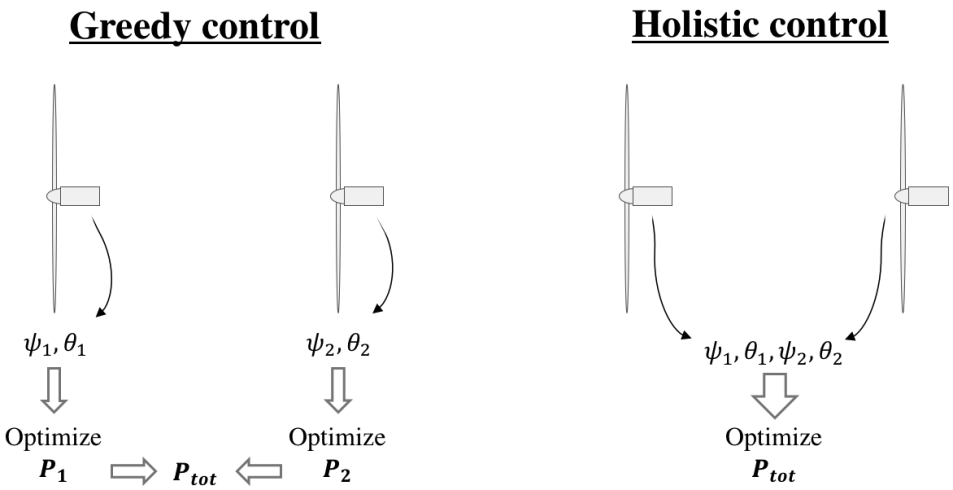

Figure 4.11: Diagram showing how the wind farm parameters of the greedy (left) and holistic (right) control approaches are optimized. The greedy approach optimizes the wind turbine settings for the wind turbines individually, whilst the holistic approach optimizes all four wind farm settings at once.

In the second approach, as shown on the right side of Figure 4.11, the upstream wind turbine will try to optimize its own performance by means of yaw-control and down-regulation. However, it will also take into account the effects of its own wake on the downstream wind turbine. This is done by optimizing the total power of the wind farm using the yaw angles and down-regulation levels of both wind turbines at the same time. As a result, with this approach the overall performance of the wind farm is optimized and is therefore referred to as *holistic* control. Note that both control approaches refer to the control of the wind farm. On an individual level, the downstream turbine is greedy in both the greedy and holistic wind farm control approaches, as its input settings have no effect on the operations of the upstream turbine. Therefore, it will always tries to optimize its own operation within the upstream turbine wake conditions. The next section defines the baseline case which will be used to quantify the potential of the holistic wind farm control approach.

4.2.1 Baseline Case

The baseline case is defined by making using the greedy control approach mentality. The first step consists in determining the greedy wind turbine settings of the upstream wind turbine. The yaw angle of the upstream wind turbine ψ_1 is defined with respect to the wind direction used as input for the LES and is defined as positive in the clockwise direction. The wind direction of the LES has been set to be aligned with the lay-out of the two wind turbines, as shown in Figure 4.12. As the upstream wind turbine wishes to optimize its own operations, the upstream wind turbine yaw input is set to $\psi_{1,LES_{FW}} = 0°$ such that it is aligned with the LES input wind direction. The LES FrameWork (LES_{FW}) notation is used to indicate that the angle is measured with respect to the wind turbine alignment, which is the same as the input wind direction used for the LES. The level of down-regulation is introduced by a constant blade pitch angle offset. The blade pitch angle is defined as the increased angle with respect to the blade pitch angle for below-rated conditions under normal operation. This angle can only be positive. As the aim is to maximize the power of the upstream wind turbine without considerations for the loading, no level of down-regulation is introduced ($\theta_1 = 0°$) for the baseline case.

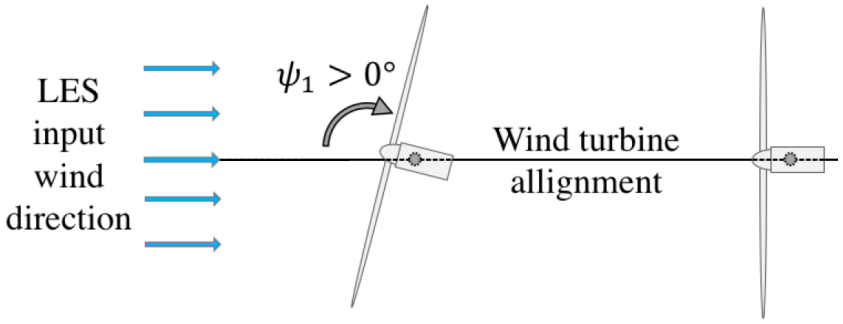

Figure 4.12: Sketch showing wind turbine alignment in order to define the LES FrameWork (LES_{FW}) and the defintion of the positive yaw angle ψ.

The second step consists in determining the greedy wind turbine settings of the downstream wind turbine. These settings are determined by means of the Wind Farm Surrogate Model (WF SM) presented in Section 3.2.4.2. Using the fixed settings of $\psi_{1,LES_{FW}} = \theta_{1,LES_{FW}} = 0°$ as inputs for the WF SM, the optimal downstream wind turbine settings are determined such that its own median power production is optimized. Similarly to the upstream turbine, the input for down-regulation is set to $\theta_2 = 0°$ as the downstream wind turbine wished to maximize its own power with no regards to the experienced loads.

The resulting yaw settings for the downstream wind turbine are shown in Figure 4.13. The x-axis represents the wind farm spacing and the y-axis represents the yaw angle of the downstream wind turbine with respect to the wind farm alignment, as defined in Figure 4.12. A small yaw angle setting for the downstream wind turbine is to be expected, as the presence of the upstream turbine will introduce a deflection in the wake

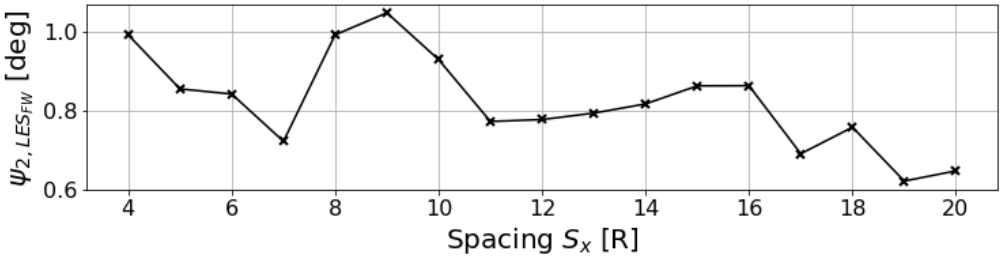

Figure 4.13: Baseline yaw angle of downstream wind turbine with respect to turbine alignment in function of spacing.

even with no yaw misalignment, as discussed in P. A. Fleming et al. (2014). However, yaw settings shown in Figure 4.13 are partly a result of the bias present in the flow direction. Although the input wind direction of the LES is set to be aligned with the wind farm lay-out, a bias is observed. The bias is shown for each spacing in Figure 4.14, where a spacing of 0R corresponds to the location of the upstream wind turbine. Each boxplot consists of 23 10-minute samples of the mean observed wind direction in the horizontal plane at the hub location. The flow data is obtained from the LES with no wind turbine present.

It can be observed that there is a negative bias within the flow direction, which causes the downstream wind turbine to set a positive yaw angle in order to align itself with the flow. The bias is originates from the inherent variability of the turbulent flow and the time constrained nature of the LES. If the LES were to be continued for an infite amount of time, the median value of the wind direction would converge to input wind direction of the LES. If the LES would be re-run with a different turbulence seed, a different bias would potentially be observed.

This negative bias is also present upstream of the first wind turbine ($S_x < 0R$). This means that in order for the upstream wind turbine to be aligned with the flow and maximize its power production, it would have to be yawed slightly positively ($\psi_{1,LES_{FW}} \approx +0.8°$). However, a positive upstream turbine yaw angle $\psi_{1,LES_{FW}} > 0°$ falls outside of the parameter domain of the WF SM, as discussed Section 3.1.5. As a result, the upstream

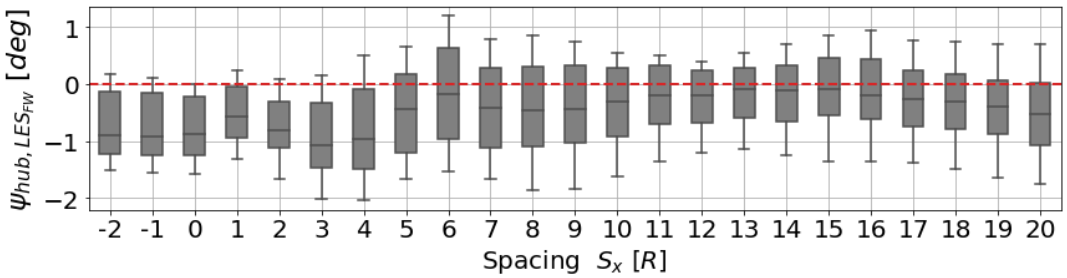

Figure 4.14: Distribution of the 10-minute mean wind direction at hub height in function of spacing for the LES with no wind turbine present.

wind turbine yaw angle is kept at $\psi_{1,LES_{FW}} = 0°$. However, this essentially means that the upstream wind turbine is not operating under its possible optimal conditions and is thus not *truly* under greedy control. The differences in power and loads between the *true* greedy approach ($\psi_{1,LES_{FW}} = +0.8°$) and *assumed* greedy approach ($\psi_{1,LES_{FW}} = 0°$) are determined by means of the upstream Wind Turbine Surrogate Model (WT SM) in Appendix I. Within this analysis, it can be seen that the differences are minimal in terms of the power and loads of the upstream wind turbine. As a result, the study is continued with the *assumed* greedy approach corresponding to $\psi_{1,LES_{FW}} = 0°$. The results presented in the following sections are presented with respect to the settings obtained from the above described baseline. This is not relevant for the downstream wind turbine yaw angle ψ_2, as the other inputs are already set to $\psi_1 = \theta_1 = \theta_2 = 0°$.

The power and loads are quantified by the normalized difference between the holistic control approach and the baseline, as shown for the power channel P in Equation (4.4). The same normalization procedure is used for all other channels. Note that this normalization process is applied individually for each spacing. Therefore, the results of the holistic optimization will show the gain or losses proportionally to the baseline at a given spacing. For example, a gain of +3% might correspond to a +40kW increase at a spacing of 4R, whilst only corresponding to a $+5kW$ increase at a spacing of 20R.

$$\Delta P\ [\%] = \frac{P_{holistic} - P_{baseline}}{P_{baseline}} \cdot 100 \tag{4.4}$$

Having defined the baseline, the optimization can be performed for the different holistic control approach cases. The next sections described the two scenarios that are used for this analysis and quantifies the results with respect to the hereabove defined baseline.

4.2.2 Holistic Optimization Cases

Two optimization cases are considered and compared with the baseline for the purpose of this study. Both cases consists in optimizing the median wind farm power by means of the wind farm settings using the holistic control approach. However, the first case does not consider the loads in the optimization process. The second case does consider the wind turbine loading by implementing a load constraint within the optimization process. The loads are included within the optimization in the form of upper boundary that limit the allowable loading increase with respect to the baseline case. The constraints is set at +5%, such that the relation shown in Equation (4.5) is respected for every load channel in the constrained optimization, where $M_{x,constr.}$ and $M_{x,BL}$ correspond to the constrained holistic optimization case and baseline case loads.

$$M_{x,constr.} \leq 1.05 \cdot M_{x,BL} \tag{4.5}$$

4.2.3 Surrogate Model Error Correction

Section 4.1.3 discussed the surrogate model error distributions over the whole input domain for each channel. The surrogate model error associated with a specific input set can also be retrieved. For each set of inputs listed in the table of input cases in Table 3.5, the empirical CDF can be constructed and the corresponding surrogate model error can be determined. However, it is highly unlikely that the optimized wind farm settings will be the exact same as the discrete cases for which empirical CDFs have been constructed. Therefore, a lookup table is used to interpolate the error linearly between the different cases. Furthermore, as each case is associated with a CDF, a fifth column is added to the look up table in order to be able to interpolate the error associated with each percentile of the distribution.

This approach has the downside that the load constraints are applied to the surrogate model outputs before the surrogate model error is added. As a result, the error-corrected surrogate model outputs might overshoot the constraints. This is seen and discussed more elaborately in Section 4.2.4.3 with the main shaft torque of the of the downstream wind turbine. A possible way to circumvent this issue would be to include the surrogate model in the optimization process, rather than adding it after the optimization has been performed. However, this has not been implemented due to a lack of data. As the look up table interpolates linearly between data points, it is requires that there is a fairly high concentration of data points across the domain. Only 4 LES simulations were used to model the upstream wind turbine combined yaw and blade inputs, one can not assume a linear behaviour of the error between these data points. Therefore, the error is only added after the optimization, as nearly all of the resulting optimized wind farm control settings fall within a well populated region in the parameter domain.

The surrogate model error is used to correct the surrogate model outputs for the optimized wind farm control settings. The results presented within the next sections have been corrected by means of the look up table, unless especially stated otherwise.

4.2.4 Optimization Results

The error-corrected results of the median power optimization of the different cases are shown in Figure 4.15. The baseline (BL), the unconstrained, and the load constrained optimization cases are shown in black, blue, and red respectively. This color code is maintained throughout the chapter. The upper left plot shows the difference in total power between the cases, normalized according to Equation (4.4) and error-correct by means the look up table. The corresponding shaded areas display the 25^{th} to 75^{th} percentile range. Note that these percentiles are not included during the optimization, as only the median total power values have been taken into account for this optimization process. The lower left plot shows the ratio of the power of the downstream wind turbine (P_2) over the power of the wind farm (P_{tot}) in percentage. A dotted line is added at 50%, which represents the scenario in which both turbine would produce the same amount of power. The plots in the right column from top to bottom correspond to the optimized ψ_1,

ψ_2, θ_1, θ_2 settings respectfully, which are measured with respect to the control settings of the base line case inputs.

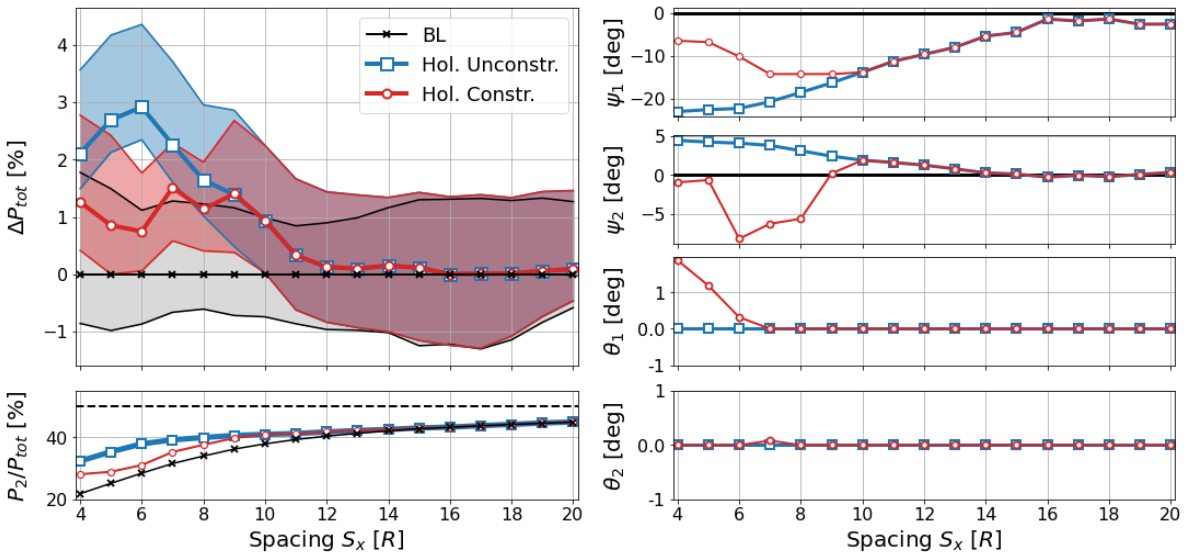

Figure 4.15: Optimization results of the unconstrained case (blue) and constrained case (red) with respect to the baseline (black). The median total power (solid lines) and interquartile ranges (shaded areas) are shown in the upper left plot. The proportional power production of the downstream turbine is shown in the lower left plot. The right columns shown the corresponding optimized wind farm settings.

By quickly going over the plots in Figure 4.15, two conclusions can be made. Firstly, whether it be unconstrained or constrained holistic case, both appear the be nearly converged with the baseline around $S_x = 12R$. Secondly, the constrained holistic case converges with the unconstrained holistic case at $S_x = 9R$, indicating that the load constraints were only relevant at close spacings.

The above optimization is performed such that the median value of the power output is maximized. One could also optimize a range of percentiles, rather than solely the median value. For example, the percentiles ranging from the 25^{th} to the 75^{th} percentiles could be used for the optimization. The optimizer would try to maximize the sum of the power output corresponding to each percentile within the range. A further step could consists in assignment different weights to the different percentiles, such that focus of the optimization could be set to be the median value, or the more extreme values. These approaches were implemented, but yielded very similar results to the median-based optimization presented in Figure 4.15. The reasons for these similar results is due to the stiffness of the single Weibull distribution. This is illustrated in Figure 4.16, where the total power P_{tot} distributions are shown for $S_x = 10R$ and $\theta_1 = \psi_2 = \theta_2 = 0°$. The orange and purple curves corresponds to the downstream turbine yaw input of $\psi_1 = 0°$ and $\psi_1 = -10°$ respectively. Note that these distributions have not been corrected by the surrogate model error. Despite the different inputs resulting in a difference in

median of approximately 1%, the shape of the distribution is very similar. It appears the distribution is merely shifted to the right. As a result of this unchanging shape, the median optimization result in very similar optimization results as the range optimizations. This issue could be mitigated by including the surrogate model error within the optimization process, by using more flexible distributions (e.g. double Weibull), and by using more a more flexible type of surrogate model (e.g. neural network). The results presented within this chapter are on basis of median optimization.

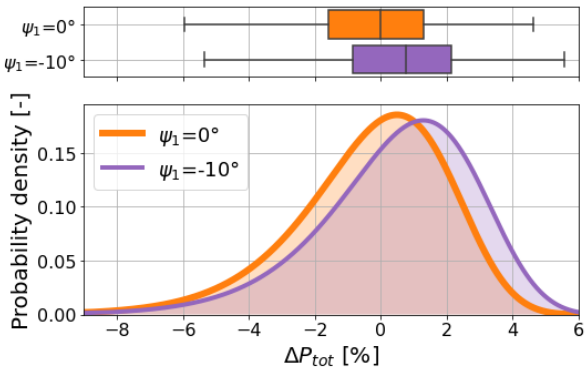

Figure 4.16: Total power ΔP_{tot} distributions and box plots for $S_x = 10R$, $\theta_1 = \psi_2 = \theta_2 = 0°$. The distributions corresponding to the $\psi_1 = 0°$ and $\psi_1 = -10°$ inputs are shown in orange and purple respectively. Note that these distributions are not error-corrected by the surrogate model error.

The following sections will dissect Figure 4.15 in three main parts by looking at the wind farm control settings in Section 4.2.4.1, the corresponding power outputs in Section 4.2.4.2, and the experienced wind turbine loading in Section 4.2.4.3.

4.2.4.1 Wind Farm Control Settings

Considering the yaw angles of the upstream wind turbine ψ_1, both optimization cases result in negative yaw angles. The negative upstream yaw is to be expected, as yawing counterclockwise has generally been observed to be more effective than yawing clockwise, as discussed in Section 2.5.2.2. However, the constrained case is yawed less significantly at closer spacings than the constrained case. The reduction in the yaw angle setting is a result of the load constraint. As will be discussed more elaborately later in this chapter in Section 4.2.4.3, yawing the upstream wind turbine to larger angles induces an increased blade root flapwise bending moment. Thus, for the constrained case the yaw angle is limited to smaller values.

The yaw angle of the downstream wind turbine ψ_2 is positive for the unconstrained case. Similar results have been observed in Hulsman et al. (2019). The positive yaw angle of the downstream turbine is a result of the wind turbine trying to realign itself with the skewed flow originating from the deflected wake of the upstream wind turbine and the bias in the flow direction discussed in Section 4.2.1. This is illustrated in Figure 4.17,

where each boxplot represents 23 10-minute samples of the observed mean wind direction with respect to the LES framework (see Figure 4.13) in the horizontal plane at the hub location. The spacing of 0R corresponds to the location of the upstream wind turbine. It can be seen that flow direction from the yawed scenario ($\psi_{1,LES_{FW}} = -15°$) induces a large skew in the wind direction than is not observed for the scenario with no input yaw angle ($\psi_{1,LES_{FW}} = 0°$). The difference in flow direction between these two scenarios over the spacing correspond with the positive downstream turbine yaw angles ψ_2 observed in Figure 4.15. This is a result of the wind turbine realigning itself with the flow. Furthermore, it can be observed in Figure 4.17 that the incoming flow of the upstream wind turbine ($S_x < 0R$) is slightly less negative for the yawed LES scenario. As the thrust force is at an negative angle with the incoming flow, it slightly changes the mean flow direction in the positive direction.

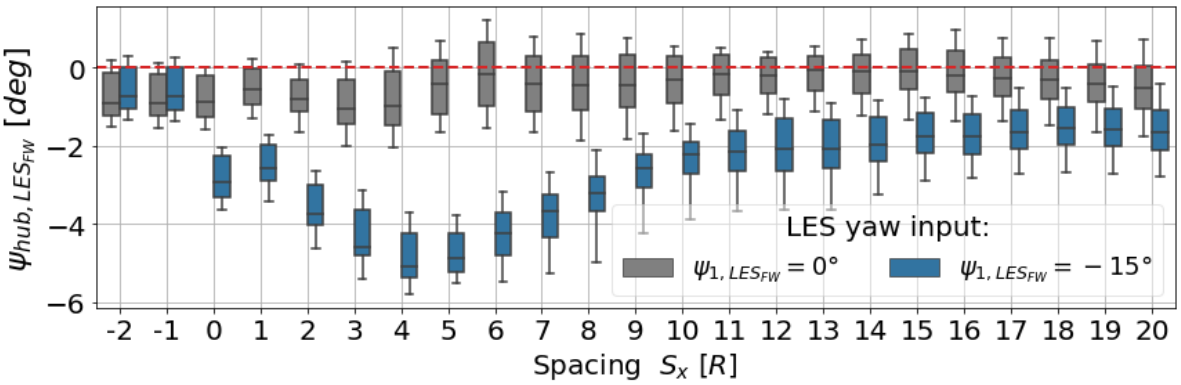

Figure 4.17: Distribution of the 10-minute mean wind direction at hub height in function of spacing for the LES with with a wind turbine with no yaw angle (gray) and a yaw angle of $\psi_{1,LES_{FW}} = -15°$ (blue).

In contrast, for the constrained case, the upstream wind turbine is yawed negatively. This negative yaw is a result of the wind turbine trying to reduce its loading. As discussed in the previous paragraph, the unconstrained downstream wind turbine will try and align itself with the flow by yawing positively. However, as the wind turbine is aligned with the flow, its aerodynamic loading will be increased. For the constrained case, the downstream wind turbine has the opposite behaviour by yawing even further away from the flow in order to reduce is loading such that the constraint are respected.

At a spacing of $S_x = 10R$, the yaw inputs converge to the inputs of the unconstrained case, resulting in the similar power output observed in Figure 4.15. The spacing at which the constrained case becomes the same as the unconstrained case depends on the severity of the constraint, as shown in Figure 4.18. It can be observed that the more severe the constraint, the larger the spacing at which the constrained and unconstrained cases merge. As for the upstream wind turbine blade pitch angle, no definitive relation between the magnitude of the input and the severity of the constraint can be observed.

The most severe constrained case has the smallest input at $S_x = 4R$, but the largest input at $S_x = 6R$.

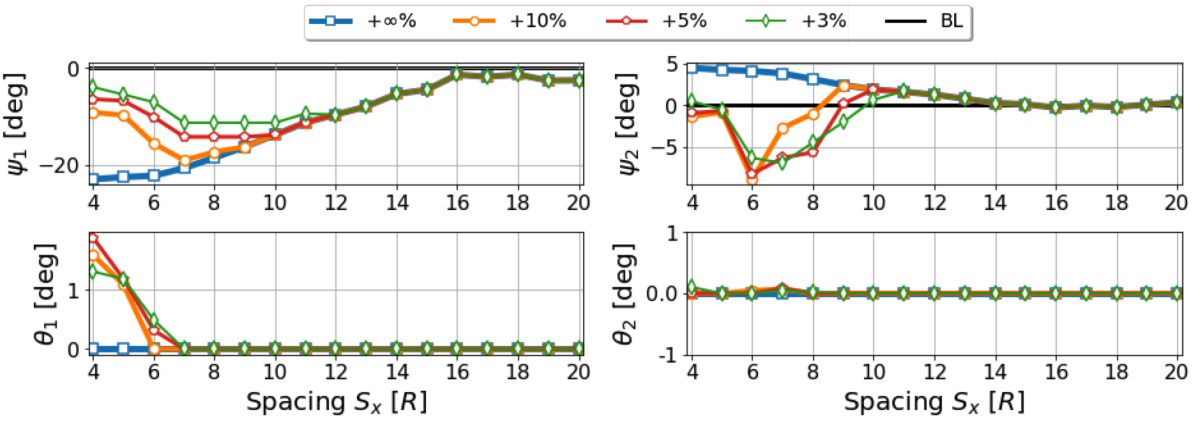

Figure 4.18: Optimized wind farm settings for the unconstrained case (blue), +10% case (orange), +5% case (red), and +3% case (green).

Further considering the resulting blade pitch angle settings of the wind farm shown in Figure 4.15, it appears that the only case that results in some form of down-regulation is the constrained case. However, the input is only given for spacings $S_x < 7R$. Observing the unconstrained case, it appears that the optimizer ideally wants the yaw the upstream wind turbine by approximately $\psi_1 = -20°$ for $S_x \leq 8R$. However, as is discussed in Section 4.2.4.3, the blade root flapwise bending moment increases with increasing negative yaw input ψ_1, causing it to reach the constraint. As a result the upstream yaw angle has to be reduced. The reduced yaw angle ψ_1 results in an general increase in aerodynamic loading, which might result in the other loads hitting the constraint. As a result, the the upstream wind turbine has to down-regulate itself in order to meet the constraints for the other loads. Down-regulation has only been observed for constraints smaller than approximately 35%. For the downstream wind turbine, the down-regulation of the upstream wind turbine will result in a more turbulent wake Annoni et al. (2016). However, this will accelerate the wake recovery, which could potentially increase the power production of the downstream wind turbine. Interestingly, the larger down-regulation inputs of the upstream wind turbine at $S_x = 5R, 6R$ correspond to the very limited yaw inputs of the downstream wind turbine. This might be due to the fact that the upstream wind turbine experienced power losses due to its down-regulation level, which as a result urges the downstream wind turbine to yaw into the flow direction in order to make up for the power losses of the upstream turbine.

It could have been speculated that the upstream turbine would have down-regulated itself ($\theta_1 > 0°$) for the unconstrained holistic case as well in an attempt increase the available energy in the flow for the downstream wind turbine. However, it appears that yaw-control is a more effective approach in this scenario. As down-regulation appears to be only applied when loads are considered, yaw-control can be assumed to be the dominating strategy to optimize the wind farm power for this particular wind farm with

the given conditions. The reason down-regulation is not used in the power optimization could be attributed to the its deteriorative effect on the effectiveness yaw-based control. The yaw-based wake deflection is caused by the misalignment of the thrust force with the flow-direction, resulting in a cross-wind thrust force component deflecting the wake from its original direction. If a level of down-regulation is introduced, the reduced thrust might render the yaw-control approach less efficient.

The overall blade pitch input of the downstream wind turbine is kept at approximately $\theta_2 = 0°$ for all cases for nearly every spacing. A negligible amount of blade pitch angle input $\theta_2 < 0.1°$ is observed for the constrained case at S_x. The only reason the downstream wind turbine would down-regulate itself is to reduce its loading, as there is no other wind turbine behind it. However, from the results it appears that yaw away from the flow direction is a more effective way to constrained the investigated loads.

4.2.4.2 Wind Farm Total Power

Analysing the results of the power optimization in the upper left corner of Figure 4.15, it can be seen that the largest differences between the cases occur at spacings $S_x < 12R$. For these spacings, power gains of up to +3% and +1.5% can be observed for the unconstrained and constrained holistic cases respectively. Similar results have been obtained in a similar study considering yaw-based control in Hulsman et al. (2019). Moreover, the unconstrained case shows signs of a slight positive skew as the medians lay closer to the Q25 boundary than Q75. A positive skew a beneficial, as it indicated that the probability of getting an exceptionally large increase in power is greater than getting an exceptionally large decrease in power. However, as will be concluded from Figure 4.20 later on in the chapter, the percentiles are misleading and the distributions actually have a relatively long tail on the left side. For larger spacings $S_x > 12R$, both holistic cases converge with the baseline. Thus, it appears that the holistic wind farm control strategies are only relevant at small inter-turbine distances. This outcome is not unexpected, as the wake interactions are larger at closer spacings, thus making the control strategies more relevant for these scenarios.

The constrained case merges with the unconstrained case at $S_x = 9R$, indicating that the load constraint is no longer relevant for larger spacings. The spacing at which the unconstrained case converges with the constrained case depends on the severity of the imposed load constraint. This is visualized in Figure 4.19, where the power outputs corresponding to different load constraints are shown. It appears that a more severe constraint (+3%) results in a later convergence with the unconstrained case (+∞%), whilst a looser constraint (+10%) results in earlier convergence. Thus, the wind turbine appear to experience larger increases in loads at closer spacings than at larger spacings, as will be discussed to a further extent at the end of the chapter.

The full distributions of the resulting total power P_{tot} of the baseline for the constrained and unconstrained cases are shown for the spacings $S_x = 4R$, $6R$, $8R$ and $10R$ in Figure 4.20. The distributions are constructed by using the look up table to error-correct the Weibull distributions obtained with the surrogate models. The distribution are

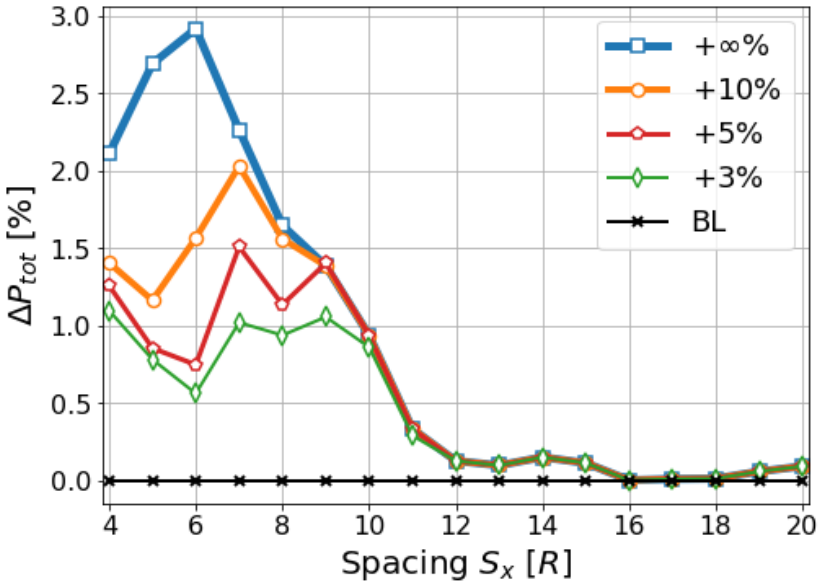

Figure 4.19: Median total power gains with respect to the baseline for the unconstrained case (blue), +10% case (orange), +5% case (red), and +3% case (green).

plotted in function of the expected percentage decrease and increase in total power with respect to the median total power of the greedy control approach, as defined in Equation (4.4). As a result, the median value of the BL is located at $\Delta P_{tot} = 0\%$, as indicated by the vertical line within the distribution.

The general tendencies observed in Figure 4.15 are also visible in the full power distributions in Figure 4.20. The unconstrained and constrained optimization cases converge to the baseline case with increasing spacing. Furthermore, the constrained case is converged with the unconstrained case at $S_x = 10R$. The distribution for larger spacings ($S_x > 12R$) are not shown, as all three distributions become nearly identical. Interestingly, the distributions of the baseline and constrained cases show some signs of bimodality for smaller spacings, indicating that there could be two main operating regimes. However, the signs of disappear when the spacing is increased. Furthermore, it was observed in Figure 4.15 that there appears to be a positive skew in the distributions. This skew is less visible when the whole distributions are inspected. It actually appears that the distribution have a relatively long tail on the left side, indicating that there is a risk of observing very low, or even negative, power gains.

Observing the proportional power output of the upstream and downstream wind turbine in the lower left plot in Figure 4.15, it appears that at small spacings the holistic approaches increase total power by increasing the relative power of the downstream turbine. At larger spacings, the greedy and holistic curves start to overlap. As the spacing increases to values considerably larger than 20R, it can be speculated that both curves would converge to a value of 50% as the wake interactions between the turbine would disappear. A similar result can be observed in the wind farm inputs over spacing in

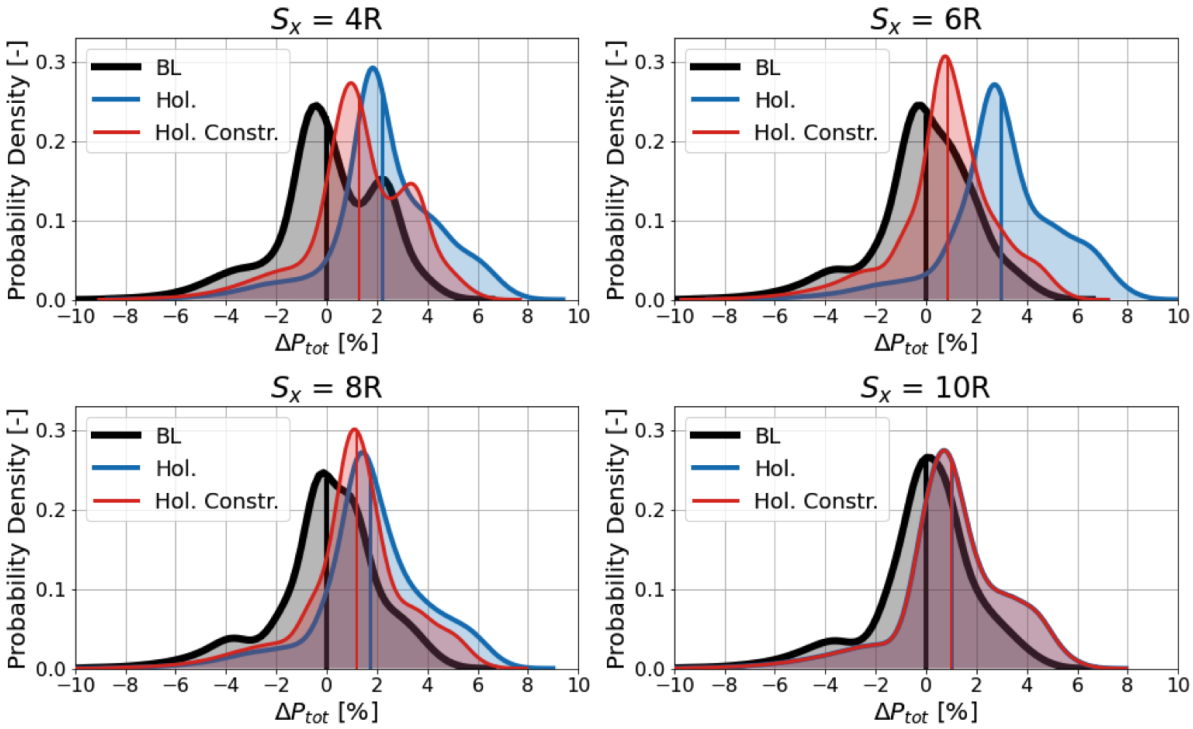

Figure 4.20: Distribution of total power for the baseline (black), unconstrained case (blue) and constrained case (red) for the spacings $S_x = 4R$, $6R$, $8R$, and $10R$.

the right column of Figure 4.15. At larger spacings, both the upstream and downstream yaw inputs of the holistic approach converge to the values of the greedy approach, thus showing again that the holistic wind farm control strategies are only relevant at small spacings for this scenario.

4.2.4.3 Wind Farm Loads

The analysis so far did not discuss the loads corresponding to the obtained optimized inputs. Figure 4.21 shows the median loads for the upstream (top) and downstream turbine (bottom) for the unconstrained (left) and constrained (right) cases. The constraint are indicated by the light brown shaded area. Note that the loads have not been corrected by the surrogate model output. As a result, the effects of the constrains can be clearly visualized. It can be seen that the limiting loads are the blade root flapwise bending moment $M_{BRF,1}$ for the upstream wind turbine and the main shaft torque $M_{MS,2}$ for the downstream wind turbine. The loads that are subject to the constraint also correspond to the channels with the highest surrogate model variance, as can seen in the Table 4.4. This is an unfortunate result, as it could mean that in some cases the constraint is applied to a load which in reality lower than the constraint, or vice versa. This is discussed in Appendix J, where the median loads corrected by the surrogate model error are shown. It is observed that the constraint is actually applied to severely on the upstream wind turbine blade root flapwise bending moment $M_{BRF,1}$, whilst being applied too loosely

for downstream wind turbine main shaft torque $M_{MS,2}$ at times. A possible way to circumvent this issue would be to include the surrogate model in the optimization process, as explained in Section 4.2.3.

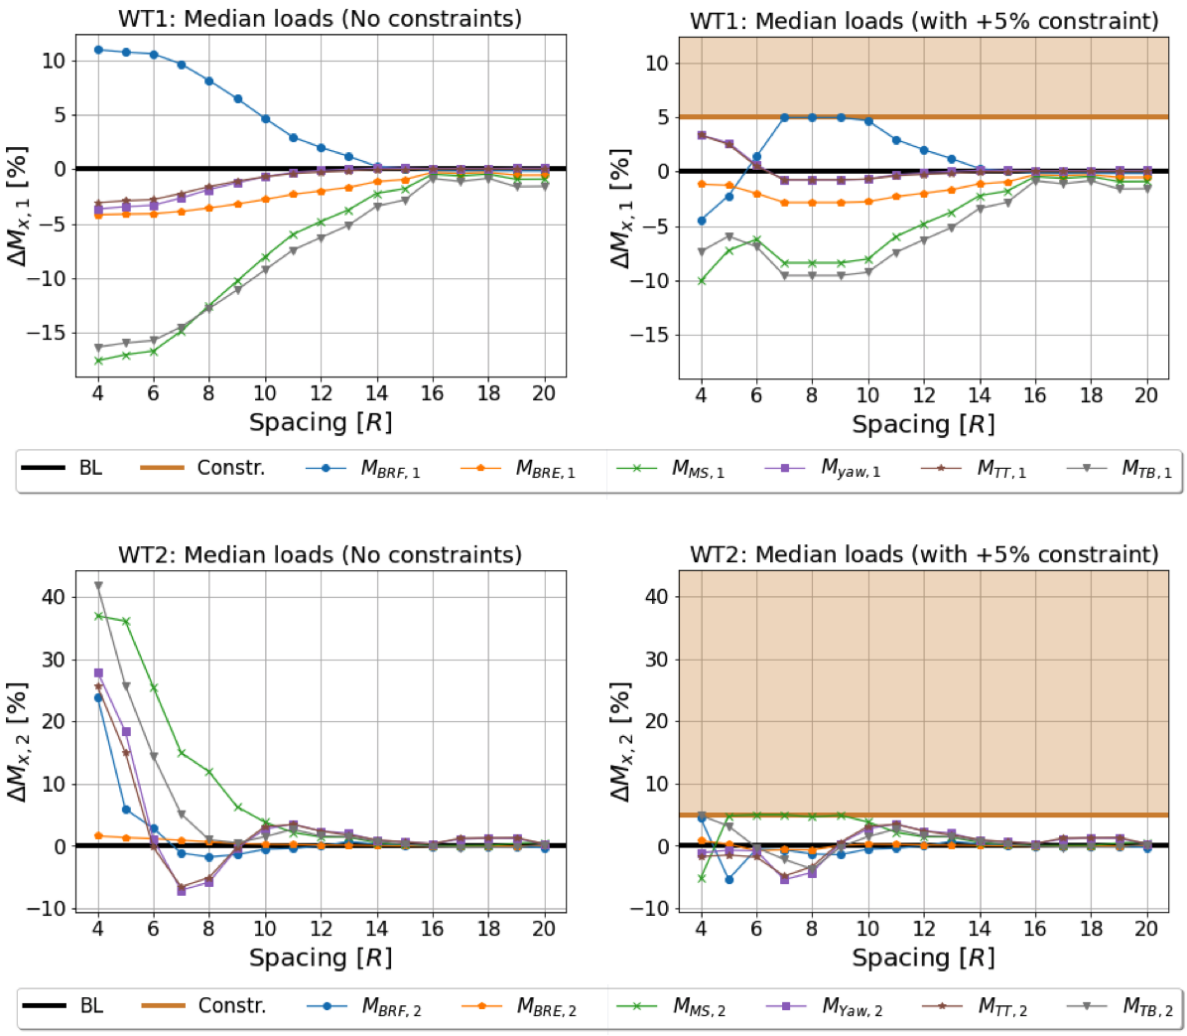

Figure 4.21: Median loading of upstream wind turbine (upper row) and downstream wind turbine (lower row) for the unconstrained (left) and constrained (right) cases. The constraint is indicated by the shaded area. Note that the shown outputs are not corrected by the surrogate model error such that the effect of the constraints is clearly visualized.

Upstream Wind Turbine

The loading of the unconstrained case of the upstream wind turbine in the upper left plot in Figure 4.21. Note that, as described in Section 4.2.1, the load changes are normalized for each distance individually, meaning that the larger loads differences are not a *direct* result of the closer spacing. Rather, largest loads with respect to the baseline correspond to the largest wind farm control settings differences between the the baseline and optimization case.

A decrease in loading with respect to the baseline is observed in all cases, except for the blade root flapwise bending moment $M_{BRF,1}$. A similar observation is made for the V27 wind turbine in Ennis et al. (2018), where the flapwise moment is observed to increase for negative yaw angles. A possible explanation for this increase is that due to the presence of shear, the blades will experience a change of angle of attack as they rotate. The amplitude of the change in angle of attack increases with the magnitude of the yaw-misalignment (Damiani et al., 2018). The periodical change in angle of attack causes a periodical change in blade loading, causing the increase flapwise bending moment. Due the increased size of the oscillations, the short-term equivalent is increased. Other moments are decreased due to the larger yaw angle, as the aerodynamic loads are reduced due the misalignment with the flow direction.

As the blade root flapwise bending moment is proportional with the magnitude of the yaw angle and the other loads are inversely proportional with the magnitude of the yaw angle, a trade-off has to be made when considering the power optimization with load constraints. Decreasing the yaw angle will reduce the upstream turbine bald root flapwise bending moment, however it will results in an increase of all the other upstream turbine loads and a decrease in the overall power production of the wind farm. This is observed in the constrained case of the upstream wind turbine in the plot in top right corner of Figure 4.21. It can be observed that the blade root flapwise bending moment $M_{BRF,1}$ is reduced as it reaches the constraint from $S_x = 7R$ to $S_x = 9R$. However, as discussed in Appendix J, the constraint applied to the $M_{BRF,1}$ channel is too severe, as the whole curve is shifted downwards when the surrogate model correction is applied. Considering the other loads, all of them experience an increase in loading due to the constraint. Most certainly the tower top yaw moment $M_{yaw,1}$, which is nearing the constraint at $S_x = 4R$.

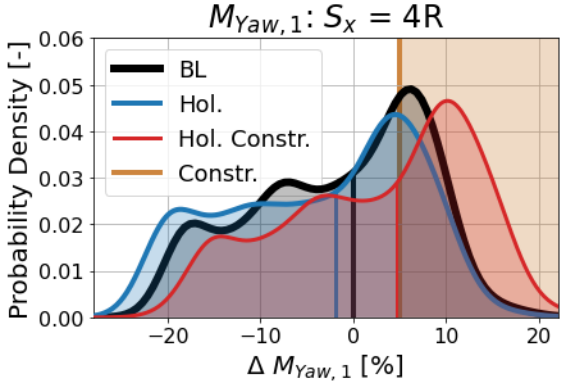

Figure 4.22: Distribution of upstream tower tower top yaw moment at $S_x = 4R$ for the different optimization cases.

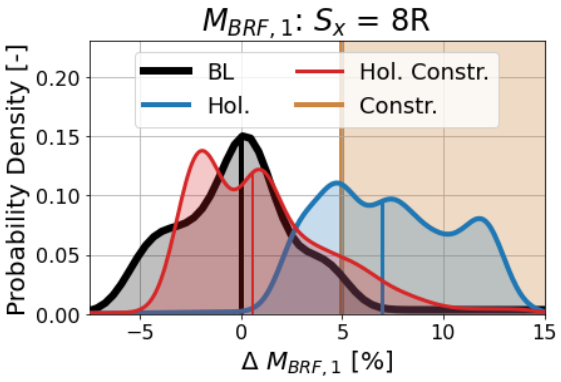

Figure 4.23: Distribution of blade root flapwise bending moment at $S_x = 8R$ for the different optimization cases.

A closer look is taken at the distribution of the $M_{yaw,1}$ channel at $S_x = 4R$ in Figure 4.22. The distributions of the baseline, unconstrained and constrained cases are shown. It can be observed that the three distributions appears to have the same overall shape, however

they are shifted with respect to one another. The increased tower top yaw moment $M_{yaw,1}$ of the constrained case is a result of the wind turbine decreasing its yaw offset with the flow direction, therefore increasing its aerodynamic loading. The unconstrained case is yawed more negatively, resulting in an increase wind turbine misalignment with the flow. As a result, its general aerodynamic loading is reduced. The used method the limit the wind turbine loading appears to not be the optimal method, as it causes a general undesired increase in all the other loads. In order to make the constraint more effective, two improvements could be made to the model. Firstly, individual constraint should be set for each loading, as some load might be more critical than other for a given wind turbine. Secondly, if more calibration data would be available, the surrogate model error should be added to the optimization process. To illustrate the latter point, the distribution of the blade root flapwise bending moment $M_{BRF,1}$ is shown for $S_x = 8R$ in Figure 4.23. This load is observed to be the constrained load at $S_x = 8R$ in Figure 4.21. However, as the surrogate model error associated with the inputs is negative, the whole curve is shifted downwards. As a result, the blade root flapwise bending moment $M_{BRF,1}$ has actually been over constrained by the optimizer. A similar observation is made in Appendix J with the median loading. Implementing the surrogate model error within the optimization process could potentially resolve this issue.

Downstream Wind Turbine

The downstream wind turbine experiences an increase for all measured loads for the unconstrained case, as can be seen in the lower left plot in Figure 4.21. Once again, the largest load increases correspond to the largest wind farm inputs with respect to the baseline, as shown in Figure 4.15. The increased loading of the downstream wind turbine originates from an increased aerodynamic loading due to the turbine yawing itself into the flow direction, as discussed previously by means of Figure 4.17. When the load constrained is enforced, as shown in lower right plot in Figure 4.21, the wind turbine reduces its loading by yawing away from the flow direction. Thus, whilst the downstream turbine yaws positively into the flow to increase its power production for the unconstrained case, the downstream turbine yaws negatively away from the flow in order to decrease the experienced loading.

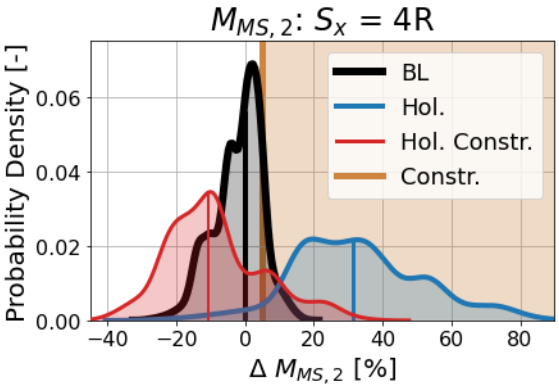

Figure 4.24: Distribution of downstream main shaft torque at $S_x = 4R$ for the different optimization cases.

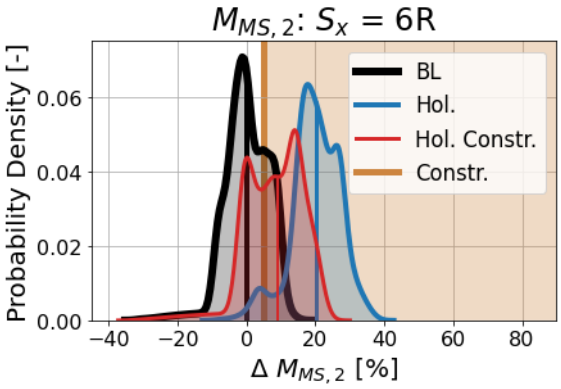

Figure 4.25: Distribution of downstream main shaft torque at $S_x = 6R$ for the different optimization cases.

Some loads, such as the main shaft torque $M_{MS,2}$, experience a drastic increase for the unconstrained case. The corresponding distributions is shown in Figure 4.24 for $S_x = 4R$. Besides having one of the highest median load increases, it also has a heavy tail ranging up to values larger than +80%. This heavy tail disappears as soon as the spacing is increase a little, as seen in for $S_x = 6R$ in Figure 4.25. The spacing of $S_x = 6R$ also corresponds to the location where the constraint is not respected by the constrained case, as a result of the surrogate model correction. It appears that the surrogate model error is positive and large enough to push the resulting load over the +5% constrain. The following section discusses the effects of putting the constraint on extreme loads, such that the long right tails of the main shaft torque $M_{MS,2}$ seen in Figure 4.24 can be limited.

Extreme Loads

The current optimization constraint have been put on the *median* value of the fatigue loads. However, as observed in Figure 4.24, some loads might have heavy right tails. Thus, it might be to the wind farm controller's best interest to limit the extreme loads rather than the median loads. The fixed load constraint of +5% has been set on the 50^{th} (Q50 or median), 75^{th} (Q75) and 95^{th} (Q95) percentiles for the following analysis. The resulting wind farm settings are shown in Figure 4.26. The unconstrained and the median (Q50) constrained cases is similar the cases analysed previously and is thus indicated using the same color code, namely blue and red respectively. The constrained extreme load cases Q75 and Q95 are indicated in orange and green respectively.

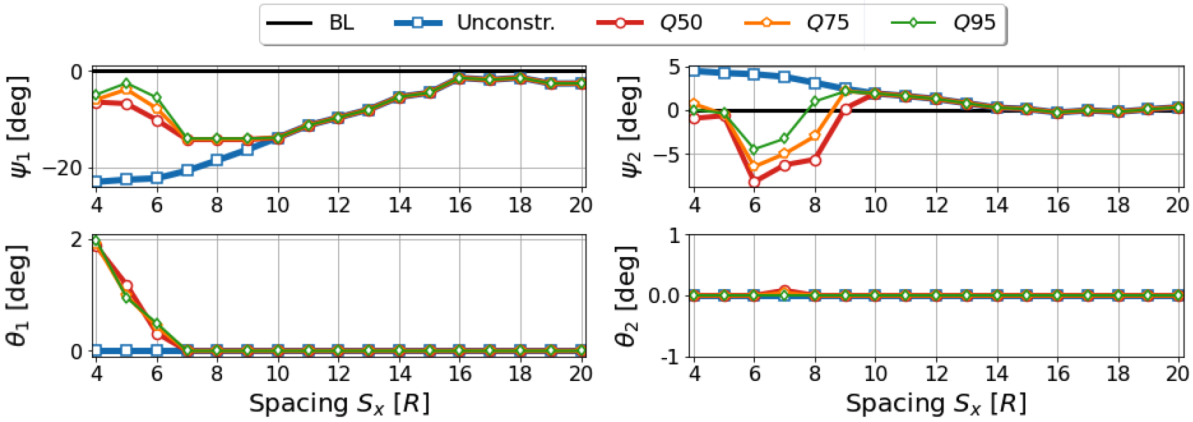

Figure 4.26: Wind farm settings with respect to the baseline for the unconstrained case (blue) and +5% constrained cases. The constraints are placed on Q50, Q75, and Q90 for the red, orange and green curves respectively.

The main differences appear to be for the yaw settings. For the upstream wind turbine, the largest differences with the previously analysed median constrained case (red) occur at spacings smaller than $S_x = 7R$. Ideally, as shown by the unconstrained case (blue), yaw offsets of more than 15° are beneficial for the total power production of the wind farm at those spacings. It thus appears that the Q95 case (green) is the most affected by the constraint for $S_x < 7R$, as its yaw offset is the smallest compared to the other cases. This means that the heavy tailed loads of the upstream wind turbine mostly occur at those spacings.

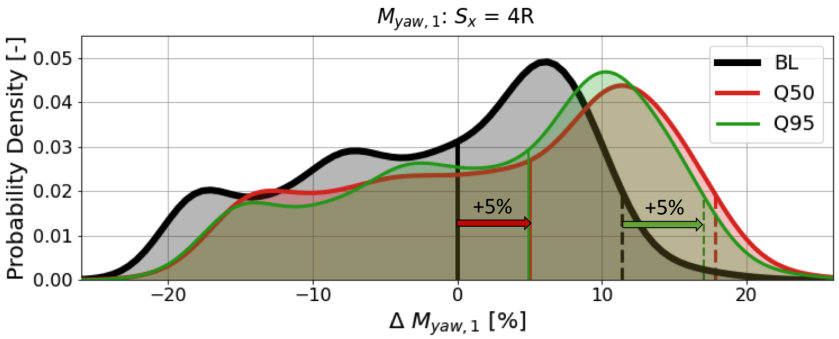

Figure 4.27: Distribution of upstream tower top yaw moment at $S_x = 4R$ for the the baseline (black), Q50 constrained case (red) and Q95 constrained case (green). The solid and dotted vertical lines show the median and 95^{th} percentile respectively.

The distribution of the upstream tower top yaw at a spacing of $S_x = 4R$ can be plotted once again as done previously in Figure 4.22, excluding the unconstrained case and including the Q95 constrained case. The resulting distributions are shown in Figure 4.27, with the baseline (black), Q50 (red), and Q95 (green) cases. The solid and dotted vertical lines show the median and 95^{th} percentile respectively. It can be observed that the Q50

case has a heavier tail than the Q95 case, as it is only bounded by its median value with respect to the baseline. However, the Q95 case is limited by the 95^{th} of the baseline, which results in a distribution with a slightly shorter tail. This shorter tail is obtained by decreasing the yaw angle of the upstream wind turbine ψ_1. Due to the smaller yaw angle of the Q95 case, one would expect a smaller power production of the wind farm. This is indeed seen in Figure 4.28, where it can be observed that the Q95 case experiences a smaller power gain compared with the other constrained cases.

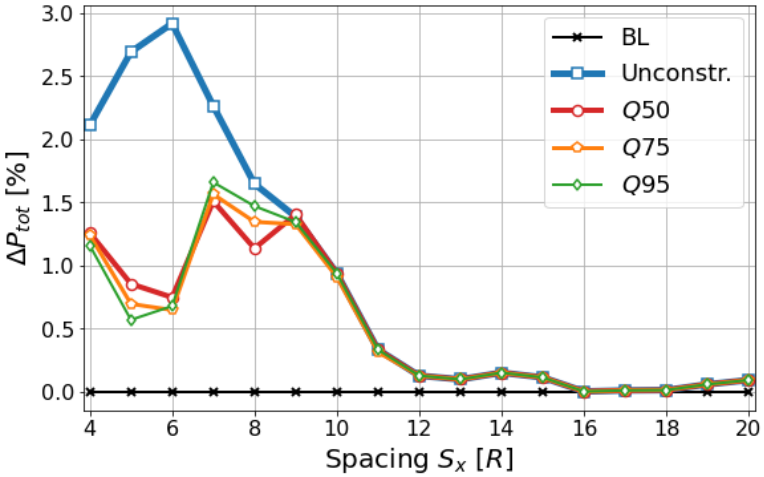

Figure 4.28: Median total power gains with respect to the baseline for the unconstrained case (blue) and +5% constrained cases. The constraints are placed on Q50, Q75, and Q90 for the red, orange and green curves respectively.

For a larger spacing ($S_x = 7R, 8R$), it appears that the Q95 case has a larger power production than the other constrained cases. Referring back to Figure 4.26, it can be observed that this is due to the fact that the Q95 case is able to yaw the wind turbine more into the flow than the other constrained cases, therefore increasing the power production of the downstream wind turbine. Thus, the opposite reasoning than for the upstream wind turbine applies here. It appears that yawing the downstream wind turbine into the flow will results in larger median loads, but in less heavy tails. As a result, the the Q95 case can be yawed more into the flow direction, and therefore increasing the wind farm power production.

As concluded in Section 4.2.4.1, it appears that down-regulation is only relevant at close spacings. As can be seen in Figure 4.26, the difference in down-regulation input for the Q50, Q75 and 95 cases are hard to differentiate, implying that the driving mechanisms is once again yaw-control. Furthermore, even when extreme loads are considered, only a negligible amount of down-regulation input is given to the down-stream wind turbine.

CHAPTER 5

Conclusions and Recommendations

The performed study focused on analyzing the potential of yaw-based wake-steering and constant blade pitch control down-regulation by means of statistical analyses of the wind turbine outputs. Data-driven polynomial-based surrogate models, calibrated through data obtained from a range of LES and aeroelastic simulations, have been used to optimize the wind farm operations in terms of power production and loads.

The holistic wind farm control approaches have been observed to return power gains with respect to the baseline at spacings smaller than $12R$. The largest power gains occurred at spacings ranging from $5R$ to $8R$, where median power gains of up to nearly $+3\%$ have been measured for the unconstrained case. For the constrained cases, median power gains ranging from $+1\%$ to $+2\%$ were measured, depending on the severity of the imposed constraint. More severe constraint result in smaller power gains and affect a larger range of spacings. For spacings larger than $12R$, no notable power increases have been observed. All cases, including the baseline case, showed wind farm power distributions with heavier tails on the left side, indicating that an exceptionally small power production with respect to the median value is more likely than an exceptionally large power production. The presented results are very similar to the ones obtained in Hulsman et al. (2019), where PCE-based surrogate models have been used to evaluate the potential of yaw-control within the same two wind turbine wind farm. Similar optimal yaw inputs and power gains are observed over the range of spacings.

Yaw-based wake-steering showed to be the driving mechanisms to increase the wind farm power production, particularly when loads are not considered. The general tendency is to yaw the upstream wind turbine negatively to deflect the flow and yaw the downstream wind turbine positively such that it aligns itself with the local flow direction. A different approach is observed when load constraints are considered. The upstream wind turbine can not yaw as much as it does for the unconstrained case, as the blade root flapwise bending moment is limited by the constraint. As a result, the flow-misalignment of the upstream wind turbine is reduced, which causes an increase all the other loads as its power capture is increased. As a result, the wind turbine has to down-regulate itself in order to limit the increases in loading. This means that wind turbine combines down-regulation with yaw-control such that the it can maximize the amplitude of the yaw angle whilst balancing the increase blade root flapwise loading with all the other loads. The combined wind farm control strategies are thus especially relevant at close spacings when loads are considered. For the downstream wind turbine, the predominant

approach to find the right balance between loading and power production appeared to be yaw-control, as no down-regulation input has been observed. The main shaft torque is subject to the constraint and is limited by introducing an offset with the flow direction.

There are some important considerations to be made when interpreting these results. Firstly, the loads of the upstream and downstream turbine that have been subject to the constraints, namely the flapwise moment and main shaft torque respectively, also correspond to the load channels with the largest spread of surrogate model errors. As the surrogate model error is not included within the optimization process, it is possible that in some cases the constrained is enforced on loads that should not be limited, and vice versa. In order to mitigate this issue, the surrogate model error could be added to the optimization process, such that the constrained it directly applied to the corrected value. This could potentially be done by constructing an additional surrogate model for the surrogate model error, therefore allowing it to be included within the optimization process.

Secondly, the same load constraint has been used to limit the output each load channel. As a consequence, as was observed in the case of the upstream wind turbine for the flapwise blade moment, the constraining of one particular load could cause an increase in all other loads. This issue could be avoided by placing individual constraints on each load channel, based upon the known load envelop of a specific wind turbine model. Furthermore, as mentioned in the first point, including the surrogate model error within the optimization process would allow to constrain the different percentiles more accurately.

Thirdly, the uncertainty associated with the *ghost* wind turbine approach has to be taken into account. The ghost turbine approach consists in extracting the dynamic inflow at given cross sections of the LES and using those as dynamic inflow for the Flex5 aeroelastic simulations. However, although Flex5 takes the blade deflection into account, the blades do no actually deflect within the incoming flow, as the inflow is determined beforehand. Furthermore, Flex5 takes the axial induction into account within its code, but the induction has no influence on the incoming flow. The outputs of the ghost wind turbine approach are compared with the directly coupled LES/actuator line outputs, as done in Appendix C. The power channel of the Flex5 code appears to be on the conservative side, whilst the fatigue loads are overestimated for the most part. However, even the LES/actuator line coupling is, albeit a very good one, only a representation of reality. Therefore, these quantitative differences offer an approximation of the uncertainty associated the constructed models.

Finally, the potential use of down-regulation might have been underestimated due to its implementation method. Blade pitch controllers have a relatively high refresh rates compared to yaw-control, as the inertial forces are much smaller. Thus, it is unrealistic to assume that a blade pitch controller would fix the blade pitch angle for a period of 10 minutes. Other approaches, such as delta control, might render the down-regulation approach more effective.

The following points discuss some of the limitations of the current study case and gives recommendations on how the focal point of this research could be extended:

- **Inflow conditions:** The current study has been performed using single below-rated wind flow condition, similar to the one used in Hulsman et al. (2019). The wind farm control strategies should be investigated under different environmental conditions by varying the wind direction, wind speed, shear exponent and turbulence intensity using different turbulence seeds.
- **Wind turbine model:** The V27-225kW wind turbine has been used the LES actuator line model, which is a relatively small and stiff wind turbine compared with modern-day wind turbines. As a result, the obtained results can not be completely extrapolated to new wind turbines and wind farms, as different physics are at play for larger wind turbines due to the different viscous effects, differences in structural stiffness and larger wind speed variations over the rotor area. However, by assuming Reynolds number similarity, fluid scaling laws allow to extrapolate the energy losses due to wake effects to larger wind turbines van der Laan et al. (2020).
- **Dynamic wind farm control strategies:** The presented work was restricted to two static control strategies. Dynamic control strategies, such as the sinusoidal thrust control strategy or helix approach, described in Munters and Meyers (2018c) and Frederik et al. (2020) respectively, focus on increasing the wind turbine wake turbulence such that wake recovery is accelerated. The possibility of using data-driven surrogate models to evaluate dynamic wind farm control strategies should be investigated.
- **Window size statistics:** The study has been conducted by means of the conventionally accepted 10-minute statistics. Due to the limited LES simulation time, a rolling window with 95% has been used to create the 23 10-minute samples, which as a result are not statistically independent. Smaller windows would allow to reduce the overlap and match the time windows used for grid operations of for example Denmark and Germany.

The bullet points below show an overview of recommendations and possible future works for the construction, performance enhancements and uses of the surrogate models:

- **Surrogate model error:** The surrogate model error, which is the error between the empirical CDF and the surrogate model generated CDF, has been quantified in Section 4.1.3. This error has been used to correct the outputs of the wind farm optimization by means of a look up table. However, this resulted in an improper applications of the load constraints, as described in Section 4.2.4.3. An alternative approach would be to construct a separate surrogate model directly returns the surrogate model error associated with a set of wind farm control parameters. Under the form of a surrogate model, the surrogate error could be estimated *during* the optimization process, rather than *after*.

- **Sample distribution fit:** The distributions of the 23 10-minute samples for each set of wind farm control settings have been assumed to be a simple Weibull distribution. It was observed that the Weibull distribution is to stiff to allow for an optimization based upon a range of percentiles rather than only the median value. A potential solution could be to include the surrogate model error during the optimization process, or to use a more complex distribution during the fitting process (e.g. double Weibull). Furthermore, the used of a PCE-based surrogate model brings its own limitations, which are discussed in the next point.
- **Polynomial chaos expansion:** PCE based surrogate models are very powerful tools, however their applications are limited to low-dimensional data, i.e. a limited amount of features. As discussed in Section 3.3, the PCE-based approach suffers from the curse of dimensionality. Increasing the number of features by adding wind turbines or environmental parameters (e.g. wind speed, turbulence intensity) would drastically increase the required computational power. Alternative types of surrogate models, such as artificial neural networks or others, should be considered for more complex systems.
- **LES database:** Extending the LES database would allow for better performing surrogate models and would allow to increase the sample density within the current domain, or even to extend the domain. For example, the *greedy* control approach could not be implemented as defined in Section 4.2.1 due to the loss of data of the positively yaw LES.

APPENDIX A
Combined control strategies LES

The possible combinations are determined through the following procedure. Two LES scenarios with pure yaw-control are selected, where the first case ($case_{1,ref}$) has a larger yaw-angle than the second case ($case_2$). $Case_2$ is combined with an increasing blade pitch input θ_1 until the resulting thrust is similar to the time-averaged thrust of $case_{1,ref}$. The corresponding thrust of the scenarios are determined with Flex5 by means of the ghost turbine approach. This procedure is repeated for all pairs of LES with pure yaw-control. The approach is then repeated by inverting the roles of yaw-control and down-regulation. Twenty combinations are obtained through this method.

Four scenarios are selected out of the twenty possible combinations, as shown in Figure A.1. These combinations were selected due to their alignment in terms of both thrust *and* power. *Combination 1* is aligned with the most down-regulated case in terms of power. *Combinations 2* and *3* are aligned with one another in terms of power. Finally, *Combination 4* is aligned with the second to most yawed case.

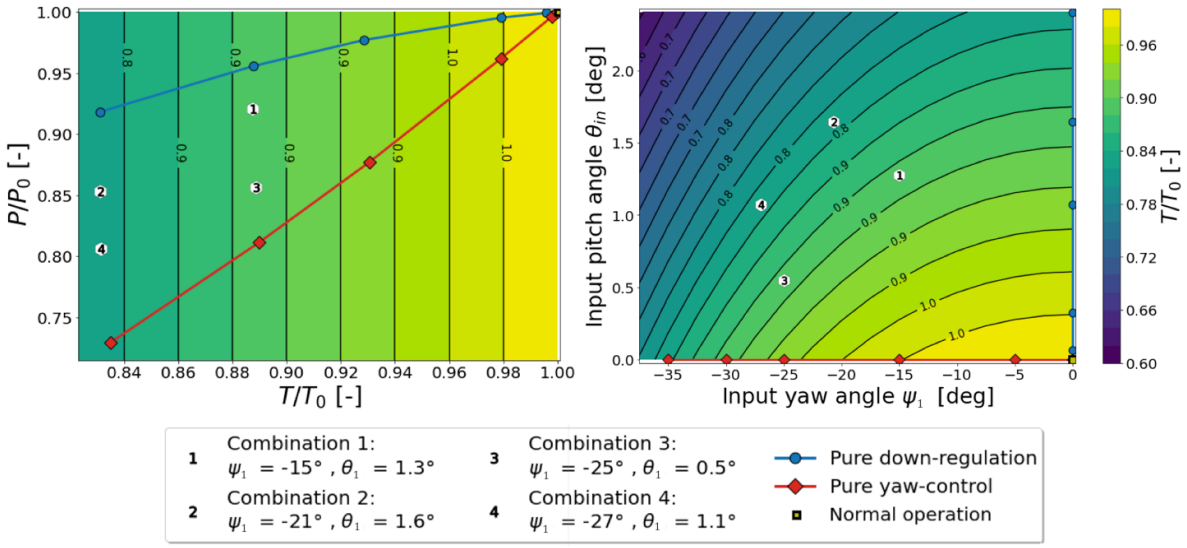

Figure A.1: Selected yaw-control and down-regulation scenarios. Left: Selected combined scenarios in terms of normalized power and thrust. Right: Selected combined scenarios in terms of the used yaw angle input ψ_1 and blade pitch input θ_1 settings.

APPENDIX B
Sweep analysis of V27

The behaviour of the V27 wind turbine is modelled for a range of wind speeds in no turbulence and zero gravity conditions. The inflow velocity is increased every two minutes in steps of starting from $4m/s$. Once the inflow velocity of $25m/s$ is reached, the inflow velocity is reduced by $1m/s$ every two minutes. The wind speed at hub height used over the simulation are shown in Figure B.1. The time-averages and their corresponding standard deviations are plotted in function of the hub wind speed.

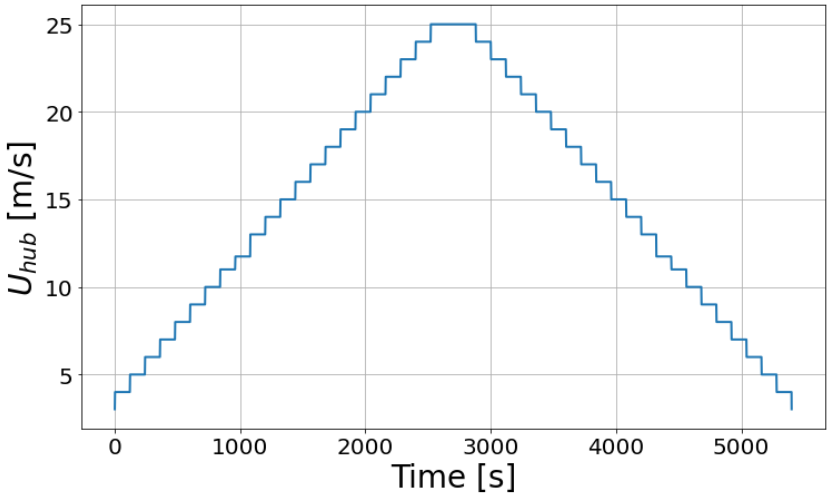

Figure B.1: Hub wind speed over sweep simulation.

The results of the sweep simulation for each selected channel, rotor rotational speed and blade pitch angle are shown in Figures B.2-B.11. The figures show hub wind speed against the the four minute time-averages and their corresponding standard deviations. The rated wind speed of the V27 is $U = 11.74m/s$. A sharp change is observed in the Figure B.5 at the rated wind speed, as the blade pitch controller is activated. This causes a drop in thrust, as can be seen in Figure B.3, as the blade pitch angle is moved away from its design point. Accordingly, as sharp decrease in the blade root moments is also observed in Figures B.6-B.7.

The main shaft torque shown in Figure B.8, is directly related to the produced electrical power, as can be observed by comparing it with figure B.2. The tower loads are shown in Figures B.9-B.11

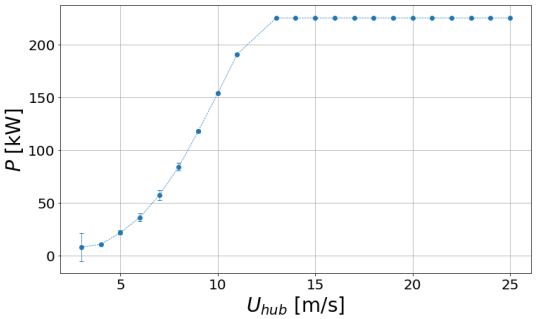

Figure B.2: Electrical power.

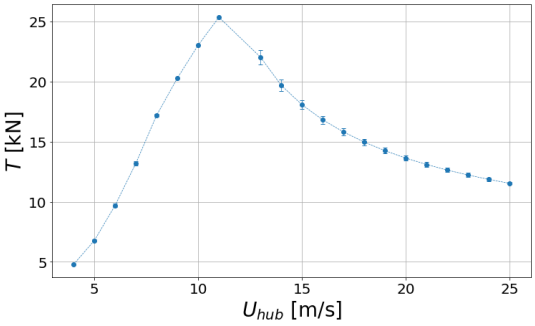

Figure B.3: Rotor thrust.

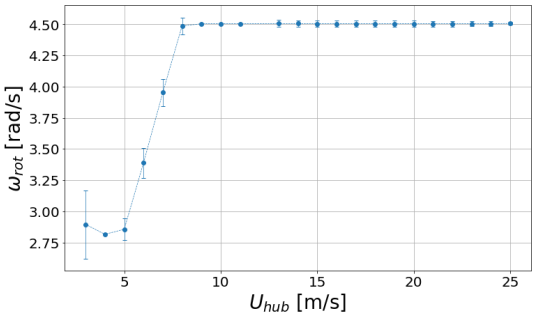

Figure B.4: Rotor rotational speed.

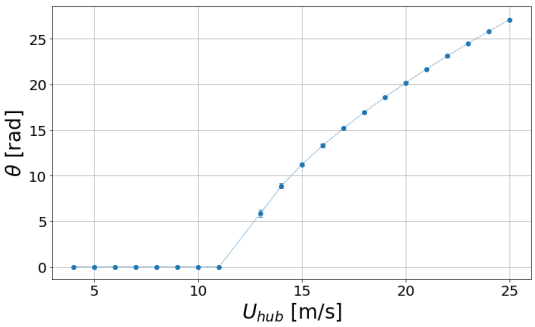

Figure B.5: Blade pitch angles.

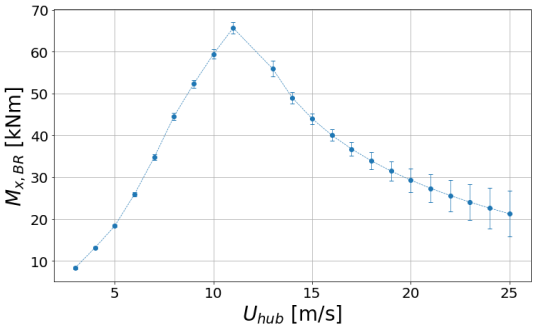

Figure B.6: Blade root flapwise bending moment.

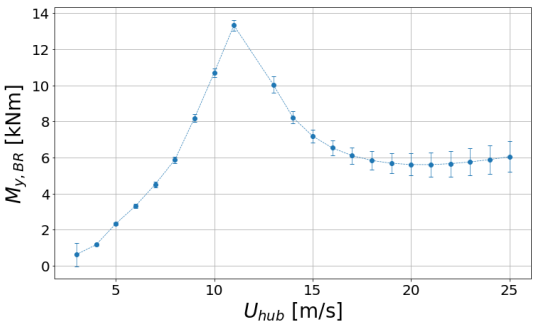

Figure B.7: Blade root edgewise bending moment.

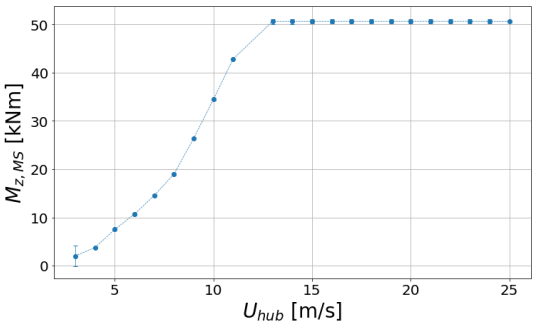

Figure B.8: Main shaft torque.

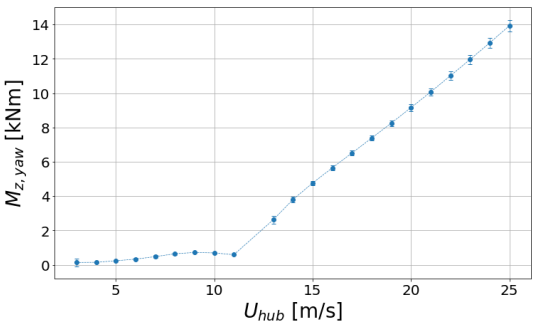

Figure B.9: Tower top torsion.

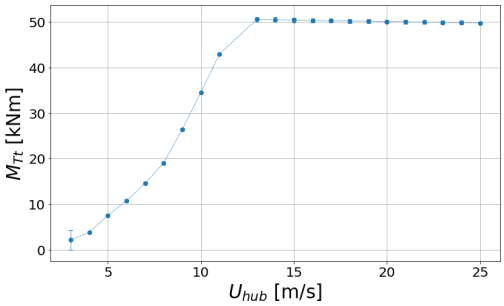

Figure B.10: Tower top bending moment.

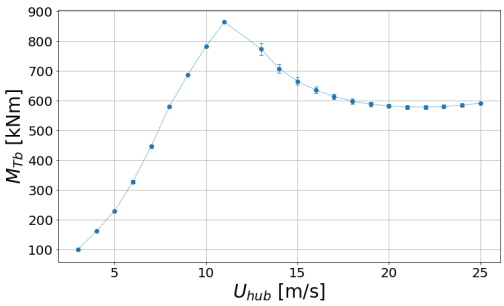

Figure B.11: Tower bottom bending moment.

APPENDIX C

Direct and indirect LES and Flex5 coupling comparison

A quantitative analysis of the differences between the "*ghost*" wind turbine approach (indirect coupling) and the LES wind turbine approach (direct coupling) is conducted in data generated through the process described in Section 3.2.1. The mean values for the power P channel or DEL for the other channels are determined for each 10-min window. This is done for both the data generated through the ghost wind turbine approach and the LES wind turbine approach. The error is then defined by the difference between the two methods, normalized by the median 10-min window value obtained for the *normal* LES, as shown in (C.1).

$$\epsilon[\%] = \frac{X_{ghost} - X_{LES}}{X_{LES,normal,median}} \cdot 100 \tag{C.1}$$

Three error distribution from the different channels are shown for three LES cases in Figure C.1. The selected LES cases are:

- Normal; no control input
- Down-regulation: $\theta_1 = 2.4^o$
- Yaw-control: $\psi_1 = -30^o$.

These LES cases were chosen such that the effect of the pure control strategy on the error can be visualised.

A closer look is taken at the effect of pure yaw-control and pure down-regulation on the error of the power P in Figures C.2-C.3. Here it can be seen that there appears to be a relation between the error and the magnitude of the thrust. Larger thrust values lead to larger errors. One could speculate that the error originates from a simplified thrust model.

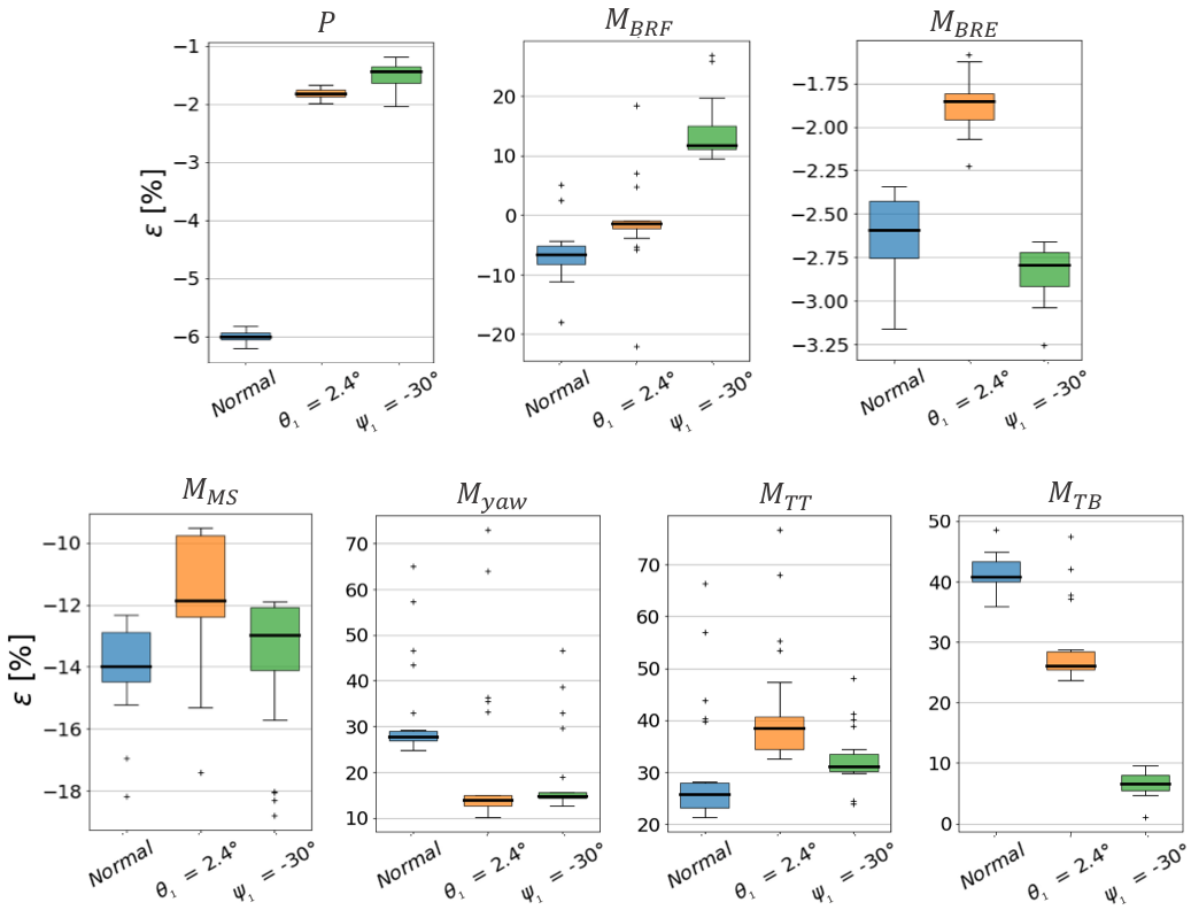

Figure C.1: Error distribution of different channels for *normal* LES ($\psi_1 = \theta_1 = 0°$), down-regulated LES ($\psi_1 = 0°, \theta_1 = 2.4°$), and yawed LES ($\psi_1 = -30°, \theta_1 = 0°$)

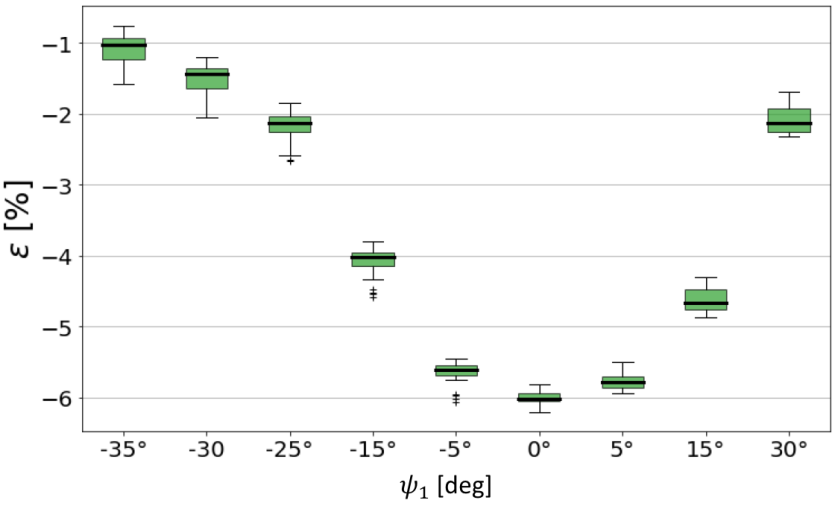

Figure C.2: Error distribution of power channel for pure yaw-control LES.

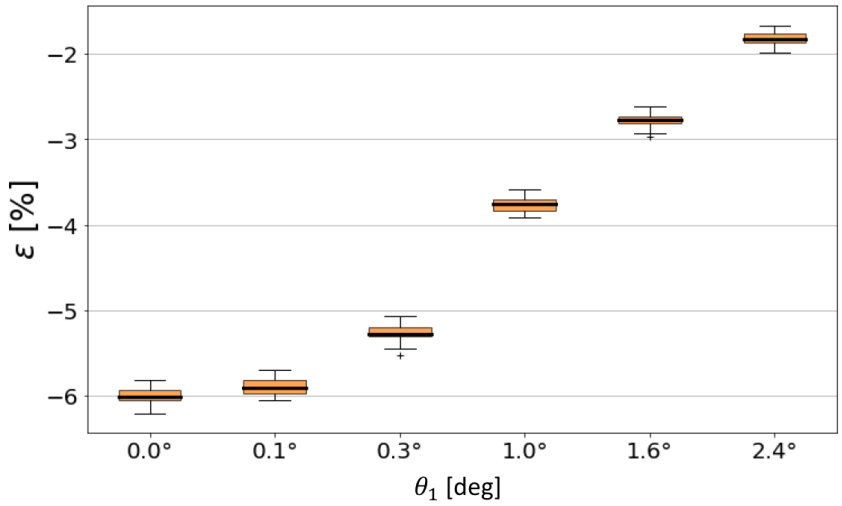

Figure C.3: Error distribution of power channel for pure down-regulation LES.

APPENDIX D
Channel frequency analysis.

A structural frequency analysis has been conducted for the load channels 2-7 listed in Table 3.6 for the V27 in a turbulent flow with no yaw or blade pitch input controls. The time series outputs of the channels are converted into the frequency domain. The resulting signals are smoothed out using the Welch method with a window of $10s$ and overlap of $5s$. The rotor rotational excitation frequencies xP and relevant structural eigenfrequencies of the Flex5 degrees of freedom are indicated in Figures D.1-D.6. This analysis could be continued be looking at the coupled eigenfrequencies in order to identify the remaining relevant peaks.

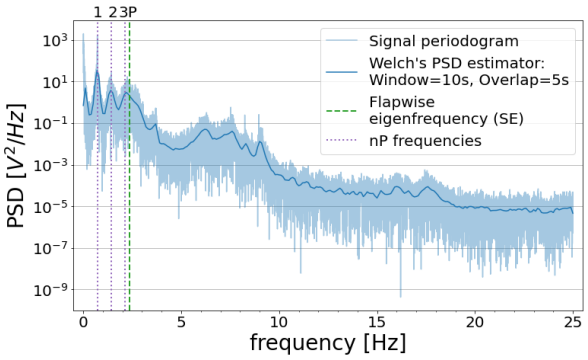

Figure D.1: Blade root flapwise bending moment.

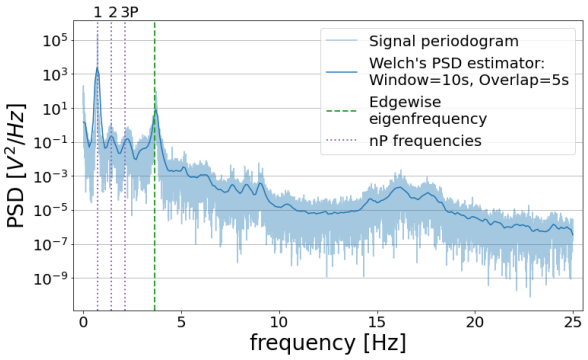

Figure D.2: Blade root edgewise bending moment.

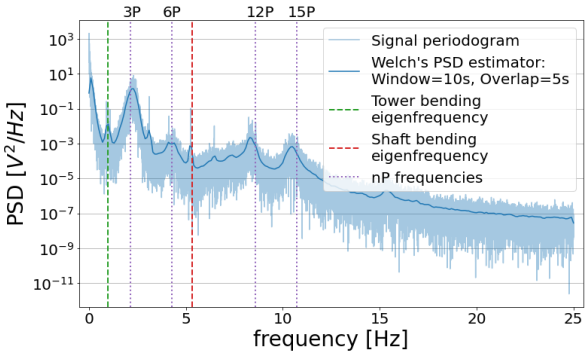

Figure D.3: Main shaft torsion.

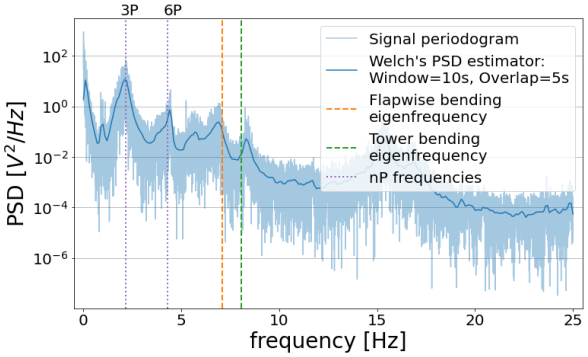

Figure D.4: Tower top torsion.

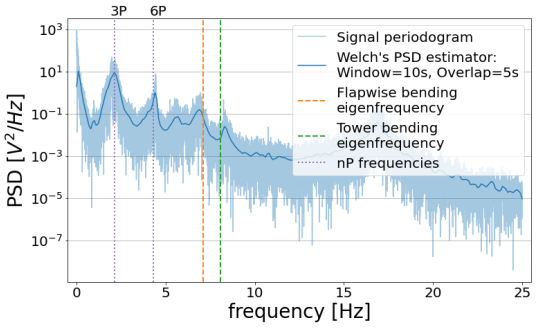

Figure D.5: Tower top bending moment.

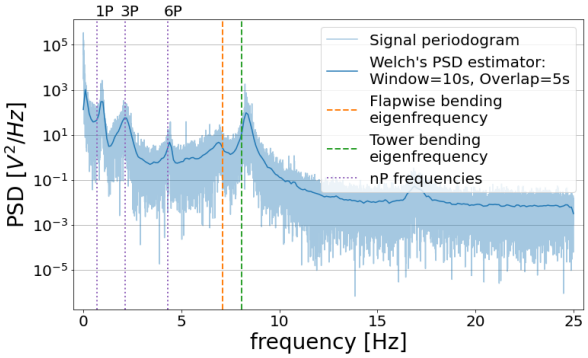

Figure D.6: Tower bottom bending moment.

APPENDIX E
Rotor Equivalent Wind Speed

The yaw-misalignment of the wind turbine is determined by means of the hub height wind velocity. However, this time signal is subject to local gusts that do not per se affect the whole rotor, but might by local turbulence at the hub. It is thus desired to filter out these gust at the hub that do not represent the general flow over the whole rotor.

A temporal low-pass filter is designed to remove gusts with a spatial scale that is smaller than the rotor itself. The low-pass filter is created by applying a centered rolling window over the time-signal with a given window size. The size of the window can be determined in two different ways.

E.1 REWS & Hub wind speed

The first approach consists in determining the Rotor Equivalent Wind Speed (REWS) by integrating the the local velocity at each control point on the rotor over the total rotor surface. This creates a spatial filter by averaging out the local gusts that happen at each control point over the time series, creating a smoother signal. The smoothing effect is shown in Figure E.1. The used REWS is determined by means of control points on 10 radial stations and 19 azimuthal stations, as shown in Figure E.2.

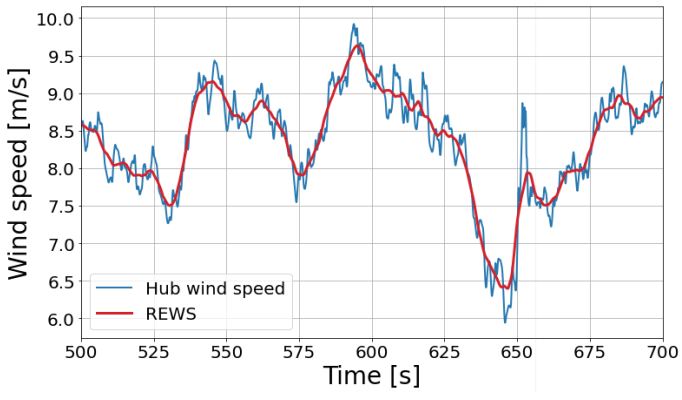

Figure E.1: Timeseries of hub wind speed (blue) and REWS (red).

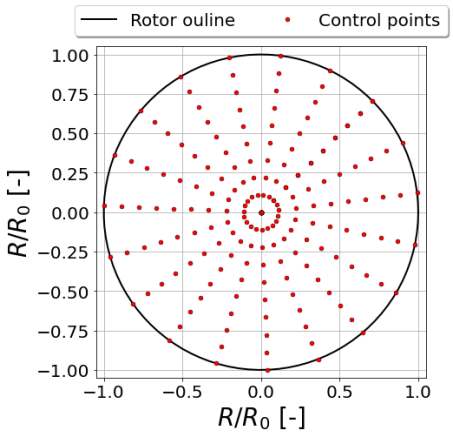

Figure E.2: Distribution of the rotor control points over the rotor surface area.

In order to determine the window-size that should be applied on the velocity at the

hub, the correlation coefficient of the REWS and filtered hub velocity is determined for a range of window-sizes. Figure E.3 shows the relation between the chosen window size for the hub velocity low-pass filter and the corresponding correlation coefficient. A maximum is observed for a window size of $6.53s$.

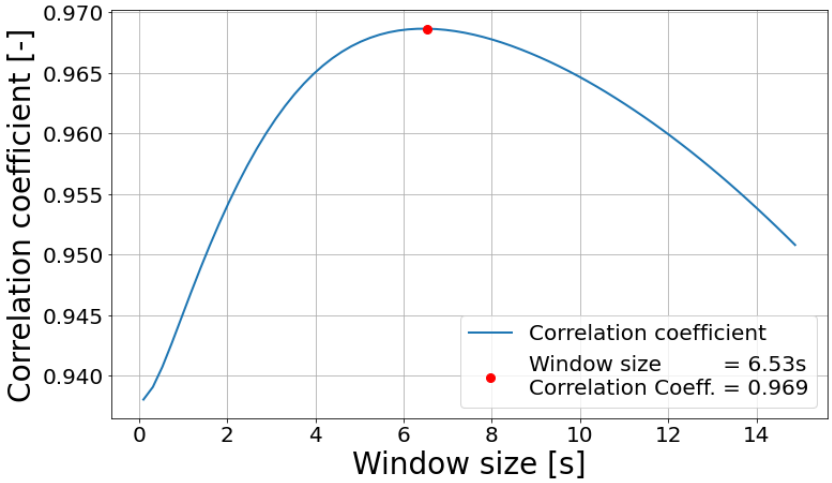

Figure E.3: Correlation coefficient of hub wind speed and REWS in function of rolling window size.

A parametric study is conducted on the effect of the number of radial and azimuthal control points and their effect on the obtained window size and correlation coefficient. As can be seen in Figure E.4, the window-size appears to converge for an increasing number of control points to the value of approximately $6.50s - 6.75s$, which is in the same range as the previously determined value.

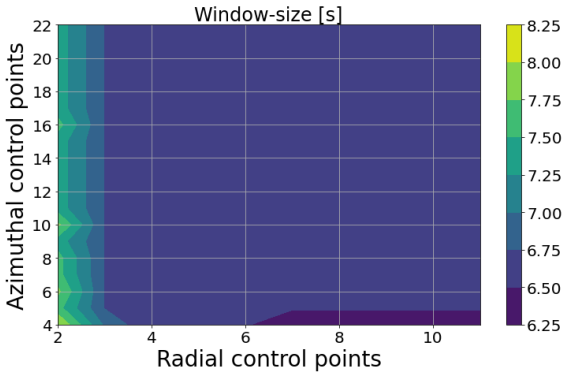

Figure E.4: Convergence of window size in function of radial and azimuthal control points.

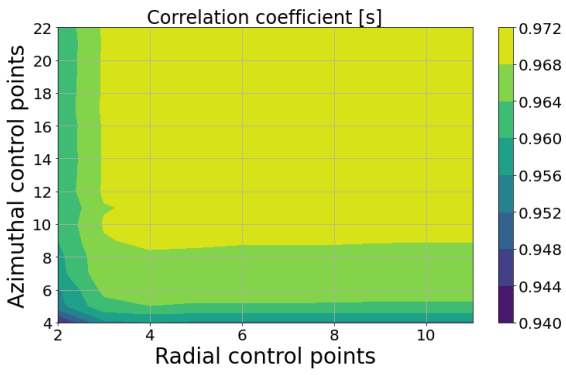

Figure E.5: Convergence of correlation coefficient in function of radial and azimuthal control points.

E.2 Power & Hub wind speed

A second method to determine the appropriate window size method consists in determining the correlation between the power and the filtered hub wind speed. The window size corresponding to the maximal correlation coefficient is the desired window size. The analysis can be conducted for both LES turbine or Ghost turbine, as shown in Figure E.6 and Figure E.7. As can be seen from the figures, this occurs at approximately 10 seconds. Thus, changes in the flow smaller than 10 seconds are irrelevant to the power output compared to larger fluctuations and can be filtered out using the corresponding window. This value is slightly larger than the one obtained from the previous approach.

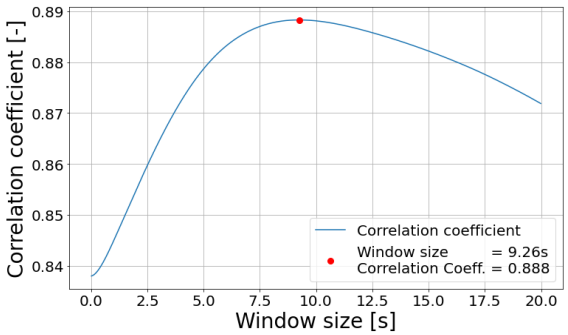

Figure E.6: Correlation coefficient in function of window size using the hub wind speed and power with the ghost turbine approach.

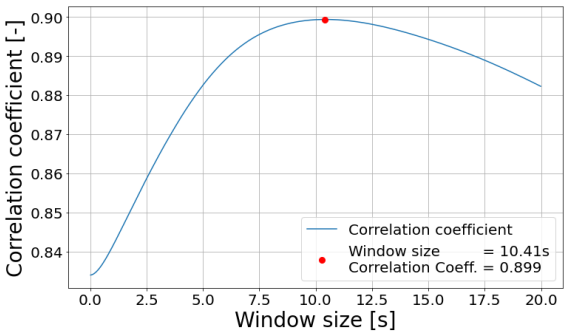

Figure E.7: Correlation coefficient in function of window size using the hub wind speed and power with the LES turbine approach.

APPENDIX F
Orthogonal Polynomial Families

The orthogonal polynomial families of most prevalent distribution are well-known and presented in Table F.1. Each polynomial family is orthogonal with respected to a certain distribution over a predefined interval. Hermite and Legendre polynomial families are discussed within the next section.

Table F.1: Overview of the classical orthogonal families and their corresponding distribution and interval.

Distribution	**Density function**	**Polynomial Family**	**Weight function**	**Interval**
Normal	$\frac{1}{\sqrt{2\pi}}e^{-x^2/2}$	Hermite $He_n(x)$	e^{-x^2}	$[+\infty, \infty]$
Uniform	$1/2$	Legendre $P_n(x)$	1	$[-1, 1]$
Beta	$\frac{(1-x)^\alpha(1+x)^\beta}{2^{\alpha+\beta+1}B(\alpha+1,\beta+1)}$	Jacobi $P_n^{\alpha,\beta}(x)$	$(1-x)^\alpha(1+x)^\beta$	$[-1, 1]$
Exponential	e^{-x}	Laguerre $L_n(x)$	e^{-x}	$[0, \infty]$
Gamma	$\frac{x^\alpha e^{-x}}{\Gamma(\alpha+1)}$	Generalized Laguerre $L_n^\alpha(x)$	$e^{-x}x^\alpha$	$[0, \infty]$

F.1 Hermite Polynomials

The Hermite polynomials $He_n(x)$ are a sequence of polynomials within a Hilbert space which are orthogonal with respect to the the Gaussian weight function on the interval $(-\infty, \infty)$,

$$\int_{-\infty}^{\infty} He_m(x)He_n(x)\phi(x)dx = n!\,\delta_{mn}, \tag{F.1}$$

where $\phi(x)$ is the standard normal distribution, shown in (F.2).

$$\phi(x) = \frac{e^{-x^2/2}}{\sqrt{2\pi}} \tag{F.2}$$

The Hermite polynomials are the solution to the differential equation,

$$y'' - xy' + ny = 0, \quad n \in \mathbb{N}. \tag{F.3}$$

Given the initial Hermite polynomial $He_0 = 1$, the sequence of higher order Hermite polynomials can be generated by means of their three-term recurrence relation,

$$He_{n+1}(x) = xHe_n(x) - nHe_{n-1}(x). \tag{F.4}$$

The first few Hermite polynomials are plotted in Figure F.1. The larger the maximal chosen order of the polynomial, the more their sum will resemble a Gaussian distribution.

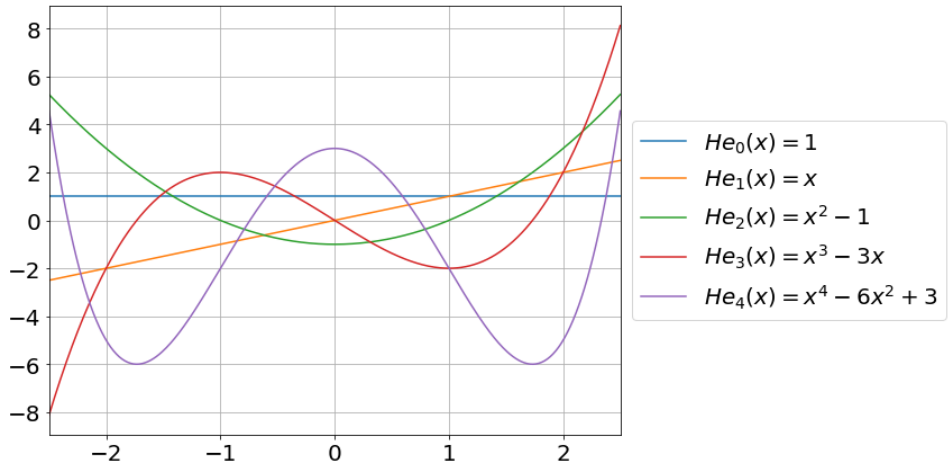

Figure F.1: Hermite polynomials of order ≤ 4.

F.2 Legendre Polynomials

In a similar way to the Hermite polynomials, the Legendre polynomials $P_n(x)$ are a sequence of polynomials within a Hilbert space which are orthogonal with respect to the the uniform weight function on the interval $(-1, 1)$,

$$\int_{-1}^{1} P_m(x)P_n(x)dx = \frac{2}{2n+1}\delta_{mn}. \tag{F.5}$$

The Legendre polynomials are the solution to the differential equation,

$$(1-x^2)y'' - 2xy' + n(n+1)y = 0, \quad n \in \mathbb{N}. \tag{F.6}$$

Given the initial Hermite polynomial $P_0 = 1$, the sequence of higher order Hermite polynomials can be generated by means of their three-term recurrence relation,

$$(n+1)P_{n+1}(x) = (2n+1)xP_n(x) - nP_{n-1}(x). \tag{F.7}$$

The first few Legendre polynomials are plotted in Figure F.2. Similarly to the Hermite polynomials for the Gaussian distribution, the Legendre polynomials will resemble a uniform distribution when their values are summed up. As the maximal order is increased to infinity, the summed up Legendre polynomials will converge to a uniform distribution.

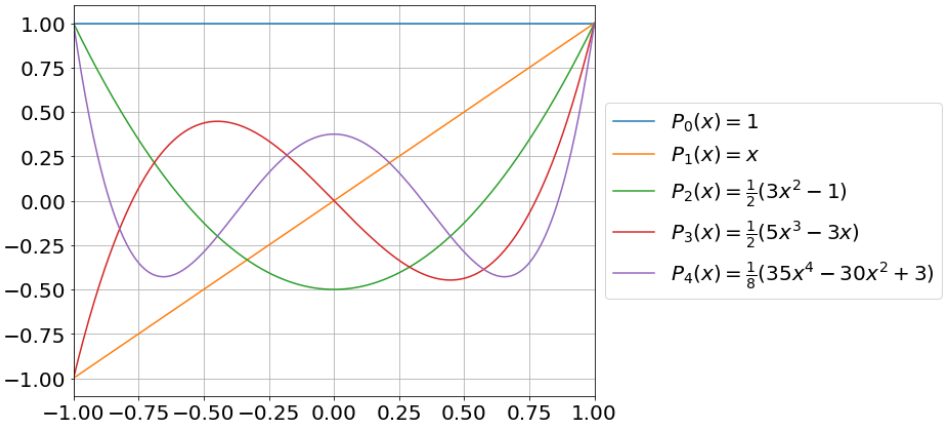

Figure F.2: Legendre polynomials of order ≤ 4.

APPENDIX G
Double Weibull parameters analysis

The shape and scale parameters for the double Weibull fit of the downstream wind turbine tower top yaw moment $M_{yaw,2}$ at $S_x = 5R$ and $\psi_1 = \theta_1 = 0°$ are shown in Figure G.1. Irregularly large values can be observed for the shape parameters (left column). The corresponding single and double Weibull distributions are shown for $\psi_2 = -10°$ and $\theta_2 = 0.9°$ in Figure G.2, where the second shape parameter is $k_2 = 47$. It can be observed that the large shape value corresponds with a vertical plateau in the CDF and too a high density concentration of data points in the PDF. Here, the power of the double Weibull is demonstrated as it fits the sample distribution very well. However, when the downstream yaw angle is changed from $\psi_2 = -10°$ to $\psi_2 = -5°$, a very different behaviour is observed. Figure G.1 shows that this correspond to a jump from $k_2 = 47$ to $k_2 = 5.5$. The resulting distribution is shown in Figure G.3. It can be observed that the double Weibull converged to a single Weibull. This unstable behaviour of the distributions parameters are very difficult to capture by the polynomial-based surrogate models.

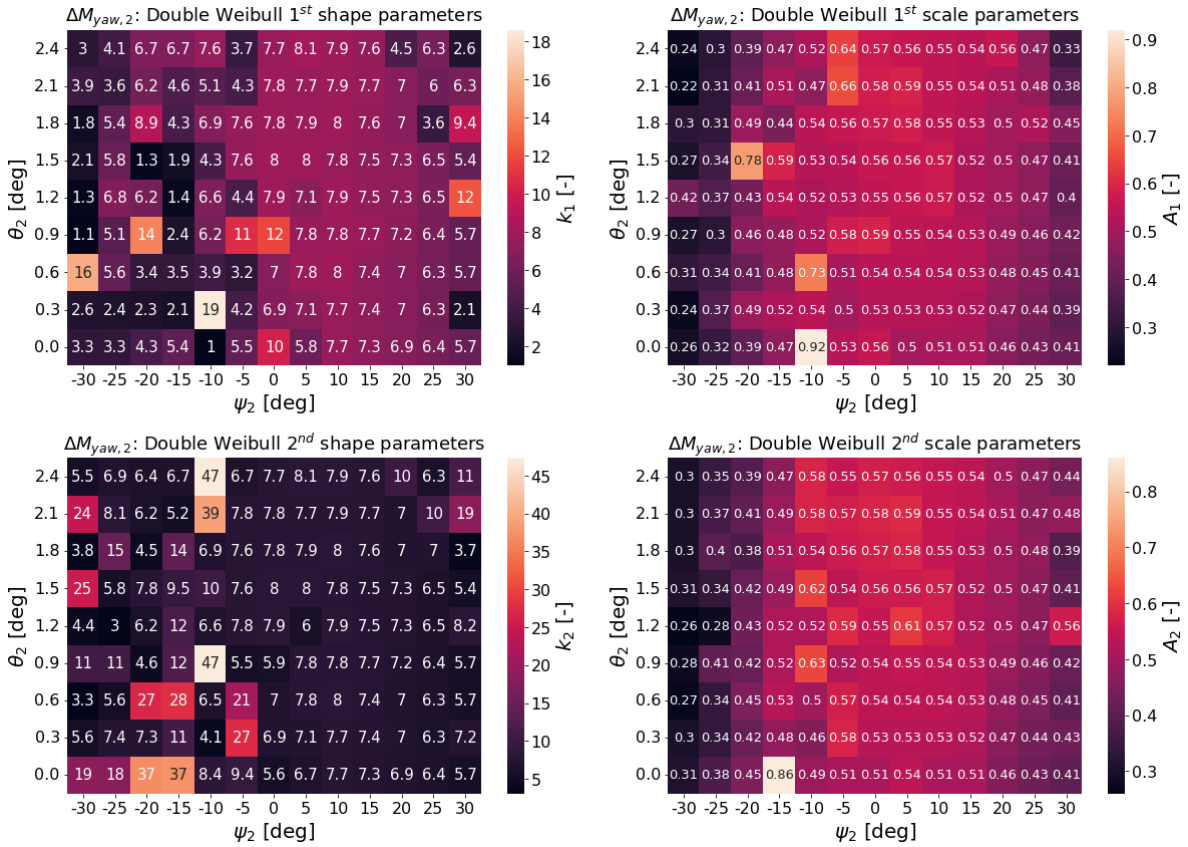

Figure G.1: Shape parameter (left) and scale parameter (right) of the first (upper row) and second (lower row) Weibull distribution of the double Weibull for $S_x = 5R$ and $\psi_1 = \theta_1 = 0°$ for the downstream wind turbine tower top yaw moment channel.

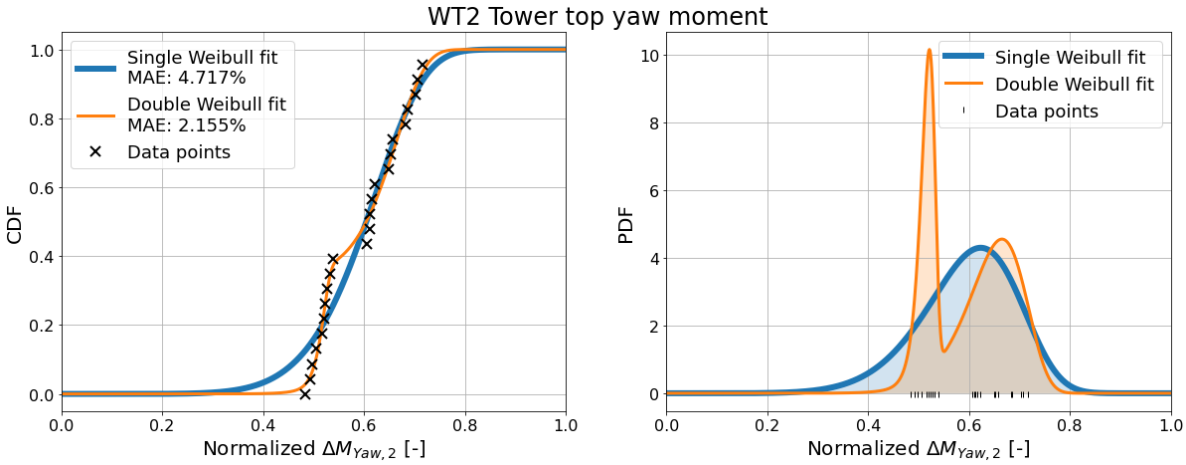

Figure G.2: Weibull curve with large second shape parameter $k_2 = 47$.

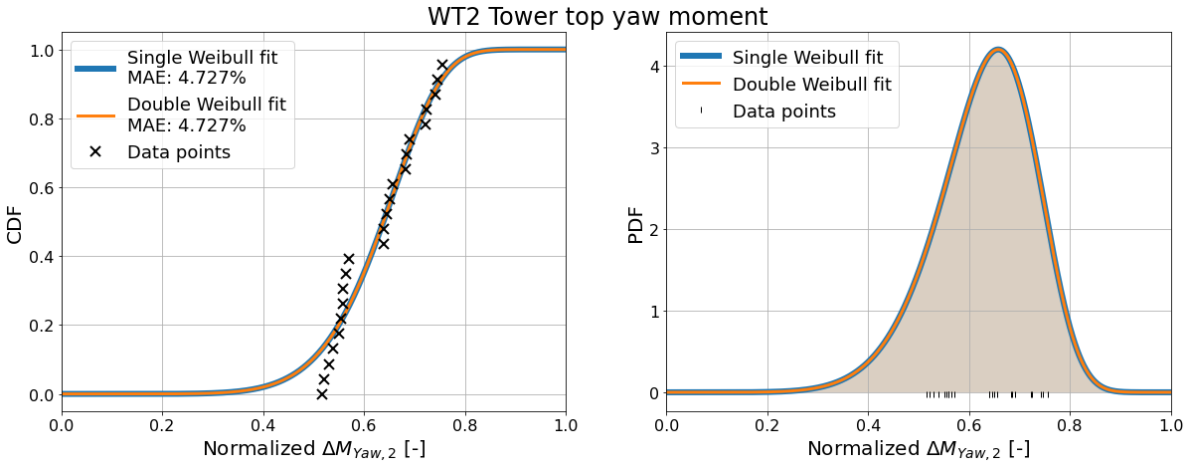

Figure G.3: Weibull curve with average sized second shape parameter $k_2 = 5.5$.

APPENDIX H

Surrogate model error of 4-fold cross-validation for P_2 channel

The surrogate model error quantifies the offset between of the surrogate model generated CDF and the empirical CDF. The error is determinde as explained in Section 4.1.1. In Section 4.1.3, it is observed that the test surrogate model error distribution appears to be bimodel, whilst the train error distribution only appears to show a single mode. The train and error distributions plotted in Figure 4.10 are an agglomeration of the 4-fold cross-validation errors. As a result, it could not be observed whether the bimodel nature of the test error was due to a particular combined LES scenario. Figure H.1 shows the train and test error for each fold separately for the power channel P_2 of the downstream turbine. It can be observed that the only test error distribution that does not show a bimodal behaviour is the upper left one, corresponding to yaw and blade pitch angles of $\psi_1 = -15°$ and $\theta_1 = -1.3°$. This scenario also corresponds to the smallest yaw angle for the combined LES scenarios. It can be speculated that due the the larger yaw angles of the other combined LES scenarios, the wake is deflected further away from the downstream wind turbine. As a result, the overlap between the rotor area and the wake is reduced. However, due to wake meandering this overlap will vary continuously. If this overlap is small due to the large yaw angle, its proportional variation will be larger than the one for the smaller yaw angles. This bimodel behaviour is only observed for the power channels, which have a relatively small error compared to the other channels.

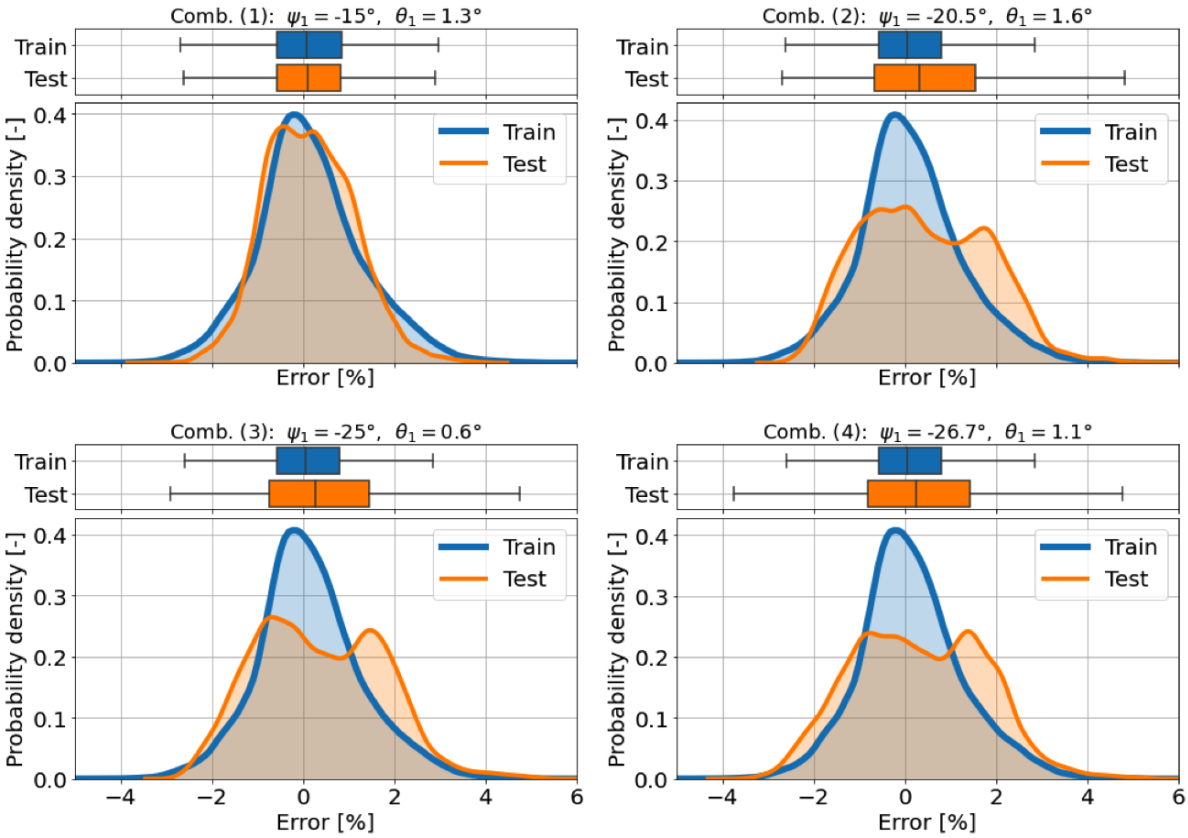

Figure H.1: The error distribution of the each fold of the 4-fold cross-validation train (blue) and test (orange) data for downstream turbine power channel P_2 surrogate models.

APPENDIX I

Upstream wind turbine greedy control cases comparison

The baseline scenario used to compare the different optimization results with is determined in Section 4.2.1. However, it was assumed that greedy input for the upstream wind turbine is $\psi_1 = 0°$, as this is the input wind direction used for the LES. However, due to the bias present in the wind direction, the upstream Wind Turbine Surrogate Model (WT SM), described in Section 3.2.4.2, returns that the true optimal input for the upstream wind turbine corresponds to $\psi_1 = +0.8°$. This input is not in the input domain of the Wind Farm Surrogate Model (WF SM) described in Section 3.2.4.3. As a result, the greedy input of the upstream wind turbine has to be assumed to be $\psi_1 = 0°$ in order to be able to continue the analysis. In order to justify this assumption, the difference in power and loads between the cases with $\psi_1 = 0°$ (Case 1) and $\psi_1 = +0.8°$ (Case 2) have to be determined. The error is measured by quantifying the difference between the two CDFs generated with the WT SM with the mentioned inputs, as done in in Section 4.1.1. The error is then normalized according to (I.1), where X is the chosen channel.

$$X_{norm} = \frac{X_{case2} - X_{case1}}{X_{case1}} \tag{I.1}$$

The mean and standard deviation of the errors for each channel are shown in Table I.1. It appears that the assumed greedy case (Case 1) has a slight tendency to estimate smaller channels outputs, for the exception of the tower top yaw ($M_{yaw,1}$) and bending ($M_{TT,1}$) moments. However, the magnitude of these errors are all smaller in magnitude than the surrogate model errors shown in Table 4.4, making these difference negligible. Furthermore, if these effects are considered negligible for the upstream wind turbine operations, it can be roughly assumed that the same can be said about the upstream turbine wake. As a result, it is assumed that the conditions for the downstream wind turbine between Case 1 and Case 2 do not differ by a relevant amount.

Table I.1: Mean and standard deviation of the errors for each channel between the *true* greedy and the *assumed* greedy control approaches.

Channel	$\boldsymbol{\Delta Case}$ [%]
P_1	0.011 ± 0.002
T_1	0.004 ± 0.001
$M_{BRF,1}$	0.161 ± 0.028
$M_{BRE,1}$	0.174 ± 0.001
$M_{MS,1}$	0.240 ± 0.02
$M_{yaw,1}$	-0.050 ± 0.027
$M_{TT,1}$	-0.010 ± 0.014
$M_{TB,1}$	0.490 ± 0.008

APPENDIX J

Median optimization results corrected with surrogate model error

The median values of the loads obtained from the optimization for the constrained and unconstrained cases with the addition of the surrogate model error are shown in Figure J.1. In the case of the downstream wind turbine (WT2), the main shaft torque $M_{MS,2}$ slightly passes the constraints at $S_x = 6R$ and $S_x = 13R$. This is a result from the addition of the surrogate model error after the optimization process. As the surrogate model error is not considered *during* the optimization process and is only added to the results with the look up table *after* the optimized wind turbine settings are obtained, some overshoot is possible. The median values of the loads obtained from the optimization without the addition of the surrogate model error can be found in Figure 4.21, where the effect of the constraints are clearly visible on the blade root flapwise bending moment $M_{BRF,1}$ of the upstream wind turbine and on the main shaft torque $M_{MS,2}$ of the downstream wind turbine. As a result, it can be observed that the $M_{BRF,1}$ channel has been constrained too severely when the surrogate model error is considered. For the $M_{MS,2}$ channel, as the associated surrogate model is relatively large, the load has been under- and over-constrained depending on the spacing.

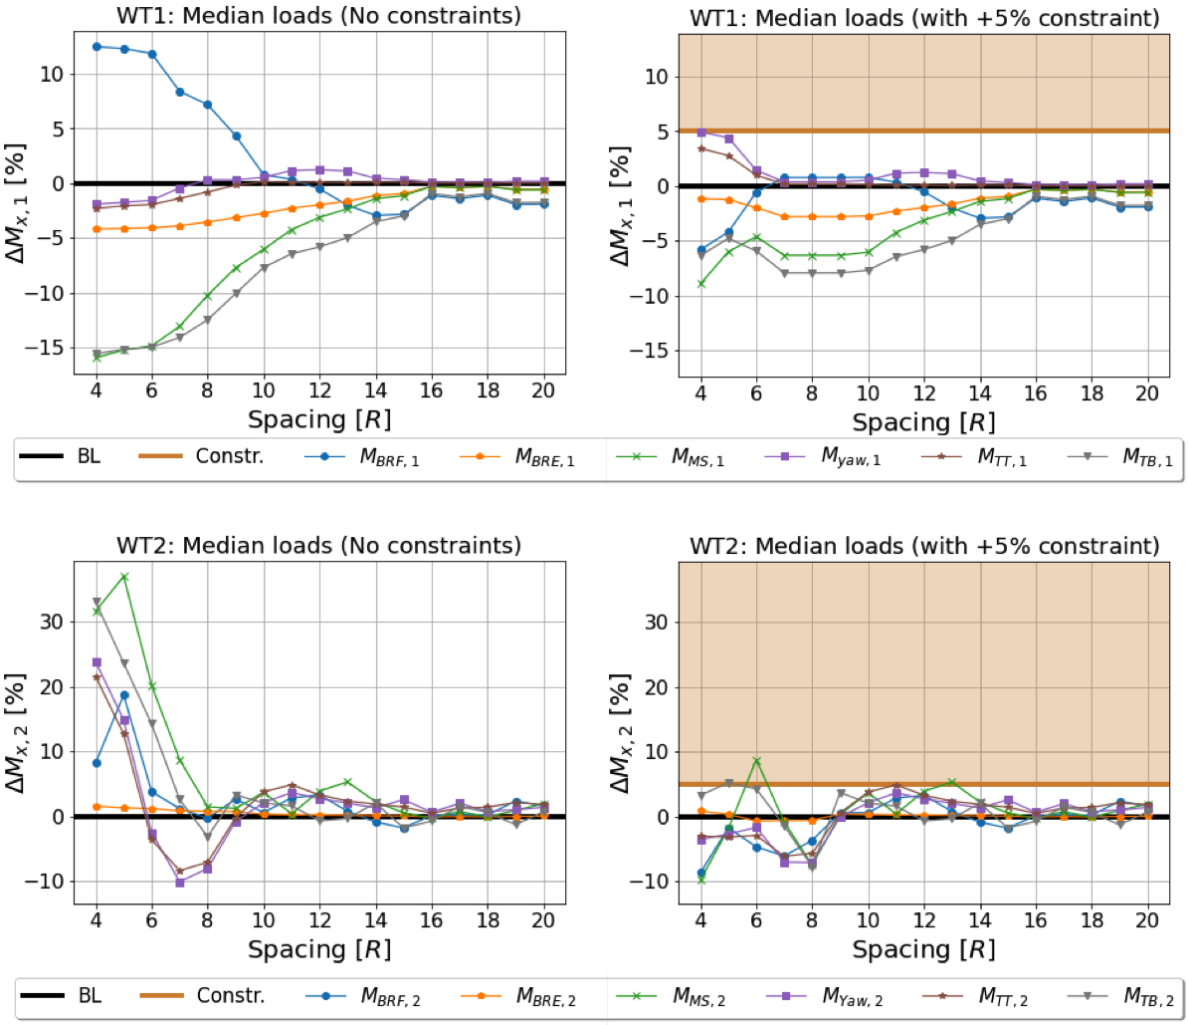

Figure J.1: Loading of upstream wind turbine (upper row) and downstream wind turbine (lower row) for the unconstrained (left) and constrained (right) cases with the surrogate model error correction. The constraint is indicated by the shaded area.

Bibliography

Adaramola, M. S., & Krogstad, P. Å. (2011). Experimental investigation of wake effects on wind turbine performance. *Renewable Energy*, *36*(8), 2078–2086. https://doi.org/10.1016/j.renene.2011.01.024

Ainslie, J. F. (1988). Calculating the flowfield in the wake of wind turbines. *Journal of Wind Engineering and Industrial Aerodynamics*, *27*(1-3), 213–224. https://doi.org/10.1016/0167-6105(88)90037-2

Allaerts, D., & Meyers, J. (2015). Large eddy simulation of a large wind-turbine array in a conventionally neutral atmospheric boundary layer. *Physics of Fluids*, *27*(6). https://doi.org/10.1063/1.4922339

Annoni, J., Bay, C., Taylor, T., Pao, L., Fleming, P., & Johnson, K. (2018). Efficient Optimization of Large Wind Farms for Real-Time Control, In *Proceedings of the american control conference*. https://doi.org/10.23919/ACC.2018.8430751

Annoni, J., Gebraad, P. M., Scholbrock, A. K., Fleming, P. A., & Wingerden, J. W. V. (2016). Analysis of axial-induction-based wind plant control using an engineering and a high-order wind plant model. *Wind Energy*, *19*(6), 1135–1150. https://doi.org/10.1002/we.1891

Annoni, J., Seiler, P., Johnson, K., Fleming, P., & Gebraad, P. (2014). Evaluating wake models for wind farm control, In *Proceedings of the american control conference*, Institute of Electrical; Electronics Engineers Inc. https://doi.org/10.1109/ACC.2014.6858970

Barthelmie, R. J., Frandsen, S. T., Nielsen, M. N., Pryor, S. C., Rethore, P. E., & Jørgensen, H. E. (2007). Modelling and measurements of power losses and turbulence intensity in wind turbine wakes at middelgrunden offshore wind farm. *Wind Energy*, *10*(6), 517–528. https://doi.org/10.1002/we.238

Bastankhah, M., & Porté-Agel, F. (2014). A new analytical model for wind-turbine wakes. *Renewable Energy*, *70*, 116–123. https://doi.org/10.1016/j.renene.2014.01.002

Bastankhah, M., & Porté-Agel, F. (2016). Experimental and theoretical study of wind turbine wakes in yawed conditions. *Journal of Fluid Mechanics*, *806*, 506–541. https://doi.org/10.1017/jfm.2016.595

Bastankhah, M., & Porté-Agel, F. (2019). Wind farm power optimization via yaw angle control: A wind tunnel study. *Journal of Renewable and Sustainable Energy*, *11*(2). https://doi.org/10.1063/1.5077038

Bay, C. J., King, J., Fleming, P., Mudafort, R., & Martínez-Tossas, L. A. (2019). Unlocking the Full Potential of Wake Steering: Implementation and Assessment of a Controls-Oriented Model. *Wind Energy Science Discussions*, (May), 1–20. https://doi.org/10.5194/wes-2019-19

Blatman, G., & Sudret, B. (2011). Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, *230*(6), 2345–2367. https://doi.org/10.1016/j.jcp.2010.12.021

Boersma, S., Doekemeijer, B. M., Gebraad, P. M., Fleming, P. A., Annoni, J., Scholbrock, A. K., Frederik, J. A., & Van Wingerden, J. W. (2017). A tutorial on control-oriented modeling and control of wind farms, In *Proceedings of the american control conference*, Institute of Electrical; Electronics Engineers Inc. https://doi.org/10.23919/ACC.2017.7962923

Boorsma, K. (2012). *Heat and flux. analysis of field measurements* (technical report). ECN.

Bossanyi, E. A. (2000). The Design of closed loop controllers for wind turbines. *Wind Energy*, *3*(3), 149–163. https://doi.org/10.1002/we.34

Boussinesq, J. (1877). *Essai sur la théorie des eaux courantes*. Paris: Imprimerie Nationale.

Burton, T., Jenkins, N., Sharpe, D., & Bossanyi, E. (2011). *Wind Energy Handbook*. J. Wiley. https://doi.org/10.1002/9781119992714

Churchfield, M. J., Fleming, P., Bulder, B., & White, S. M. (2015). Wind turbine wake-redirection control at the Fishermen's Atlantic City Windfarm, In *Proceedings of the annual offshore technology conference*. https://doi.org/10.4043/25644-ms

Churchfield, M. J., Lee, S., Michalakes, J., & Moriarty, P. J. (2012). A numerical study of the effects of atmospheric and wake turbulence on wind turbine dynamics. *Journal of Turbulence*, *13*, 1–32. https://doi.org/10.1080/14685248.2012.668191

Crespo, A., Hernandez, J., Fraga, E., & Andreu, C. (1988). Experimental validation of the UPM computer code to calculate wind turbine wakes and comparison with other models. *Journal of Wind Engineering and Industrial Aerodynamics*, *27*(1-3), 77–88. https://doi.org/10.1016/0167-6105(88)90025-6

Damiani, R., Dana, S., Annoni, J., Fleming, P., Roadman, J., van Dam, J., & Dykes, K. (2018). Assessment of wind turbine component loads under yaw-offset conditions. *Wind Energy Science*, *3*(1), 173–189. https://doi.org/10.5194/wes-3-173-2018

Dilip, D., & Porté-Agel, F. (2017). Wind turbine wake mitigation through blade pitch offset. *Energies*, *10*(6). https://doi.org/10.3390/en10060757

Dimitrov, N. (2019). Surrogate models for parameterized representation of wake-induced loads in wind farms. *Wind Energy*, *22*(10), 1371–1389. https://doi.org/10.1002/we.2362

Dimitrov, N., Kelly, M., Vignaroli, A., & Berg, J. (2018). From wind to loads: wind turbine site-specific load estimation using databases with high-fidelity load simulations. *Wind Energy Science Discussions*, 1–39. https://doi.org/10.5194/wes-2018-18

Eldred, M. S. (2009). Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design, In *Collection of technical papers - aiaa/asme/asce/ahs/asc structures, structural dynamics and materials conference*. https://doi.org/10.2514/6.2009-2274

Elorza, I., Calleja, C., & Pujana-Arrese, A. (2019). On wind turbine power delta control. *Energies*, *12*(12), 2344. https://doi.org/10.3390/en12122344

Emanuel, G. (2015). *Analytical fluid dynamics*. Taylor & Francis, CRC Press. https://books.google.dk/books?id=x2WvswEACAAJ

Ennis, B. L., White, J. R., & Paquette, J. A. (2018). Wind turbine blade load characterization under yaw offset at the SWiFT facility. *Journal of Physics: Conference Series*, *1037*(5). https://doi.org/10.1088/1742-6596/1037/5/052001

Feinberg, J., & Langtangen, H. P. (2015). Chaospy: An open source tool for designing methods of uncertainty quantification. *Journal of Computational Science*, *11*, 46–57. https://doi.org/10.1016/j.jocs.2015.08.008

Fleming, P. A., Gebraad, P. M., Lee, S., van Wingerden, J. W., Johnson, K., Churchfield, M., Michalakes, J., Spalart, P., & Moriarty, P. (2014). Evaluating techniques for redirecting turbine wakes using SOWFA. *Renewable Energy*, *70*, 211–218. https://doi.org/10.1016/j.renene.2014.02.015

Fleming, P. A., Gebraad, P. M., Lee, S., van Wingerden, J. W., Johnson, K., Churchfield, M., Michalakes, J., Spalart, P., & Moriarty, P. (2015). Simulation comparison of wake mitigation control strategies for a two-turbine case. *Wind Energy*, *18*(12), 2135–2143. https://doi.org/10.1002/we.1810

Fleming, P., Annoni, J., Churchfield, M., Martinez-Tossas, L. A., Gruchalla, K., Lawson, M., & Moriarty, P. (2018). A simulation study demonstrating the importance of large-scale trailing vortices in wake steering. *Wind Energy Science*, *3*(1), 243–255. https://doi.org/10.5194/wes-3-243-2018

Frederik, J. A., Doekemeijer, B. M., Mulders, S. P., & van Wingerden, J. W. (2020). The helix approach: Using dynamic individual pitch control to enhance wake mixing in wind farms. *Wind Energy*, *23*(8), arXiv 1912.10025, 1739–1751. https://doi.org/10.1002/we.2513

Gebraad, P. M., Teeuwisse, F. W., Van Wingerden, J. W., Fleming, P. A., Ruben, S. D., Marden, J. R., & Pao, L. Y. (2014). A data-driven model for wind plant power optimization by yaw control, In *Proceedings of the american control conference*, Institute of Electrical; Electronics Engineers Inc. https://doi.org/10.1109/ACC.2014.6859118

Gebraad, P. M., Teeuwisse, F. W., Van Wingerden, J. W., Fleming, P. A., Ruben, S. D., Marden, J. R., & Pao, L. Y. (2016). Wind plant power optimization through yaw control using a parametric model for wake effects - A CFD simulation study. *Wind Energy*, *19*(1), 95–114. https://doi.org/10.1002/we.1822

Gebraad, P. M., & Van Wingerden, J. W. (2014). A control-oriented dynamic model for wakes in wind plants, In *Journal of physics: Conference series*, Institute of Physics Publishing. https://doi.org/10.1088/1742-6596/524/1/012186

Ghanem, R. G., & Spanos, P. D. (1991). *Stochastic finite elements: A spectral approach*. Berlin, Heidelberg, Springer-Verlag.

Göçmen, T., Laan, P. V. D., Réthoré, P. E., Diaz, A. P., Larsen, G. C., & Ott, S. (2016). Wind turbine wake models developed at the technical university of Denmark: A review. Elsevier Ltd. https://doi.org/10.1016/j.rser.2016.01.113

Herp, J., Poulsen, U. V., & Greiner, M. (2015). Wind farm power optimization including flow variability. *Renewable Energy*, *81*, 173–181. https://doi.org/10.1016/j.renene.2015.03.034

Howland, M. F., Lele, S. K., & Dabiri, J. O. (2019). Wind farm power optimization through wake steering. *Proceedings of the National Academy of Sciences*, *116*(29), 14495–14500. https://doi.org/10.1073/pnas.1903680116

Hulsman, P. (2018). Power production and turbine response during yaw-based wind farm control - DTU Findit. Retrieved November 5, 2019, from https://findit.dtu.dk/en/catalog/2438297445

Hulsman, P., Andersen, S. J., & Göçmen, T. (2019). Optimizing Wind Farm Control through Wake Steering using Surrogate Models based on High Fidelity Simulations. *Wind Energy Science Discussions*, 1–34. https://doi.org/10.5194/wes-2019-46

International Electrotechnical Commission, I. (2005). *IEC 61400-1 third edition 2005-08 Wind turbines – Part 1: Design requirements* (technical report).

IRENA. (2019a). *Future of wind: Deployment, investment, technology, grid integration and socio-economic aspects (a global energy transformation paper)* (technical report). International Renewable Energy Agency. www.irena.org/publications/2019/Oct/Future-of-wind

IRENA. (2019b). *Renewable Capacity Statistics* (technical report). International Renewable Energy Agency. https://www.irena.org/publications/2019/Mar/Renewable-Capacity-Statistics-2019

Ivanell, S., Sørensen, J. N., Mikkelsen, R., & Henningson, D. (2007). Numerical analysis of the tip and root vortex position in the wake of a wind turbine, In *Journal of physics: Conference series*. https://doi.org/10.1088/1742-6596/75/1/012035

Japar, F., Mathew, S., Narayanaswamy, B., Lim, C. M., & Hazra, J. (2014). Estimating the wake losses in large wind farms: A machine learning approach, In *2014 ieee pes innovative smart grid technologies conference, isgt 2014*, IEEE Computer Society. https://doi.org/10.1109/ISGT.2014.6816427

Jensen, N. O. (1983). A note on wind generator interaction. *Risø-M-2411 Risø National Laboratory Roskilde*, arXiv arXiv:1011.1669v3, 1–16. https://doi.org/Riso-M-2411

Jiménez, Á., Crespo, A., & Migoya, E. (2010). Application of a LES technique to characterize the wake deflection of a wind turbine in yaw. *Wind Energy*, *13*(6), 559–572. https://doi.org/10.1002/we.380

Johnson, K. E., & Thomas, N. (2009). Wind farm control: addressing the aerodynamic interaction among wind turbines, In *Proceedings of the american control conference*. https://doi.org/10.1109/ACC.2009.5160152

Kaintura, A., Dhaene, T., & Spina, D. (2018). Review of polynomial chaos-based methods for uncertainty quantification in modern integrated circuits. MDPI AG. https://doi.org/10.3390/electronics7030030

Katic, I., Hojstrup, J., & Jensen, N. O. (1986). A simple model for cluster efficiency. *Citation*, *1*, 407–410.

Kheirabadi, A. C., & Nagamune, R. (2019). A quantitative review of wind farm control with the objective of wind farm power maximization. Elsevier B.V. https://doi.org/10.1016/j.jweia.2019.06.015

Knudsen, T., & Bak, T. (2012). Data driven modelling of the dynamic wake between two wind turbines, In *Ifac proceedings volumes (ifac-papersonline)*. https://doi.org/10.3182/20120711-3-BE-2027.00128

Knudsen, T., Bak, T., & Svenstrup, M. (2015). Survey of wind farm control - Power and fatigue optimization. *Wind Energy*, *18*(8), 1333–1351. https://doi.org/10.1002/we.1760

Kragh, K. A., & Hansen, M. H. (2014). Load alleviation of wind turbines by yaw misalignment. *Wind Energy*, *17*(7), 971–982. https://doi.org/10.1002/we.1612

Kristoffersen, J. (2005). The Horns Rev Wind Farm and the Operational Experience with the Wind Farm Main Controller. *Proceedings of the Copenhagen Offshore Wind*, (October), 1–9.

Leonard, B. P. (1979). A stable and accurate convective modelling procedure based on quadratic upstream interpolation. *Computer Methods in Applied Mechanics and Engineering*, *19*(1), 59–98. https://doi.org/10.1016/0045-7825(79)90034-3

Lesieur, M. (1996). *New Trends in Large-Eddy Simulations of Turbulence* (technical report Number 1). https://doi.org/10.1146/annurev.fluid.28.1.45

Lissaman, P. B. S. (1979). Energy effectiveness of arbitrary arrays of wind turbines. *Journal of Energy*, *3*(6), https://doi.org/10.2514/3.62441, 323–328. https://doi.org/10.2514/3.62441

Mann, J. (1994). The Spatial Structure of Neutral Atmospheric Surface-Layer Turbulence. *Journal of Fluid Mechanics*, *273*, 141–168. https://doi.org/10.1017/S0022112094001886

Mann, J. (1998). Wind field simulation. *Probabilistic Engineering Mechanics*, *13*(4), 269–282. https://doi.org/10.1016/s0266-8920(97)00036-2

Martínez-Tossas, L. A., Churchfield, M. J., & Leonardi, S. (2015). Large eddy simulations of the flow past wind turbines: Actuator line and disk modeling. *Wind Energy*, *18*(6), 1047–1060. https://doi.org/10.1002/we.1747

McTavish, S., Feszty, D., & Nitzsche, F. (2013). Evaluating Reynolds number effects in small-scale wind turbine experiments. *Journal of Wind Engineering and Industrial Aerodynamics*. https://doi.org/10.1016/j.jweia.2013.07.006

Michelsen, J. A. (1992). Basis3d-a platform for development of multiblock pde solvers.

Michelsen, J. (1994). *Block structured multigrid solution of 2d and 3d elliptic pde's*. Department of Fluid Mechanics, Technical University of Denmark. https://books.google.dk/books?id=dDu-XwAACAAJ

Munters, W., & Meyers, J. (2018a). Dynamic strategies for yaw and induction control of wind farms based on large-eddy simulation and optimization. *Energies*, *11*(1), 177. https://doi.org/10.3390/en11010177

Munters, W., & Meyers, J. (2018b). Optimal dynamic induction and yaw control of wind farms: Effects of turbine spacing and layout, In *Journal of physics: Conference series*, Institute of Physics Publishing. https://doi.org/10.1088/1742-6596/1037/3/032015

Munters, W., & Meyers, J. (2018c). Towards practical dynamic induction control of wind farms: analysis of optimally controlled wind-farm boundary layers and sinusoidal induction control of first-row turbines. *Wind Energy Science*, *3*(1), 409–425. https://doi.org/10.5194/wes-3-409-2018

Murcia, J. P., Réthoré, P. E., Dimitrov, N., Natarajan, A., Sørensen, J. D., Graf, P., & Kim, T. (2018). Uncertainty propagation through an aeroelastic wind

turbine model using polynomial surrogates. *Renewable Energy*, *119*, 910–922. https://doi.org/10.1016/j.renene.2017.07.070

NREL. (2019). FLORIS. Version 1.0.0. GitHub. https://github.com/NREL/floris

Oye, S. (1996). Flex4 simulation of wind turbine dynamics. proceedings of 28th iea meeting of experts concerning state of the art of aeroelastic codes for wind turbine calculations.

Pao, L. Y., & Johnson, K. E. (2009). A tutorial on the dynamics and control of wind turbines and wind farms, In *Proceedings of the american control conference*. https://doi.org/10.1109/ACC.2009.5160195

Peña, A., & Rathmann, O. (2014). Atmospheric stability-dependent infinite wind-farm models and the wake-decay coefficient. *Wind Energy*, *17*(8), 1269–1285. https://doi.org/10.1002/we.1632

Petersen, M. (1990). *General rights Wind turbine test Vestas V27-225 kW* (technical report). Risø-M, No. 2861.

Quick, J., Annoni, J., King, R., Dykes, K., Fleming, P., & Ning, A. (2017). Optimization under Uncertainty for Wake Steering Strategies, In *Journal of physics: Conference series*. https://doi.org/10.1088/1742-6596/854/1/012036

Resor, B. R., & LeBlanc, B. (2014). *An Aeroelastic Reference Model for the SWIFT Turbines* (technical report). Sandia National Laboratories (SNL). Albuquerque, NM, Livermore, CA (United States). https://doi.org/10.2172/1503792

Rychlik, I. (1987). A new definition of the rainflow cycle counting method. *International Journal of Fatigue*, *9*(2), 119–121. https://doi.org/10.1016/0142-1123(87)90054-5

S. Dhiman, H., & Deb, D. (2020). Fundamentals of Wind Turbine and Wind Farm Control Systems, In *Studies in systems, decision and control*. https://doi.org/10.1007/978-981-15-0275-0_1

Sanderse, B. (2009). *Aerodynamics of wind turbine wakes - literature review* (technical report). ECN.

Sanderse, B., Van Der Pijl, S. P., & Koren, B. (2011). Review of computational fluid dynamics for wind turbine wake aerodynamics. *Wind Energy*, *14*(7), 799–819. https://doi.org/10.1002/we.458

Schepers, J. G., & Van Der Pijl, S. P. (2007). Improved modelling of wake aerodynamics and assessment of new farm control strategies, In *Journal of physics: Conference series*, Institute of Physics Publishing. https://doi.org/10.1088/1742-6596/75/1/012039

Schmitt, F. G. (2007). About Boussinesq's turbulent viscosity hypothesis: historical remarks and a direct evaluation of its validity. *Comptes Rendus Mécanique*, *335*(9-10), 617–627. https://doi.org/10.1016/j.crme.2007.08.004

Schröder, L., Krasimirov Dimitrov, N., Verelst, D. R., & Sorensen, J. A. (2018). Wind turbine site-specific load estimation using artificial neural networks calibrated by means of high-fidelity load simulations, In *Journal of physics: Conference series*, Institute of Physics Publishing. https://doi.org/10.1088/1742-6596/1037/6/062027

Sessarego, M., Ramos-García, N., Sørensen, J. N., & Shen, W. Z. (2017). Development of an aeroelastic code based on three-dimensional viscous–inviscid method for wind

turbine computations. *Wind Energy*, *20*(7), 1145–1170. https://doi.org/10.1002/we.2085

Smagorinsky, J. (1963). General circulation experiments with the primitive equations. *Monthly Weather Review*, *91*(3), https://doi.org/10.1175/1520-0493(1963)091<0099:GCEWTP 99–164. https://doi.org/10.1175/1520-0493(1963)091<0099:GCEWTP>2.3.CO;2

Snel, H. (1998). Review of the Present Status of Rotor Aerodynamics. *Wind Energy*, *1*(S1), 46–69.

Soize, C., & Ghanem, R. (2005). Physical systems with random uncertainties: Chaos representations with arbitrary probability measure. *SIAM Journal on Scientific Computing*, *26*(2), 395–410. https://doi.org/10.1137/S1064827503424505

Soleimanzadeh, M., Wisniewski, R., & Kanev, S. (2012). An optimization framework for load and power distribution in wind farms. *Journal of Wind Engineering and Industrial Aerodynamics*, *107-108*, 256–262. https://doi.org/10.1016/j.jweia.2012.04.024

Soleimanzadeh, M., Wisniewski, R., & Shakeri, S. M. (2010). Wind deficit model in a wind farm using finite volume method, In *Proceedings of the 2010 american control conference, acc 2010*. https://doi.org/10.1109/acc.2010.5530830

Sorensen, J. N., Mikkelsen, R. F., Henningson, D. S., Ivanell, S., Sarmast, S., & Andersen, S. J. (2015). Simulation of wind turbine wakes using the actuator line technique. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *373*(2035). https://doi.org/10.1098/rsta.2014.0071

Sørensen, J., & Shen, W. (2002). Numerical modelling of wind turbine wakes. *Journal of Fluids Engineering*, *124*(2), 393–399.

Sørensen, N. (1995a). *General purpose flow solver applied to flow over hills* (Doctoral dissertation) [Published 2003]. Risø National Laboratory. Published 2003.

Sørensen, N. (1995b). *General purpose flow solver applied to flow over hills* (Doctoral dissertation) [Published 2003]. Technical Univeristy of Denmark. Risø National Laboratory. Published 2003.

Steinbuch, M., de Boer, W. W., Bosgra, O. H., Peters, S. A., & Ploeg, J. (1988). Optimal control of wind power plants. *Journal of Wind Engineering and Industrial Aerodynamics*, *27*(1-3), 237–246. https://doi.org/10.1016/0167-6105(88)90039-6

Sudret, B. (2008). Global sensitivity analysis using polynomial chaos expansions. Elsevier. https://doi.org/10.1016/j.ress.2007.04.002

Taylor, G. I. (1938). The Spectrum of Turbulence. *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences*, *164*(919), 476–490. https://doi.org/10.1098/rspa.1938.0032

Tian, J., Su, C., Soltani, M., & Chen, Z. (2014). Active power dispatch method for a wind farm central controller considering wake effect, In *Iecon proceedings (industrial electronics conference)*, Institute of Electrical; Electronics Engineers Inc. https://doi.org/10.1109/IECON.2014.7049333

Troldborg, N. (2009). *Actuator Line Modeling of Wind Turbine Wakes* (Doctoral dissertation). http://orbit.dtu.dk/files/5289074/Thesis.pdf%20http://www.dtu.dk/English/Service/Phonebook.aspx?lg=showcommon%7B%5C&%7Did=228924

Troldborg, N., & Meyer Forsting, A. R. (2017). A simple model of the wind turbine induction zone derived from numerical simulations. *Wind Energy*, *20*(12), 2011–2020. https://doi.org/10.1002/we.2137

Troldborg, N., Sørensen, J. N., Mikkelsen, R., & Sørensen, N. N. (2014). A simple atmospheric boundary layer model applied to large eddy simulations of wind turbine wakes. *Wind Energy*, *17*(4), 657–669. https://doi.org/10.1002/we.1608

van der Laan, M., Andersen, S., Kelly, M., & Baungaard, M. (2020). *Fluid scaling laws of idealized wind farm simulations* [unpublished]. unpublished.

Vermeer, L. J., Sørensen, J. N., & Crespo, A. (2003). Wind turbine wake aerodynamics. Elsevier Ltd. https://doi.org/10.1016/S0376-0421(03)00078-2

Vreman, A. W. (2004). An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications. *Physics of Fluids*, *16*(10), 3670–3681. https://doi.org/10.1063/1.1785131

Wiener, N. (1938). The homogeneous chaos. *American Journal of Mathematics*, *60*(4), 897–936. http://www.jstor.org/stable/2371268

Zalkind, D. S., & Pao, L. Y. (2016). The fatigue loading effects of yaw control for wind plants, In *Proceedings of the american control conference*, Institute of Electrical; Electronics Engineers Inc. https://doi.org/10.1109/ACC.2016.7524969