Robust energy production optimization of a wind farm using wake steering by calibrating the FLORIS model on SCADA data

Master of Science Thesis

Maarten van Beek

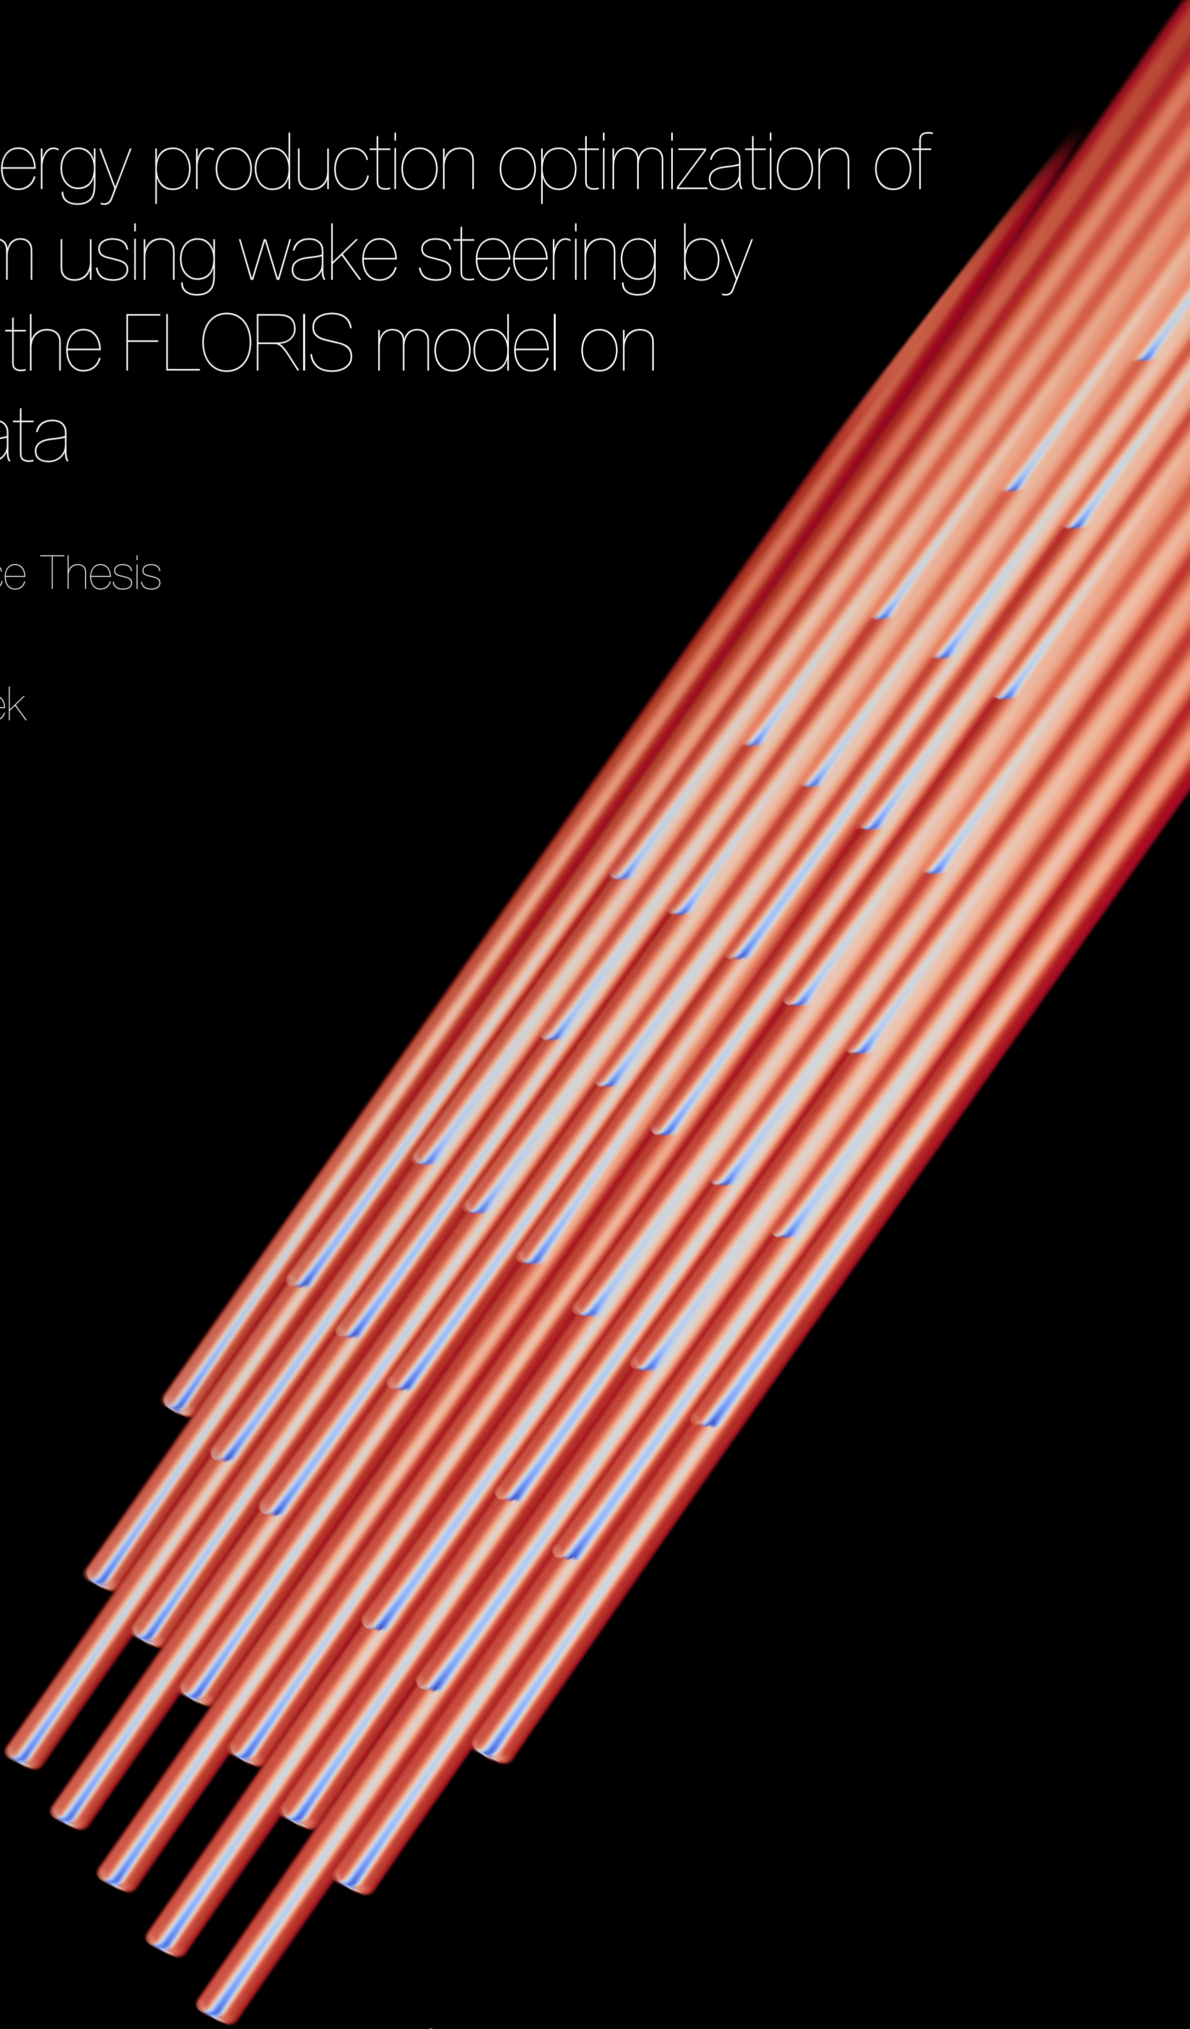

Robust energy production optimization of a wind farm using wake steering by calibrating the FLORIS model on SCADA data

Master of Science Thesis

by

Maarten van Beek

to obtain the degree of Master of Science in Aerospace Engineering at the Delft University of Technology and
Master of Science in Engineering (European Wind Energy) at Technical University of Denmark,
to be defended publicly on Thursday August 20, 2020 at 10:00 AM.

Student number:	4368134 (TU Delft)	183180 (DTU)
Project duration:	November 1, 2019 – July 31, 2020	
Thesis committee:	Dr. ir. S. Andersen,	DTU, supervisor
	Dr. ir. A. Viré,	TU Delft, supervisor
	Prof. dr. D. A. von Terzi,	TU Delft, chairman
	Prof. dr. ir. C. J. Simão Ferreira,	TU Delft, censor

Preface

This Master thesis marks the ending of a fantastic journey provided by the European Wind Energy Master program. After learning the basics of aerodynamics during my bachelor Aerospace Engineering, I decided to continue my studies in this field but applied on more sustainable machines. This program gave me the opportunity to re-discover my interests by offering courses in multiple turbine related fields. Eventually, after various conversations with different academic people, I found a project that I could work on every day with high motivation.

The basis of the project was shaped through a conversation with Jan-Willem van Wingerden, who was able to combine my interests into one form. After that, with input from my supervisors Søren Andersen (DTU) and Axelle Viré (TU Delft), I was able to form this exciting research. I owe them a big thanks in supporting me week by week, helping me to not lose sight of the bigger picture and getting me into contact with the right people. Furthermore, I would like to thank Kurt Hansen and Tuhfe Göçmen for providing me with the data sets and answering all my questions related to these. I also want to thank Bart Doekemeijer for sharing his experience with the FLORIS model and helping me shape the basis of the calibration method and Nikolai Nygaard for sharing his knowledge about his model uncertainty analysis. Finally, I would like to thank my fellow student Wessel Keijer, who helped me with his knowledge in parallel computing and interpreting my results, and all my other fellow students for the interesting discussions the past years.

I believe that this thesis contributes as a small step towards the actual use of wake steering. I hope that the calibration method proposed will prove to be useful for companies and institutions and can help them test the potential of wake steering for their wind farms.

Maarten van Beek
Delft, July 2020

Abstract

Wind farms experience significant efficiency losses due to the aerodynamic interaction between turbines. A possible control technique to reduce these losses to a minimum is yaw-based wake steering. This thesis investigates the feasibility of this technique by calibrating a surrogate model called the FLOw Redirection and Induction in Steady-state (FLORIS) model on a data set from the Lillgrund wind farm and using it to estimate the potential energy gain. The data set available is processed methodically to remove outliers and erroneous data points, resulting in a reliable and useful data set. It is used to obtain free stream wind conditions per time step and relate those to power measurements. The data set is consequently used to calibrate the tuning parameters of the FLORIS model. The calibration is done using a newly proposed method that determines the tuning parameters per combination of wind speed and turbine spacing. A difference with commonly applied calibration methods is that power measurements are used instead of predicted powers or flow field data from high-fidelity models. The performance of the calibrated model is tested through multiple uncertainty analyses. It is found that the model has a significant bias but low uncertainty by comparing the predicted wake losses with measured wake losses. This bias can potentially be reduced if atmospheric stability is taken into account. With the bias and uncertainty quantified, the FLORIS model is used to optimize the annual energy production of the Lillgrund wind farm by finding the ideal yaw angles for specific inflow conditions. A significant energy gain can be achieved when the optimal yaw angles are determined deterministically. However, the energy gain decreases drastically when uncertainty in input conditions is considered, showing that these yaw angles are not robust in terms of performance under uncertainty. More robust yaw angles can be obtained when the input uncertainty is taken into account during the yaw optimization. The energy gain achievable with these more robust yaw angles is approximately 3.4%. Therefore, it can be concluded that achieving an energy gain using yaw-based wake steering is feasible for the Lillgrund wind farm.

List of Figures

List of Tables

Contents

Nomenclature

Parameters

α	Power-yaw loss coefficient	[-]
α	Tuning parameter	[-]
α_r	Tuning parameter	[-]
$\alpha_{p,\infty}$	Free stream wind shear coefficient	[-]
β	Tuning parameter	[-]
$\Delta\theta_\infty$	Wind direction step size	[deg]
ΔU_∞	Wind speed step size	[m/s]
δ_f	Wake centerline displacement	[m]
δ_r	Wake deflection due to wake rotation	[m]
Δ_γ	Yaw misalignment angle deviation	[deg]
Δ_θ	Wind direction deviation	[deg]
ϵ	Relative wake model error	[-]
ϵ	Tuning parameter	[m]
η_f	Wind farm efficiency	[-]
γ	Wake misalignment angle	[rad]
Γ_{cr}	Circulation of the counter-rotating vortex	[m^2/s]
γ_{eff}	Effective yaw misalignment angle	[rad]
Γ_{wr}	Circulation of the wake rotation vortex	[m^2/s]
κ_f	Normalized farm energy	[-]
λ	Tip speed ratio	[-]
λ_i	Tip speed ratio	[-]
μ_U	Mean of wind speed	[m/s]
ω	Rotational speed	[rad/s]
ρ	Air density	[kg/m^3]
σ	Standard deviation in the specified direction (x, y, z) in the Euclidean space	[m]
σ_U	Standard deviation of wind speed	[m/s]
σ_γ	Standard deviation in yaw angle	[deg]
σ_θ	Standard deviation in wind direction	[deg]
θ	Initial deflection angle	[rad]
θ_∞	Free stream wind direction	[deg]
θ_{der}	Derived wind direction	[deg]
θ_{mast}	Met mast wind direction	[deg]
$\theta_{t,\infty}$	Free stream wind direction per time step	[deg]
ξ_∞	Free stream wind veer coefficient	[-]
A	Scale parameter	[-]
a	Axial induction	[-]
a_d	Tuning parameter	[-]
b_d	Tuning parameter	[-]
C_0	Velocity deficit at start far-wake region	[-]
C_P	Power coefficient	[-]
C_T	Thrust coefficient	[-]
D	Rotor diameter	[m]
E	Annual energy production	[Wh]
E_{base}	Baseline annual energy production	[Wh]
E_{nw}	No-wake annual energy production	[Wh]
E_{opt}	Optimized annual energy production	[Wh]
G_f	Energy gain	[-]
I_∞	Free stream turbulence intensity	[-]
I_{rotor}	Turbulence intensity just before the rotor	[-]
k	Block length	[time steps]
k	Shape parameter	[-]
k_a	Tuning parameter	[-]
k_b	Tuning parameter	[-]
k_y	Lateral wake expansion coefficient	[-]
k_z	Vertical wake expansion coefficient	[-]
L_{model}	Modeled wake loss	[-]
L_{obs}	Observed wake loss	[-]
N	Number of bootstrap evaluations	[-]
N	Number of discretization points	[-]
N_γ	Number of yaw misalignment discretization points	[-]

N_θ	Number of wind direction discretization points	[-]
p	Number of uncertain parameters	[-]
P_0	Turbine power	[W]
P_f	Farm power	[W]
P_w	Weighted average power	[W]
P_γ	Turbine power with yaw misalignment	[W]
P_{FLORIS}	Power predicted by FLORIS	[W]
P_{gross}	Power without wake effects	[W]
P_{net}	Power with wake effects	[W]
$P_{SCADA,68}$	68% confidence interval of measured power	[W]
$P_{SCADA,mean}$	Mean of measured power	[W]
P_{SCADA}	Measured power	[W]
$P_{t,F}$	Predicted power per time step	[W]
$P_{t,i}$	Power per turbine per time step	[W]
P_t	Measured power per time step	[W]
$P_{gross_{model}}$	Modeled power without wake effects	[W]
$P_{gross_{obs}}$	Observed power without wake effects	[W]
$P_{net_{model}}$	Modeled power with wake effects	[W]
$P_{net_{obs}}$	Observed power with wake effects	[W]
R	Rotor radius	[m]
R^2	Coefficient of determination	[-]
S_i	First-order Sobol sensitivity index	[-]
$S_{i,CI}$	95% Confidence interval of the first-order Sobol sensitivity index	[-]
$S_{ij,CI}$	95% Confidence interval of the second-order Sobol sensitivity index	[-]
S_{ij}	Second-order Sobol sensitivity index	[-]
$S_{Ti,CI}$	95% Confidence interval of the total-order Sobol sensitivity index	[-]
S_{Ti}	Total-order Sobol sensitivity index	[-]
u_g	Gaussian velocity in the wake	[m/s]
U_∞	Free stream wind speed	[m/s]
U_{mast}	Met mast wind speed	[m/s]
U_{nac}	Nacelle wind speed	[m/s]
U_{REWS}	Rotor effective wind speed	[m/s]
$U_{t,\infty}$	Free stream wind speed per time step	[m/s]
V	Lateral velocity component	[m/s]
W	Vertical velocity component	[m/s]
x_0	Near-wake cone length	[m]
x_0	Turbine location on the x-axis	[m]
X_i	Location of turbine i on the x-axis	[m]
y_0	Turbine location on the y-axis	[m]
Y_i	Location of turbine i on the y-axis	[m]
z_{hub}	Hub height	[m]
AEP_f	Annual energy production of the farm	[Wh]
AEP_{nopyc}	Annual energy production without power-yaw loss correction	[Wh]
AEP_{oc}	Annual energy production with operative turbines correction	[Wh]
AEP_{pyc}	Annual energy production with power-yaw loss correction	[Wh]
AEP_{tc}	Annual energy production with time step correction	[Wh]

Abbreviations

ABL Atmospheric Boundary Layer
AEP Annual Energy Production

BEM Blade Element Momentum

CFD Computational Fluid Dynamics
CI Confidence Interval

DIC Dynamic Induction Control
DIPC Dynamic Individual Pitch Control
DTU Danmarks Tekniske Universitet
DWM Dynamic Wake Meandering

FAST Fatigue, Aerodynamics, Structures and Turbulence
FLORIDyn FLOw Redirection and Induction Dynamics
FLORIS FLOw Redirection and Induction in Steady-state

GCH Gauss-Curl Hybrid

LES Large Eddy Simulation
LOF Local Outlier Factor

M-O Monin-Obukhov
MAPE Mean Absolute Percentage Error
MPE Mean Percentage Error

NREL National Renewable Energy Labratory

OpenFOAM Open-source Field Operations and Manipulations
OUU Optimization Under Uncertainty

PALM PArallelized LES Model
PCE Polynomial Chaos Expansion
PDF Probability Density Function

RANS Reynolds Averaged Navier-Stokes
REWS Rotor Effective Wind Speed

SA Sensitivity Analysis
SCADA Supervisory Control And Data Acquisition
SLSQP Sequential Least SQuares Programming
SOWFA Simulator for Off/Onshore Wind Farm Applications

TI Turbulence Intensity

UQ Uncertainty Quantification
UTD-WF UTD Wind Farm

WFSim WindFarmSimluator

1

Introduction

Since the start of the new millennium, a major increase in global wind power capacity has been observed, increasing from about 17 GW in 2000 to 623 GW in 2019 [62]. This trend is expected to continue due to the common understanding that there is a need for a sustainable energy transition. Harvesting wind energy is currently divided into two sub-sectors: onshore and offshore wind energy. Though a major part, 95% of the 623 GW, is currently from the onshore sector, the offshore sector is catching up. Where the onshore sector has increased its capacity the past decade by a factor of 4, the offshore sector has increased by a factor of 13 [62]. This steep increase in offshore wind energy is, among other factors, a result of the fact that wind turbines placed offshore are often clustered in so-called wind farms. These wind farms come with advantages.

First of all, a large amount of energy can be extracted from a relatively small area. Also, it is possible to transport the extracted energy to land collectively, and both installation and maintenance costs can be reduced. However, the clustering of turbines has disadvantages as well. When wind turbines are placed in close proximity, it can result in aerodynamic interaction, reducing the farm's efficiency [13]. The aerodynamic interaction is especially present when a row of turbines is aligned with the wind direction, resulting in downstream turbines being waked by upstream turbines. The flow in a wake has a velocity deficit and increased turbulence compared to the free stream flow that the upstream turbine experiences [85]. The velocity deficit results in lower power production for the downstream wind turbines, while the increased turbulence generally leads to higher fatigue loads [85].

Multiple solutions are being researched to reduce efficiency losses due to wake-turbine interaction, such as optimization of wind farm layouts [55]. However, in this thesis, the focus will be on wind farm control. There are currently five main control concepts being researched: power de-rating, turbine re-positioning, Dynamic Individual Pitch Control (DIPC), Dynamic Induction Control (DIC), and yaw-based wake steering. Research has shown that power de-rating is a viable control concept for load mitigation and power tracking, but not so much for increasing the wind farm power output [71]. DIPC, DIC, and turbine re-positioning are new research fields with great potential. However, because there are no existing models for these concepts yet, it is chosen to use the fifth concept for this research: yaw-based wake steering. This thesis aims to evaluate the feasibility of achieving an energy gain by using yaw-based wake steering for an existing wind farm.

This control concept has shown power production efficiency gains for wind farms using both low and high-fidelity models [71] and is, therefore, a promising field for further research. The effectiveness of yaw-based wake steering depends on the wind direction and ambient conditions. In theory, every possible combination of wind input conditions can be modeled using high-fidelity models to have an optimized wind farm at all times. However, this is not feasible in terms of computational effort and time [47]. Therefore, a model is needed that is accurate enough to ensure power gains while being sufficiently fast. This model can then be used as a (near) real-time wind farm controller or for optimization problems.

This is one of the reasons why the FLOw Redirection and Induction in Steady-state (FLORIS) model [47, 98] is created. It is a data-driven parametric wind farm model that can calculate the power output of a farm using relatively few (tuning) parameters, including yaw angles. The tuning parameters need to be calibrated by fitting the FLORIS model on data, either produced by high-fidelity models or using Supervisory Control And Data Acquisition (SCADA) data from existing wind farms. In this thesis, a new calibration method is proposed, which uses SCADA data. This new method enables calibration of the model for many relevant inflow conditions, which is harder to do if data is to be produced by high-fidelity models. The calibrated FLORIS model can then be used to optimize the wind farm power output for different wind input conditions. The optimization can be made more robust by taking into account and minimizing the model and input uncertainty. This optimization type is expected to lead to a robust estimation of the energy gain with and without yaw control, which is used to evaluate the feasibility of this technique on an existing wind farm.

2

State of the art

This chapter summarizes the state of the art found in the literature about yaw-based wake steering and optimization of wind farms while taking into account uncertainty. First, in section 2.1, the background on the subjects related to this thesis are discussed. Then in section 2.2, the scientific gap between current studies and this thesis is identified. Finally, section 2.3 highlights the research objectives and questions.

2.1. Background

In this section, all relevant concepts and research topics are introduced together with an overview of the research done so far.

2.1.1. Wake

When free stream air flows past a wind turbine, the properties of this flow change. The affected flow region can roughly be described as a circular tube that increases in diameter with increasing downstream distance. This tube of air is commonly called a wake. When another downstream turbine operates in the wake of an upstream turbine, it produces less power than when it is operating in free stream air. This phenomenon is called aerodynamic coupling between individual turbines [13]. It is necessary to lower the efficiency losses due to aerodynamic coupling, in order to raise the competitiveness of a wind farm. A visualization of aerodynamic coupling can be seen in figure 2.1.

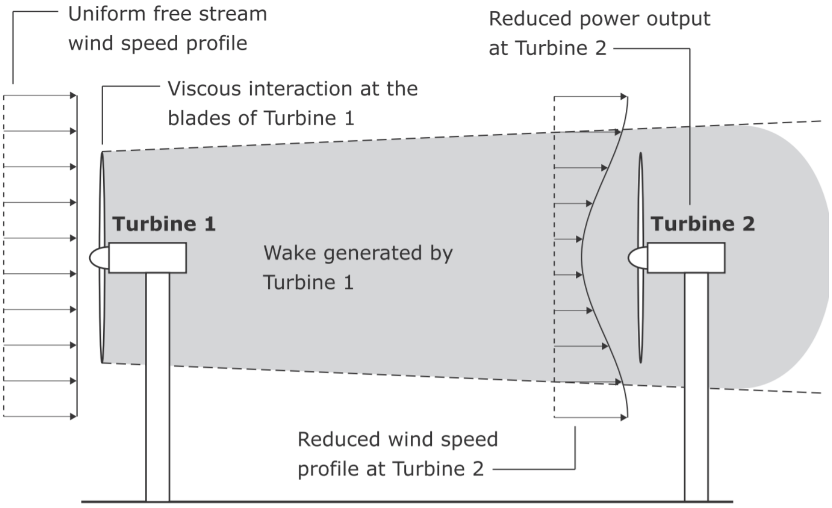

Figure 2.1: The phenomenon of aerodynamic coupling between two wind turbines aligned with the free stream wind. Taken from Kheirabadi and Nagamune [71].

The reason that the properties of the flow change is the viscous interaction of the blades of turbine 1 with the incoming wind. This interaction generates a downstream wake with a lower speed and higher turbulence intensity than the free stream wind [85]. Hence, downstream turbine 2 produces less power and will have more fatigue damage over time [85]. Note that the visualization in figure 2.1 shows the steady effects of a wake, while in reality, the wake is dynamic and hence an unsteady phenomenon.

Something that is not shown in figure 2.1 is that the wake is usually divided into two parts called the near wake region (close to the rotor) and the far wake region (further downstream). The two regions, including their typical velocity profiles, are shown in figure 2.2.

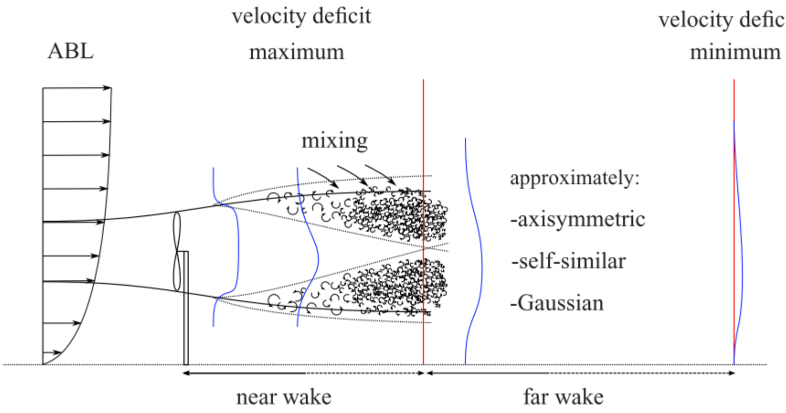

Figure 2.2: Schematic overview of the different wake regions and the mixing process. Taken from Sanderse [111].

The geometry of the turbine directly influences the flow in the near wake. Vortices created by the tips and roots of the rotor blades are present and cause the near wake region to have a high turbulence intensity and large velocity gradients. When propagating downstream, the velocity profile evens out more and becomes the far wake region. In this region, the effects of the rotor can only be seen through more large-scale effects of velocity deficit and increased turbulence intensity [45].

2.1.2. Wind farm control

As mentioned in chapter 1, five main control techniques are being researched. Three of those are visualized in figure 2.3. The oldest and thus far most intensively studied control technique shown in figure 2.3 is axial induction control. It was first proposed by Steinbuch et al. [120]. The axial induction factor indicates the flow velocity deficit of the air flowing past a wind turbine. This induction can be directly related to the thrust and power coefficient of the turbine. As shown in the top of figure 2.3, the axial induction can be controlled to lower the thrust force of turbine 1 and hence the velocity deficit for turbine 2, which results in lower electrical power production for turbine 1 and higher electrical power production of turbine 2. When applied correctly, this can result in an overall power gain of the two turbines [66]. Annoni et al. [5] showed that although this technique looks promising using low-fidelity models, with higher-fidelity models, the technique can result in efficiency losses. This discrepancy occurs due to aerodynamic phenomena that are not being captured by low-fidelity models.

The second technique, which is more novel but gaining more attention, is the re-positioning of floating wind turbines. The concept of this technique is sketched at the bottom of figure 2.3. The goal is to re-position the individual turbines in real-time to minimize the amount of wake overlap [13]. The downstream turbine is able to raise its power output when there is less wake overlap because the turbine experiences higher winds speeds. The potential of the re-positioning technique in terms of power gain looks very promising [34, 71, 108]. However, the technique can only be applied in the floating offshore wind power sector that is currently in an early development stage [13].

The third option to influence wake effects is operating a turbine that has a yaw misalignment with the incoming wind direction [65, 127]. This control technique is called yaw-based wake steering. A sketch of this technique is shown in the middle of figure 2.3. The unsteady and asymmetric loads generated along the blades of a yawed rotor, induce a force from the rotor on the flow in the crosswind direction. The force causes the wind to gain momentum and deflect in the crosswind direction. The imbalance in forces from yaw misalignment also results in the generation of trailing vortices in the wind field. These vortices influence the cross-sectional shape and also contribute to the wake deflection [37].

When an upstream turbine has a yaw misalignment with the wind flow, it operates in sub-optimal conditions, hence less power is produced than when it would be perfectly aligned. However, due to the resulting deflection of the wake, a downstream turbine may produce more power. When applied correctly, this can result in an overall increase in wind farm power, as shown through numerous studies using low-fidelity models [27, 48, 106, 109, 116] and high-fidelity models [20, 34, 37]. Also, wind tunnel experiments reported positive

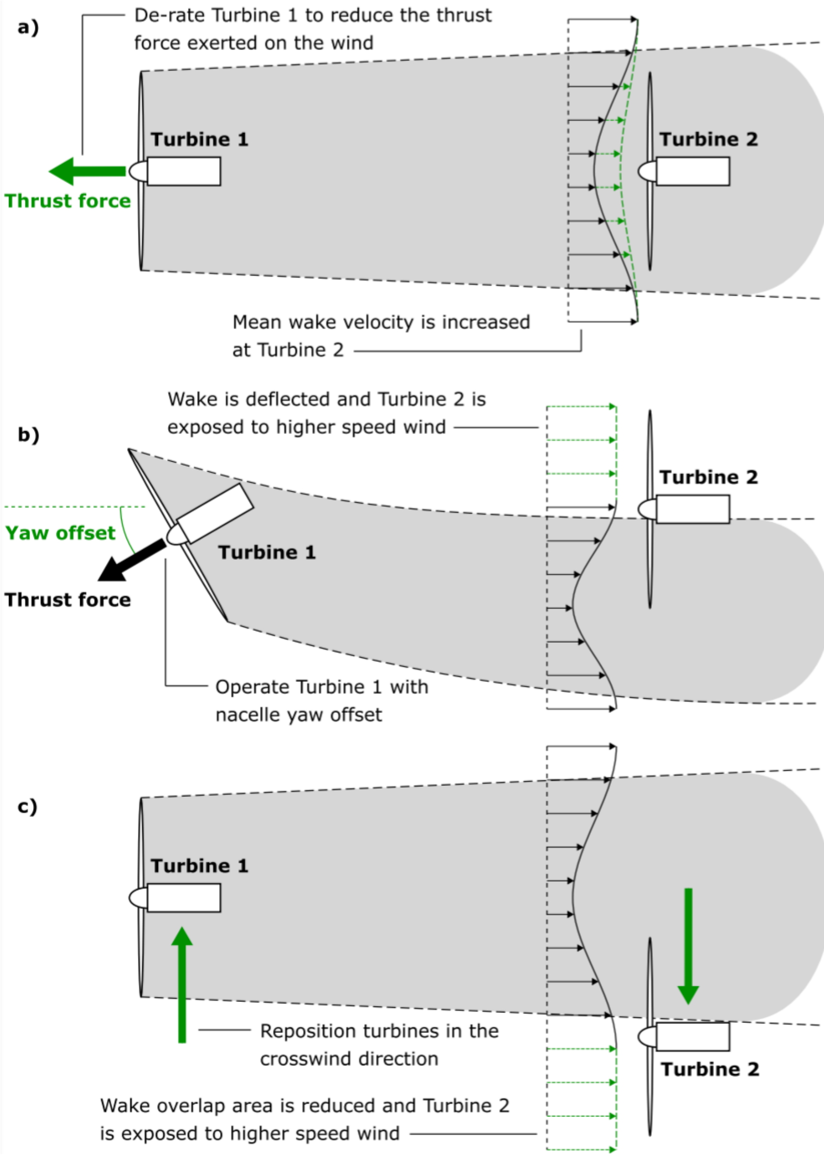

Figure 2.3: Overview of three notable wind farm control techniques proposed for mitigating the wake effect within wind farms. The schematic shows a top-view of two wind turbines aligned with the incoming free stream wind. Green elements highlight changes in turbine operation and wake conditions associated with each control concept. Taken from Kheirabadi and Nagamune [71].

gains [9, 18], and even recent field tests have shown promising results [28, 60].

Other novel control techniques are DIC and DIPC; the latter is also called the helix approach. DIC was first introduced by Goit and Meyers [54]. The idea is to dynamically regulate individual wind turbines to make them act as (vertical) flow actuators. When this control technique is applied properly, kinetic energy from higher regions in the boundary layer is transported to lower regions, which increases the total kinetic energy in the downstream boundary layer for the next turbine. This way, the total energy extraction of the wind farm can be increased by 16% compared to a reference farm in the same conditions but without control [54]. Since the introduction of this control concept, multiple studies have been done to study its potential. These studies are amongst others done by Frederik et al. [41], Munters and Meyers [93, 94].

The helix approach is recently introduced by Frederik et al. [42] and exploits DIPC to achieve a similar effect as with DIC. Dynamic variations of the wind turbine control settings lead to variations in the fixed-frame tilt and yaw moments, resulting in an enhanced wake recovery. Frederik et al. [42] have shown that the helix approach is more effective than DIC in terms of power gain while having minimal power and thrust variations.

2.1.3. Wind farm models

There are roughly three classes of wind farm models based on the level of detail of the physics modeled and the computational effort the model requires. These three classes are called low-fidelity, medium-fidelity, or high-fidelity. High-fidelity means that the physics are modeled with more detail, while low-fidelity means that more fundamental assumptions are made, and hence the physics are modeled with less detail. The higher the fidelity, the more computational effort it takes to finish the simulation. A schematic overview of a few different types of wake models is given in figure 2.4. In this figure Engineering Wake Models and Modified Reynolds Averaged Navier-Stokes (RANS) can be classified as low-fidelity, RANS and Dynamic Wake Meandering (DWM) as medium-fidelity and Large Eddy Simulation (LES) as high-fidelity.

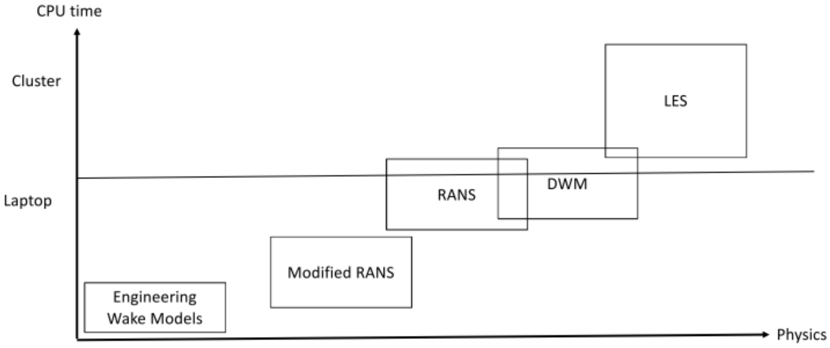

Figure 2.4: Classification of different wake models. Taken from Murcia Leon [95].

Low-fidelity models

Low-fidelity models are usually 2D parametric models. They describe the flow in a simplified way and only capture the most dominant wake characteristics. Most of these models assume the model to be steady-state or quasi-steady-state, while the wake is dynamic and unsteady in reality. The Jensen model presented by Jensen [63] is one of the first wake decay models created to estimate the velocity profile in a wind farm based on the incoming wind velocity and the axial induction per turbine. It is later extended to calculate the mean power production of a wind farm, taking into account the turbine-wake interaction by Katić et al. [68]. The model is straightforward and, therefore, easy to implement. However, the accuracy is generally lower than models that are only slightly more computationally expensive. An example of such a model is the FLORIS model, which uses the Jensen model as a basis.

The FLORIS model introduced by Gebraad et al. [47] is a data-driven parametric wind farm model that can predict the steady-state effects of yaw misalignment on the power output for a farm using relatively few tuning parameters. It is based on the Jensen model, augmented with a model for wake deflection through yaw. The model can be used to optimize a wind farm for specific ambient conditions, by finding yaw angles per turbine that result in a power gain. The FLORIS model is a vital part of this thesis and is more elaborately described in chapter 5. It has been widely used in literature to investigate the potential of yaw-based wake steering [27, 29, 50, 105, 106, 109]. Also, several studies have tried to validate the predictions by the FLORIS model by comparing their outcomes with measurements from field tests. Annoni et al. [6] showed that the Gaussian model, which is one of the surrogate models available in the FLORIS model, is best able to simulate the wake characteristics of a single turbine. Doekemeijer et al. [28] showed that the FLORIS model overestimates wake recovery because it could not take into account essential terrain effects. However, they were still able to achieve power gains with real turbines by calculating optimized yaw angles using the model.

Another low-fidelity wind farm model that exists is the Frandsen model [40], which assumes a linearly expanding wake and that the velocity deficit depends on the distance from the rotor only. One more model is the Larsen model [76]; it models the axial and radial velocity as well as the width of the wake based on the turbulent boundary conditions. Next to those models, a variant of the FLORIS model exists that is not steady-state, but quasi-steady-state. It is presented by Gebraad and van Wingerden [46] and called FLOw Redirection and Induction Dynamics (FLORIDyn). There is also a model that is similar to the FLORIDyn model called SimWindFarm [56]. A recent study by Hulsman et al. [61] proposes yet another method of creating a low-fidelity model. In this paper, they develop a fast and reliable surrogate model for yaw control using Polynomial Chaos Expansion (PCE), which can mimic the results found using the high-fidelity software EllipSys3D closely.

Medium-fidelity models

Medium-fidelity models are often 2D models that, contrary to most low-fidelity models, can capture the dynamics of the flow in the wake. This increased level of detail of the physics included in the model comes with an extra computational cost.

One of these medium-fidelity models is the DWM model first introduced by Larsen et al. [77]. This model is an engineering wake model designed to physically model both the wake deficit evolution and the unsteady wake meandering. The basic conjecture behind the model is that the large scale lateral and vertical turbulence components in the atmospheric boundary layer drive the wake transportation. The DWM model includes a stochastic model of the downstream wake meandering, a model characterizing the mean wake deficit and a model that characterizes the added wake turbulence.

The model is later calibrated, validated, and adjusted to make it fully compatible with an aero-elastic model by Madsen et al. [87]. Further validation has been done on the DWM with focus on loads [78–80]. The model is also compared to low-fidelity models, such as the Larsen and Frandsen model by Reinwardt et al. [107]. In this study is concluded that the DWM model coincides best, in most cases, with the measurement results. An advantage of the DWM model is that it calculates both the power and the loads, although the computational cost is significantly higher than, e.g., the FLORIS and Jensen model.

Other medium-fidelity models are for example the 2D Ainslie model [2] and the WindFarmSimluator (WFSim) first introduced by Boersma et al. [14]. Both are dynamic models, and the latter model either neglects or estimates terms involving the vertical dimension of the Navier-Stokes equation to decrease the computational cost. Next to that is the sparsity of and structure in the system matrices exploited to reduce the computational cost even more. This partial inclusion of the vertical dimension results in improved results when they are compared to data from high-fidelity models. At the same time, the calculation is still fast enough to have the potential of being used as an online wind farm controller.

High-fidelity models

High-fidelity models use the differential relations of fluid mechanics. Some high-fidelity wind farm models use LES, through which they are able to solve temporally and spatially-filtered forms of the three-dimensional Navier-Stokes equations [71]. The main disadvantage is that simulations using high-fidelity models can take up to weeks on a cluster of processors.

One of the high-fidelity models that uses LES is the Simulator for Off/Onshore Wind Farm Applications (SOWFA) model [20, 99]. It is a Computational Fluid Dynamics (CFD) solver based on Open-source Field Operations and Manipulations (OpenFOAM) libraries which is coupled with the Fatigue, Aerodynamics, Structures and Turbulence (FAST) tool [67], which is an open-source wind turbine simulator developed by National Renewable Energy Labratory (NREL). The usability of SOWFA is enhanced by creating a super controller that makes it possible to test various wind plant control techniques, which is presented in the study by Fleming et al. [33]. Though able to simulate the flow in a wind farm realistically, SOWFA is too computationally intensive to be used as an online wind farm controller. Therefore it is often used to validate the results of lower-fidelity models such as the FLORIS model, for example, done by Doekemeijer et al. [27, 29], Gebraad et al. [49].

Multiple other high-fidelity wind farm models use three-dimensional LES flow models. Examples are a model developed at UT Dallas, UTD Wind Farm (UTD-WF) [89], SP-Wind [3] developed in Leuven, and PArallelized LES Model (PALM), a model developed at the Institute of Meteorology and Climatology at Leibniz Universität Hannover, called PALM [88]. Within these models, different approaches are used to model the turbine. For example, the UTD-WF model has an option for actuator disc and actuator line theory. In contrast, the SP-Wind model always uses actuator disc theory, which reduces the computational cost of the simulations.

Next to these wind farm models are also other general-purpose CFD solvers being used for wind farm control. For example, EllipSys3D, which is developed at DTU/Risø [118] and, OpenFOAM [1], which is open-source software. The performance of these two solvers are compared by Cavar et al. [19]. They found that Ellipsys3D is about 2-6 times faster than OpenFOAM while achieving the same order of accuracy on similar or identical computational meshes.

High-fidelity models are generally able to capture more complex aerodynamic phenomena than the lower-fidelity models. This allows for accurate flow predictions, though the computation times are often long.

2.1.4. Optimization under uncertainty

The FLORIS model is usually used to optimize the wind farm using a deterministic approach (only one output value instead of a probability density function) by using inputs in which uncertainty is present. Additionally, every model has a certain prediction error. An uncertainty analysis can be done to capture this model prediction error. This uncertainty analysis aims to perform an optimization while including uncertainty.

In this thesis, Uncertainty Quantification (UQ) is used to improve the certainty of the expected outcome. Taking into account uncertainty in the design process of wind farms is called Optimization Under Uncertainty (OUU). This combination of UQ and optimization techniques is, among other things, used to optimize wind farm layouts [55] or wake steering [105, 106]. These studies take into account different sources of uncertainty. González et al. [55] takes into account the uncertainty in the wind direction while Quick et al. [105] takes into account the uncertainty of the yaw angles.

Another study aiming to demonstrate the influence of uncertainty in wind direction, due to dynamic changes and measurement errors on active wake steering, is done by Rott et al. [109]. This study found that including the uncertainty in wind direction can successfully increase the performance of the reference wind farm.

In the more recent study of Simley et al. [116], OUU is applied by taking both yaw position uncertainty and inflow direction variability into account. In this study is found that inflow direction variability is the dominant source of uncertainty. Quick et al. [106] builds further upon this prior work on OUU by taking into account uncertainty in turbine yaw positions and inflow direction, speed, shear, and turbulence intensity during the optimization of turbine yaw offsets for wake steering strategies. They found that OUU results in wake steering strategies that are more conservative than the deterministic approach in terms of yaw positions. The OUU approach produces up to about 4% more power than wake steering strategies formulated using the deterministic approach when the input uncertainties are taken into account [106].

2.1.5. Data analysis

An important part of this thesis is the analysis of SCADA data to obtain a useful and reliable data set. Data sets come with invalid data points due to, for example, a temporary broken or missing sensor or a turbine that is in a transition mode. Therefore, the SCADA data available should be intensively studied. The wind farm considered in this thesis is the Lillgrund wind farm. The data of Lillgrund has been used for multiple analyses found in the literature. Dahlberg [23] investigates the power performance of individual turbines, as well as the whole wind farm. Göçmen and Giebel [51] estimate the Turbulence Intensity (TI) using the Rotor Effective Wind Speed (REWS), which is derived using a 1 Hz SCADA data set. Gaumond et al. [43, 44] checks the accuracy of several engineering models by comparing their wake loss predictions with wake losses obtained through SCADA data.

All these studies had to process their data sets before using it in their research to draw conclusions. Most of them followed the analysis methods described by Hansen [57]. These methods involve checking the quality of the data, for example, by creating a scatter plot and removing apparent outliers. Another method to reduce less apparent outliers can be done using the Local Outlier Factor (LOF) method [17]. The LOF algorithm computes a score that indicates the abnormality of observations based on the neighboring observations. Hansen [57] also shows a way to find yaw position offsets using wind direction measurements or via power deficits along a line of turbines. When the yaw position offsets are known, the nacelle positions can be corrected and used to estimate the inflow direction when, for example, data from a meteorological mast is missing. Similarly, the inflow wind speed can be obtained through the REWS of the upstream turbines. Finally, Hansen [57] emphasizes the importance of having a robust synchronization of the data from different sources of measurements.

2.2. Scientific gap

So far, in the literature, the focus has been on whether power optimization through wake steering is possible using parametric models, how accurate these parametric models are, and what the influence is on the results when uncertainty is taken into account. This thesis combines the methods and outcomes of these studies to research the feasibility of an annual energy gain for an existing wind farm while taking into account uncertainty. The main difference compared to existing studies, therefore, is that an actual wind farm is modeled of which SCADA data has been collected. This data can be used to calibrate the FLORIS model accurately with a newly proposed calibration method. This data can also be used to assess the performance of the model. Another difference between the studies found in literature and this thesis is that the focus is on Annual Energy

Production (AEP) rather than power production.

2.3. Research objective

The literature study above results in the main research objective of this thesis being:

> “To estimate the feasibility of achieving an energy gain by optimizing the use of yaw-based wake steering on a wind farm, by calibrating a parametric wind farm model on SCADA data and taking input and model uncertainty into account."

This main objective can roughly be divided into three sub-parts. The SCADA data that is going to be used consists of raw data that needs to be analyzed and processed. Therefore, the first sub-objective is to filter the raw data, resulting in a useful and reliable data set by performing an intensive analysis on the data. This involves filtering for outliers in the data, checking whether the observed behavior makes sense, and determining what data is useful to calibrate the FLORIS model.

The parametric wind farm model that will be used is called the FLORIS model. The sub-objective of the FLORIS model in this thesis is to optimize the power production of a wind farm while taking into account wake behavior, by calculating the ideal yaw position per turbine in the wind farm for specific ambient conditions. The power gain obtained using the ideal yaw positions can be determined and combined with data about the wind speed and wind direction, to estimate the annual energy gain.

Aside from this, it is important to know and increase the robustness of the optimized yaw angles obtained. Therefore, the last sub-objective is to determine and increase the robustness of the optimized yaw angles and, hence, the obtained energy gain by performing an uncertainty and global sensitivity analysis on the model.

Five research questions are created to reach these objectives. Answering these five questions will result in the achievement of the research objectives and hence completion of the thesis. These five research questions are:

1. In what way can SCADA data of a wind farm be used to calibrate the FLORIS model accurately?
2. What is the sensitivity of the tuning parameters in the FLORIS model that need to be calibrated?
3. What is the performance of the calibrated FLORIS model, and what is the optimization potential for the wind farm?
4. What is the robustness of the optimized model, and can it be increased using an uncertainty analysis?
5. Is it feasible to achieve an energy gain for a wind farm using yaw-based wake steering?

3

Data analysis

This chapter elaborates on the available data and the applied processing techniques. First, in section 3.1, it is explained what SCADA data is available. Then in section 3.2 is described what data is actually used and finally, in section 3.3, it is explained how the data is processed to obtain a reliable data set.

3.1. Available data

The first data set used is from Vattenfall and provided by Danmarks Tekniske Universitet (DTU). It contains SCADA data from turbines located in the Lillgrund wind farm, as well as measurement data from nearby meteorological masts. A second similar data set is provided by DTU of the same wind farm. The main difference is that the second data set consists of 1 Hz data compared to 10-minute averages for data set 1 and that the measurements are performed during a different measurement campaign. More information about data set 2 can be found in section 4.2.2.

The Lillgrund wind farm [64] is owned by Vattenfall and located in Öresund, the water between Malmö, Sweden, and Copenhagen, Denmark. The farm consists of 48 turbines, which are placed in rows. The original labels of the turbines are combinations of letters (rows A-G) and numbers (rows 1-8). The layout of the wind farm, including the row, column, and turbine labeling, can be seen in figure 3.1. Note that the column labeling is denoted with a 'C'.

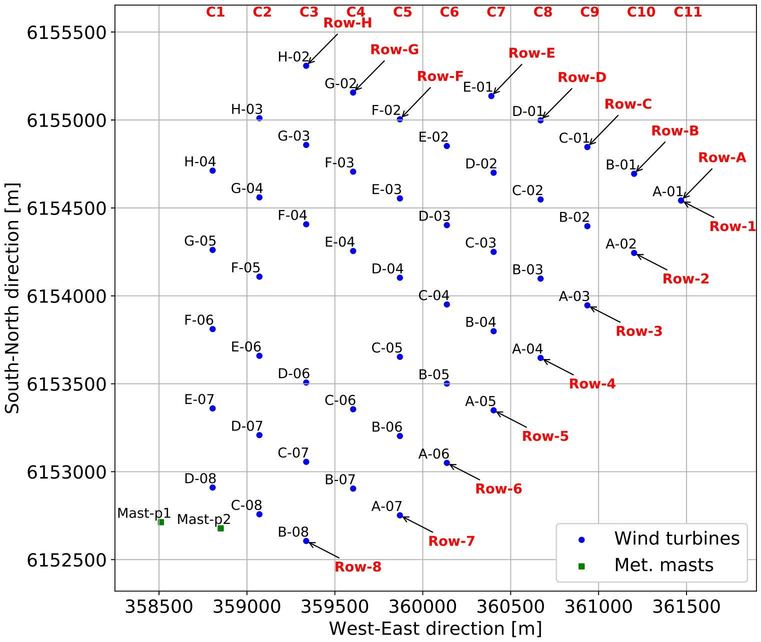

Figure 3.1: The layout of the Lillgrund wind farm including turbines, meteorological masts, column and row labelling.

The wind farm is densely packed with a spacing of 3.3D between turbines in the numbered rows, 4.3D in the lettered rows, and 4.8D in the columns, where D indicates the rotor diameter in m. The gap in the farm in row D and E is due to shallow waters, making it impossible for vessels to maneuver in this area [64].

The turbines are Siemens SWT-2.3-93 turbines [64]. These turbines have a rotor diameter of 92.6 m, a hub height of 65 m, and a rated power of 2.3 MW. The corresponding rated wind speed is around 12 m/s. The rated capacity of the farm is 110 MW. Note that the hub height of Siemens SWT-2.3-93 turbines are usually higher, but are lower for the turbines at the Lillgrund site due to height restrictions.

There have been two meteorological masts named Mast-p1 and Mast-p2 at the Lillgrund site. Mast-p1 was placed before the construction of the wind farm and removed just before the construction started. Both masts were 65 meters high and had anemometers and wind direction vanes at several levels.

The SCADA data measurement campaign was from 01-01-2008 till 31-12-2012. All turbines have a SCADA system available to record electrical power, rotational speed, local wind speed, pitch angle, and nacelle position measurements. These measurements are stored as 10-minute statistical values.

Additionally, there have been two measurement campaigns with the meteorological masts. Mast-p1 performed a measurement campaign from 01-09-2003 till 01-03-2006. It measured the wind speed, wind direction, and temperature at several heights and pressure at one height. The highest wind speed measurement is at 65 m, the highest wind direction, and temperature measurement at 61 m. The height of the pressure measurement is unknown. Mast-p2 performed its first measurement campaign from 19-05-2006 till 23-01-2008 and the second from 11-12-2008 till 16-11-2010. Mast-p2 performed the same measurements at the same heights as Mast-p1. The measurements are stored as 10-minute statistical values for both masts, similar to the SCADA data of the turbines. For the second campaign of Mast-p2, only 10-minute averaged mean values are available.

3.2. Used data

Data set 1 is large, and it is essential to separate the useful data. However, to make the calculations and findings from this data reliable, as much data as possible should be used. The electrical power measurements from the SCADA data are critical in the calibration method described in section 6.4.

Next to the electrical power, also the rotational speed and pitch angle are used. Combining these three measurements is especially relevant for the determination of the REWS as described in section 4.2. Finally, the nacelle position measurements are used to obtain the wind direction data that the FLORIS model is calibrated on. So from the SCADA data, all possible measurement types are used.

Only the wind speed and wind direction measurements are used from the measurement campaigns of the meteorological masts. They are, among other things, used to obtain a wind rose and to map the TI, as described in chapter 4. The temperature and pressure measurements are not used. This choice is made because there was not sufficient information to, for example, estimate the atmospheric stability. In theory, the wind speeds and wind directions at several heights could be used to estimate the wind shear and wind veer. However, it is chosen not to do this because the measurements are too inconsistent for this purpose.

3.3. Outlier detection and removal

Before the data becomes useful, all data should be quality checked. Erroneous observations should be identified and excluded. These erroneous observations can be the result of several things, among others:

- A defect or imperfect sensor.
- A miscalibrated sensor, e.g., after maintenance.
- A turbine being in a transition mode, e.g., start, stop, or emergency stop.
- A turbine being power de-regulated using the pitch angle.

Most of the data points provided come with a mode number. These mode numbers are listed in table 3.1 and defined by the creator of the provided data set. Note that mode number 2 is missing. These mode numbers can be used to perform the first filtering. For example, mode number 5 means that the turbine is starting or stopping, so all data points having mode number 5 are removed. Similarly, mode number 9 means that the sensor is faulty, so these data points are removed as well.

Another powerful way to identify outliers is to make a scatter plot of the data and visually identify outliers or unrealistic values. For example, in figure 3.2, a scatter plot is made from the power measurements of turbines in C1 (figure 3.1) corresponding to western winds. Some outliers are easily detected, for example, the

Table 3.1: List of mode numbers and their description

Mode number	Description
0	Not validated
1	Validated
3	Icing
4	Low nacelle wind speed
5	Event, start/stop
6	Idling
7	Unknown
8	Curtailed or stopped
9	Faulty

ones in the black circles. Other erroneous observations might be removed since they do not have a realistic value, such as a wind direction measurement of 999°.

So another way to remove erroneous observations is by removing unrealistic values. This is done for all data sets. As just mentioned, for the wind direction, all values should be in the range between 0 and 360, for power, the values should be between 0 and 2.35 MW (note that there might be an overshoot above 2.3 MW for power), for wind speeds, the values should be between 0 and 35 m/s (high values are irrelevant for this research) etc.

Finally, there may also be other outliers in the data that are harder to spot. Take, for example, the scatter plot in figure 3.2, where the power measurements of C1 are plotted against the free stream wind speed (derived in chapter 4). Note that only western winds are considered in this example. Some outliers are obvious and can be visually detected; in this figure, they can be found inside the black circles. However, there are also less obvious outliers that are hard to identify and remove manually. To make this a more automated process, the LOF method [17] is used.

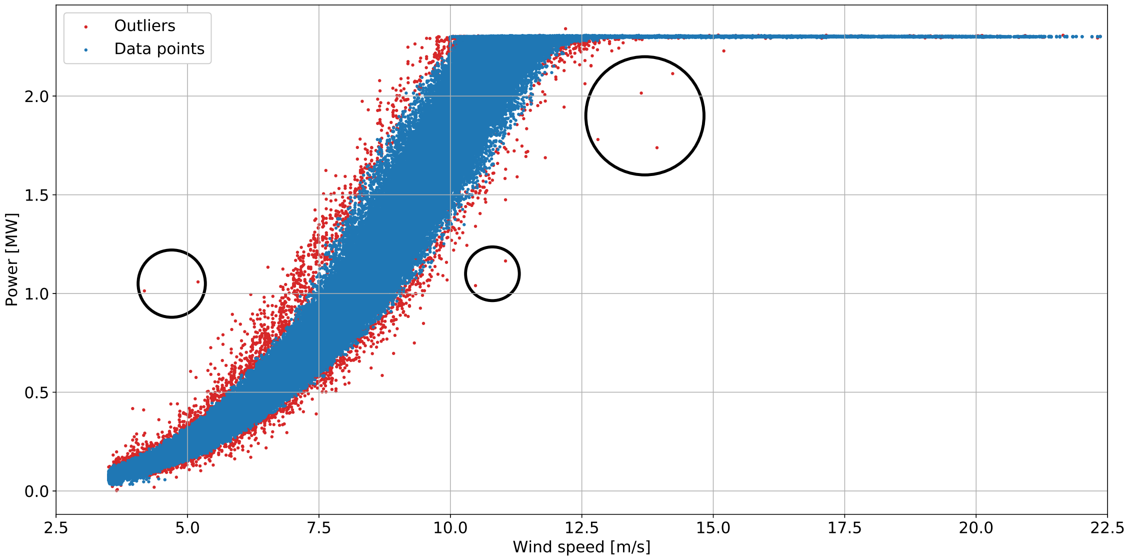

Figure 3.2: Scatter plot of the electric power versus free stream wind speed, including the observed outliers, found using the local outlier factor method. The black circles indicate apparent outliers.

The LOF algorithm computes a score that indicates the abnormality of observations based on the neighboring observations. This score is called the LOF. It indicates the local density deviation of a given data point with respect to its neighbors. The LOF method can easily be implemented through the module `sci-kit learn`. It is an open-source module created for Python, which has multiple and efficient machine learning tools that can be used for (predictive) data analysis.

The primary input for this tool is the amount of k-nearest neighbors. It is typically chosen larger than the minimum amount of cluster observations needed and smaller than the maximum number of close-by-observations that can potentially be local outliers. Unfortunately, these numbers are often not known, and the amount of k-nearest neighbors is found iterative. Typically 20 neighbors work well. Next to that, it is possible to force the LOF algorithm to have a certain percentage of contamination. This might be needed if obvious individual outliers are not captured without this forced contamination.

In figure 3.2 the result of the LOF method can be observed. The method works well and has as a result, together with all previous filters, that the remaining data set contains almost only valid and useful data points that capture the general behavior and possible under/overshoots well. This is important while calibrating the FLORIS model. The LOF method is therefore used on all relevant data sets. For example for the electric powers per turbine row or column, used to fit the FLORIS model on (section 6.4), but also to identify outliers in the free stream TI, see figure 4.7.

It is important to note that the LOF method is not 100% perfect. Some data points may be labeled as an outlier while they are not, and other points may be outliers but might not be identified as such. However, the method seems to capture the essential outliers correctly, and it can, therefore, be used confidently.

4

Free stream conditions

This chapter analyzes the general wind behavior at the Lillgrund site. First, the wind speed, wind direction, and turbulence intensity behavior are analyzed in section 4.1. Then, the free stream conditions are determined for as many time steps as possible in section 4.2. Also, in this section, the turbulence intensity is mapped on the wind speed and wind direction.

4.1. Meteorological mast 1

For this thesis, it is crucial to have an accurate estimation of the free stream wind conditions. Free stream conditions relevant for the FLORIS model are the free stream wind speed U_∞, wind direction θ_∞, turbulence intensity I_∞, shear $\alpha_{p,\infty}$ and veer ξ_∞. Due to limited availability of data, the focus will be on estimating U_∞, θ_∞ and I_∞ only. The data of Mast-p1 is first analyzed to get an overall idea of the wind behavior at the Lillgrund site. Mast-p1 performed a measurement campaign from 01-09-2003 till 01-03-2006, so before the wind farm was constructed. A quality check, as described in section 3.3, is performed on the data to obtain valid data. This data is then ready to be analyzed.

In the first part of the analysis, a wind rose is created with a wind direction bin width of 5° and a wind speed bin of 3 m/s. The wind rose can be found in figure 4.1a. Note that θ_∞ is defined as positive clockwise. It can be seen that the dominant wind direction is west-southwest. Figure 4.1b and the differences with figure 4.1a is further discussed in section 4.2.2.

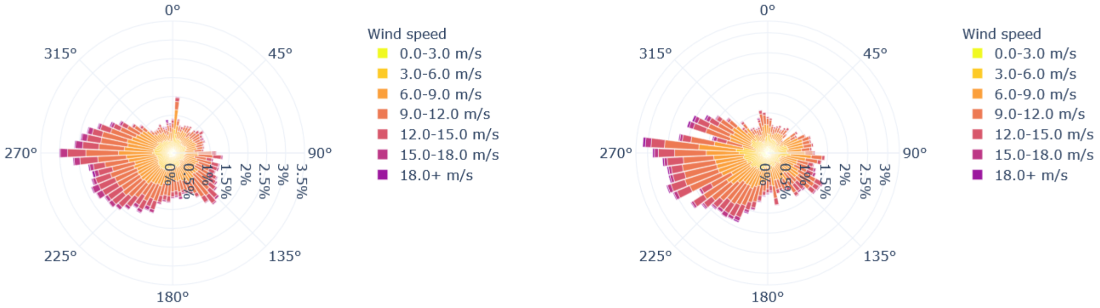

(a) Mast-p1

(b) Derived free stream

Figure 4.1: The wind rose from Mast-p1 data and the wind rose from the derived free stream data.

Secondly, probability histograms for U_∞, θ_∞ and I_∞ are created that can be found in figure 4.2. Note that the histograms are normalized such that the summed area of the bars is 1. It can be seen that the wind speed histogram has the shape of a Weibull distribution and that the most frequently occurring wind speed is around 8 m/s. In the same way, it can be seen that the most frequently occurring turbulence intensity is around 6%, and the most frequently occurring wind direction is around 270 degrees.

$$f(U) = k\frac{U^{k-1}}{A^k}\exp\left(-\left(\frac{U}{A}\right)^k\right) \tag{4.1}$$

A Weibull probability function (equation (4.1)) is fitted on the wind speed distribution following the method from Berg et al. [10], by using the first and second (non-central) moments of the wind speed data, also known as the mean and standard deviation. Through this method, the shape parameter k and scale parameter A are found to be 2.27 and 9.44, respectively. It can be seen that the function fits well. Note that all wind speeds independent of their corresponding wind direction are used to create the Weibull function.

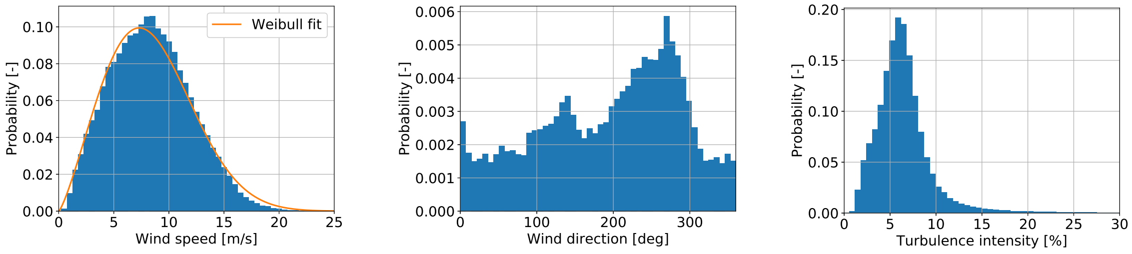

(a) Distribution of U_∞ (b) Distribution of θ_∞ (c) Distribution of I_∞

Figure 4.2: Histograms showing the probability in fractions for different variables for the Mast-p1 data. A Weibull curve is fitted on the distribution of U_∞.

4.2. Free stream conditions

The overall behavior of the wind is known through the analysis done in section 4.1. However, estimates of the free stream wind conditions are needed for every time step.

The ideal option to obtain U_∞, θ_∞ and I_∞ for time steps for which also power measurements are available, would be to use Mast-p2. Unfortunately, there are some disadvantages to using this mast. First of all, the measurement campaign of Mast-p2 is shorter than the measurement campaign of the available SCADA data, hence limiting the data that can be used to calibrate the model. Secondly, the second measurement campaign only consists of mean values, so it cannot be used to estimate I_∞. Thirdly, the mast experiences shading effects of the wind farm for wind directions from approximately 345° to 110° (clockwise). The shading effects have a result that lower wind speeds, higher turbulence intensities, and distorted wind directions are measured. Hence, it is impossible to estimate the free stream conditions accurately for wind directions that cause shading effects on the mast. Therefore, an alternative approach is used to obtain free stream wind conditions for as many time steps as possible.

4.2.1. Wind direction

First of all, a part of the data set is available that is already processed to estimate the free stream wind directions. The processing method is as follows: "I have performed a re-calibration of the yaw positions [112] and determined an undisturbed wind direction for the whole wind farm in the range 0-360 degrees; stored as channel 92." (K. S. Hansen, personal communication, January 30, 2020).

The performed steps are summarized below.

1. The nacelle positions are corrected for possible offsets.
2. Based on the nacelle positions of all operating turbines, the first estimate of θ_∞ is obtained.
3. Based on this first estimate, three to four free stream turbines that operate in an undisturbed flow on the edge of the farm are selected.
4. The θ_∞ is then estimated by averaging the nacelle positions of those free stream turbines.

The results of this method are validated by comparing the derived θ_{der} and the θ_{mast} measured by Mast-p2. As previously mentioned, this mast measures distorted wind directions for the range 345 to 110 degrees; hence these are not taken into account. All available time steps with a valid mast measurement are used. The wind directions are plotted against each other in figure 4.3. The orange line represents a binned median obtained with a wind direction bin of 5°. The median value is used, because a small change in wind direction can be the difference between, for example, 359° and 1°, hence oppositely influencing the mean value.

It can be seen that there is not a perfect match, which would be the case if all data points fit perfectly on the black line. On average, the derived wind direction is similar, but there is some uncertainty. Note the

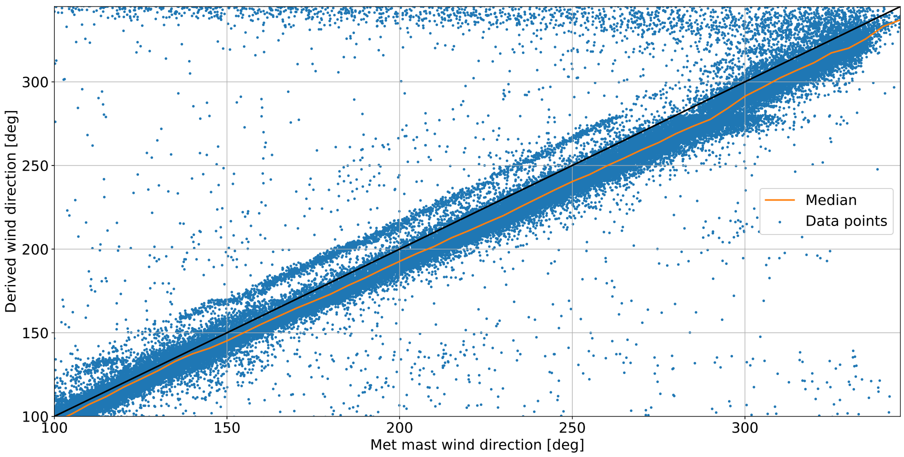

Figure 4.3: Scatter plot of θ_{der} versus θ_{mast}, including the binned median values.

upper 'line' of data points around 340° on the y-axis. An explanation for this line can be that the actual wind direction is around 340°, while the mast is still influenced slightly by the distortion from the wind farm. The difference between the observations can be due to a drift in wind direction between the meteorological mast's measurement location and the turbines used to derive the wind direction. However, this would suggest that the closer the turbines are to the meteorological mast, the closer the observations should be. In that case, the best match in observations is expected for the sector 180-270°, but it can be seen that the match is actually the closest for 100-150°. Although the data does not show a perfect match, the result of the method is still considered to be usable when the uncertainty of the wind direction is kept in mind.

4.2.2. Wind speed

The second step is to obtain U_∞ per time step. Every turbine in the farm has an anemometer that measures the nacelle wind speed U_{nac} on the hub behind the rotor. The measurements from these anemometers do not give reliable estimates of the free stream wind speeds, because the rotor slows down the wind and induces added turbulence intensity. Hence, these measurements cannot be used to estimate U_∞ without a correction factor. Therefore, a correction factor needs to be obtained that can relate the REWS U_{REWS} to U_{nac}. U_{REWS} can then, in turn, be used to find U_∞.

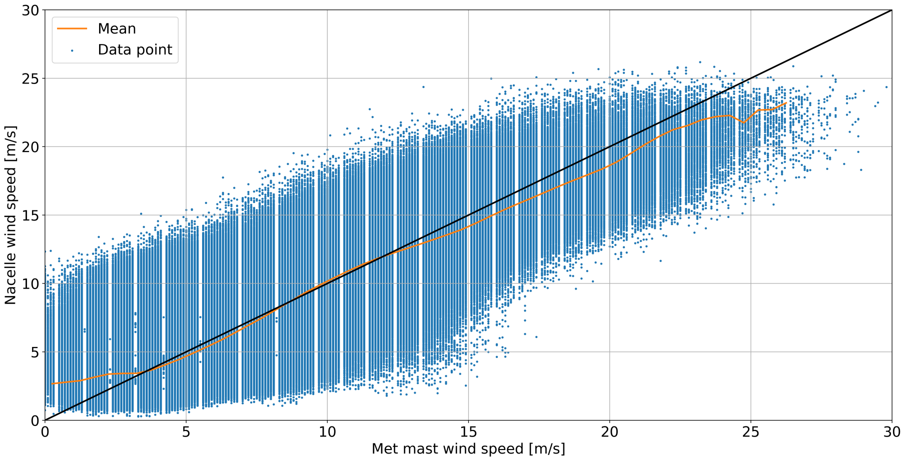

Figure 4.4: Scatter plot showing U_{nac} versus U_{mast} using the 1 Hz data of turbine C-08; the mean is also shown.

Another data set is used to find this correction factor. This one is also from Vattenfall and provided by DTU. The data set contains 1 Hz data from Mast-p2 and turbines C-07, D-07, C-08, and D-08, obtained during a measurement campaign from 06-2012 till 01-2013. The same parameters are known as from the 10-minute averaged data, for example, electrical power, rotational speed, pitch angle, etc. These three parameters are used to obtain the REWS using the method described by Göçmen et al. [52]. The same data set and method is used by Göçmen and Giebel [51].

First, the turbine measurements are compared to the mast measurements. Figure 4.4 shows that U_{nac} is lower than U_{mast} for wind speeds near cut-in (3.5 m/s) and for wind speeds above rated (12 m/s). This is a confirmation that U_{nac} cannot be used to estimate U_∞. A quality check of the 1 Hz data set is performed before it is used to derive the REWS. On top of the usual outlier detection method described in section 3.3, additional filters on specific rates are applied. For example, unrealistic high values of ω/U_{nac} are removed, where ω is the rotational speed in rad/s.

$$P = \frac{1}{2}\rho C_P(\lambda,\theta)\pi R^2 U_{REWS}^3 \tag{4.2}$$

$$C_P(\lambda,\theta) = c_1\left(\frac{c_2}{\lambda_i} - c_3\theta - c_4\theta^{c_5} - c_6\right)\exp\left(-\frac{c_7}{\lambda_i}\right) \tag{4.3}$$

$$\lambda_i = \left[\left(\frac{1}{\lambda + c_8\theta}\right) - \left(\frac{c_9}{\theta^3+1}\right)\right]^{-1} \tag{4.4}$$

The basic principle of the method is to estimate the power coefficient $C_P(\lambda,\theta)$ as a function of tip speed ratio λ and pitch angle θ in deg and combine it with the power equation to find U_{REWS} in m/s. The power equation can be found in equation (4.2) where ρ is the density of air in kg/m^3 and R the rotor radius in m. The generic C_P expression, first proposed by Heier [59], is used and can be found in equation (4.3). The C_P expression needs λ_i which is defined in equation (4.4). The same values for coefficients $c_1 - c_9$ as used by [52] and listed in Göçmen et al. [53] are used and found to give a good agreement for the turbines in Lillgrund by Göçmen and Giebel [51]. The coefficients can be found in table A.1.

Since λ_i depends on λ, which in turn depends on U_{REWS}, this method needs to be applied iteratively. It is chosen to use the Newton-Raphson method for this. Details on the Newton-Raphson method and how it is applied here can be found in appendix A, and the code implementation in van Beek [125].

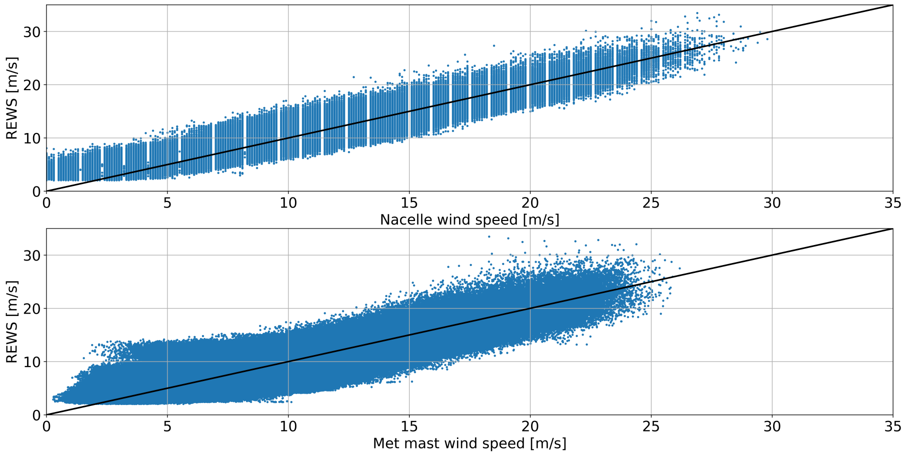

Figure 4.5: Scatter plot showing the derived U_{REWS} versus U_{nac} and U_{mast} using the 1 Hz data of turbine C-08.

The resulting U_{REWS} for turbine C-08 is compared to U_{nac} and the measurements from Mast-p2 U_{mast} in figure 4.5. This graph looks similar to the graphs found in the paper of Göçmen and Giebel [51]. It can be seen that U_{nac} compared to U_{REWS} underestimates the wind speed for low wind speeds (below 9 m/s). That is as expected and was also observed in figure 4.4. It can also be seen that U_{REWS} compared to U_{mast} consistently underestimates the wind speed for higher wind speeds (above 9 m/s). This can be explained by the fact that

U_{mast} is a point measurement while U_{REWS} is considering the wind speed seen by the whole rotor [51], which means that U_{REWS} automatically includes a geometrical averaging of the wind speed between 21.5 m and 114.5 m height. Generally, it can be concluded that these graphs show realistic results and that the method is applied correctly. The derived U_{REWS} values are therefore considered to be verified.

Now that the applied method is verified, the 1 Hz data needs to be transformed into 10-minute averages. Since some data points are removed, and the 10-minute average values should give a proper estimate, only 10-minute averages with a minimum of 90% of time steps are kept. That means that at least 540 out of 600 time steps should be available for every 10-minute average.

Table 4.1: Scores of the k-fold cross-validation using ten folds for two different models on the U_{REWS} versus U_{nac} data.

	Mean	**Std**	**a**	**b**	**c**
First order	0.97	+/- 0.01	0.2535	0.9399	0
Second order	0.97	+/- 0.02	0.5396	0.8797	0.002806

A correction factor can then be found between 10-minute average U_{REWS} and U_{nac} by fitting a first- or second-order line on the data in figure 4.6. To make sure that the best fit is used, k-fold cross-validation is applied [11]. The cross-validation is checked for two different model fits. A first-order model and a second-order model, see equation (4.5).

$$f(x) = a + bx + cx^2 \tag{4.5}$$

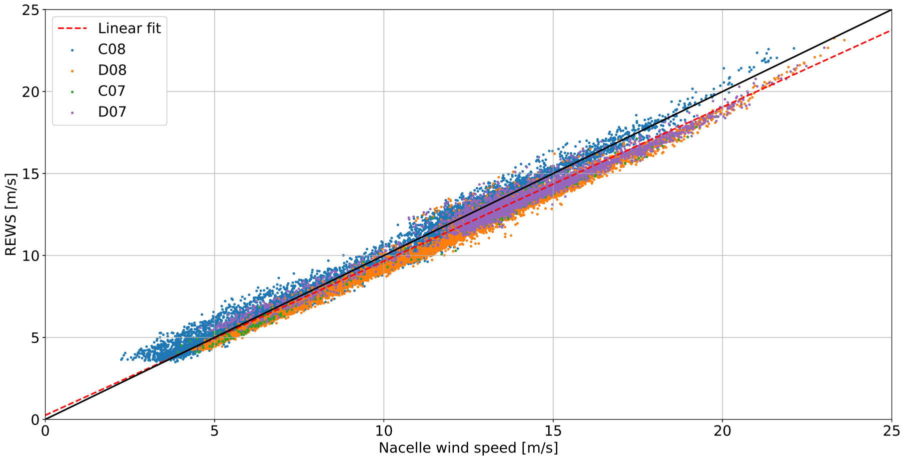

Figure 4.6: Scatter plot showing the derived 10-minute average U_{REWS} versus U_{nac} and a linear fit.

The resulting coefficient of determination R^2 score [30] and the coefficients corresponding to the fits (a-c) can be found in table 4.1. This table shows that the mean score of the ten folds is equal for the first and second-order but that the standard deviation is lower for the first-order fit. It is chosen to use the first-order fit to get the correction factor between U_{REWS} and U_{nac} because it is the simpler model, and the standard deviation is lower. When analyzing figure 4.6, it can be concluded that for low wind speeds, U_{REWS} is slightly larger than U_{nac} while for higher wind speeds, U_{REWS} is lower than U_{nac}. The reason that U_{REWS} is lower than U_{nac} is thus because U_{REWS} considers the wind speed experienced by the whole rotor, while U_{nac} is a point measurement.

When U_{REWS} is obtained for all known free stream wind directions, another wind rose can be obtained. The free stream wind rose can be found in figure 4.1b and compared to figure 4.1a. It can be seen that in general, the roses are very similar. Two things stand out. First, the peak at 10° in the rose of Mast-p1 is not visible in the free stream rose. When looking at the trend, it is plausible that it is an error on the Mast-p1 side. The second thing that stands out is the low frequency at $U_\infty = 285°$ for the free stream rose. It should be noted that the data on which the wind roses are based differ in size and that the time steps are also different. So the free stream rose corresponds to different measurement years than the Mast-p1 rose.

4.2.3. Turbulence intensity

The third and final step is to obtain I_∞ per time step. I is defined as equation (4.6), where σ_U is the standard deviation of the wind speed in m/s and μ_U the mean of the wind speed in m/s. Multiple options to estimate I_∞ have been considered.

$$I = \frac{\sigma_U}{\mu_U} \tag{4.6}$$

The first option was to use U_{nac} of which the 10-minute average standard deviation and mean values are known for every turbine and use a similar approach as used for the wind direction to estimate I_∞. However, as mentioned, the U_{nac} is measured behind the rotor, which adds turbulence. Therefore this approach would overestimate I_∞.

Another option considered was to find a similar correction factor for the standard deviation as is used for the 10-minute average mean value between U_{nac} and U_{REWS}, and obtain I_∞ in that way. This method turned out to underestimate I_∞ significantly when compared to the distribution shown in figure 4.2c.

Therefore it was chosen to use the final option, which is to make use of TI roses obtained from the Mast-p1 data. This does not result in a unique I_∞ per time step, but an estimate of I_∞ based on θ_∞, U_∞ and the season of the time step (winter, spring, summer or autumn). The reason why the season is taken into account will be explained later in this section. The outlier detection method described in section 3.3 is used to remove the invalid TI points, as can be seen in figure 4.7. Note that points above 23 m/s are seen as outliers, while they are not necessarily outliers. However, in this thesis, the TIs for high wind speeds ($\sim U_\infty$>17.5 m/s) are not relevant, since all turbines operate at rated power at all times.

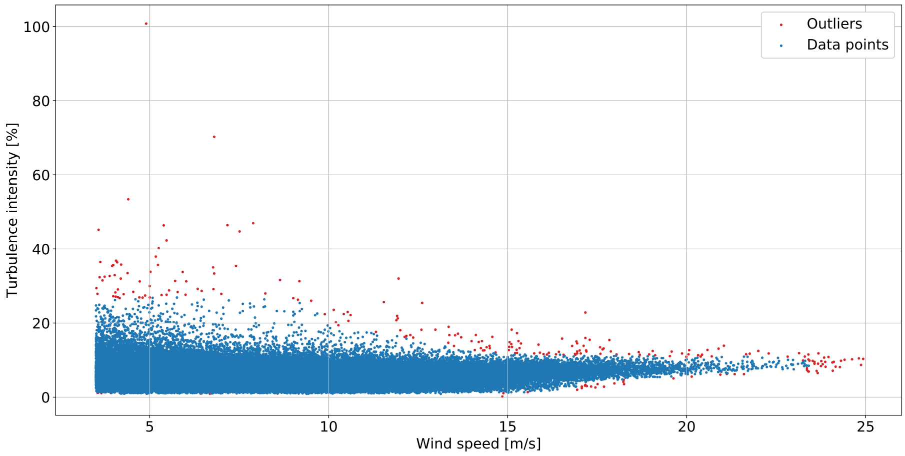

Figure 4.7: Scatter plot of the turbulence intensity versus free stream wind speed, including the observed outliers found using the local outlier factor method.

The TI data is first further analyzed. Multiple 2D distributions are shown in figure 4.8. It shows the frequency of occurrence in percentages for different combinations of variables. From figure 4.8a can be concluded that the most frequent occurring wind speed is around 8 m/s from around 270° as was also concluded from figure 4.2a. From figure 4.8b can be concluded that the highest I_∞ occur for lower wind speeds, that the average I_∞ is around 6% and that this average is slightly higher for higher wind speeds.

Finally, from figure 4.8c can be concluded that the average I_∞ is slightly higher for north-eastern and eastern winds, somewhat lower for southern winds and again slightly higher for western and north-western winds. This can be explained by the geographical location of the wind farm. The land that is closest to the farm is located to the east and north-west of the farm. Wind over land generally has a higher TI than wind over sea due to the roughness length of the surface [10]. The positions and distance from Lillgrund relative to land can be seen in figure 4.9.

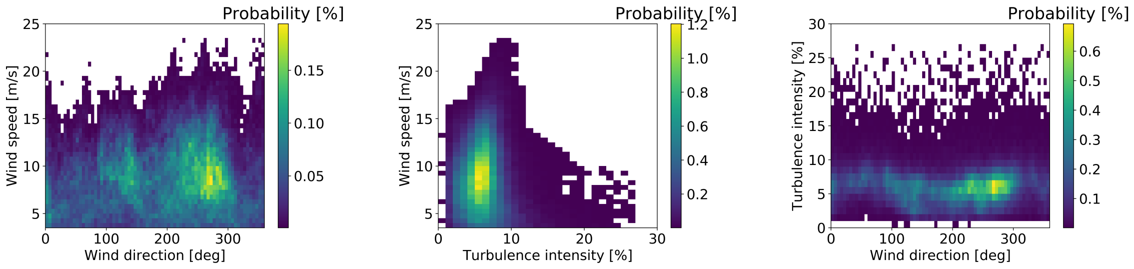

(a) Distribution of θ_∞ and U_∞ (b) Distribution of U_∞ and I_∞ (c) Distribution of I_∞ and θ_∞

Figure 4.8: Two-dimensional histograms showing the probability in percentage for different combinations of variables for the Mast-p1 data.

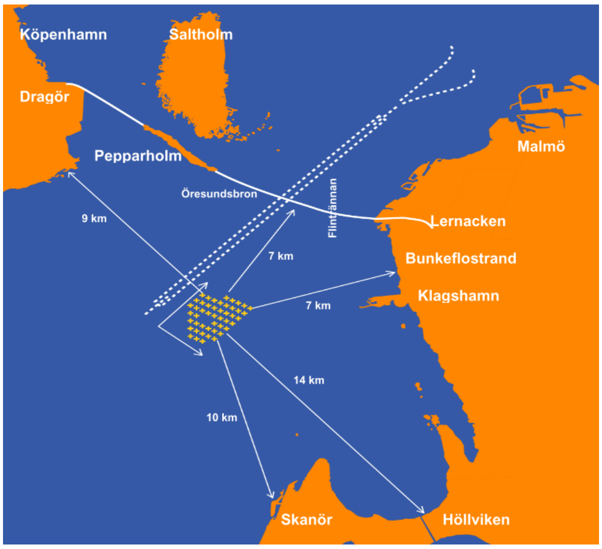

Figure 4.9: Location of the Lillgrund wind farm relative to land. Taken from Jeppsson et al. [64].

From the literature, it is known that two primary sources influence the TI. As just mentioned, one is the roughness length of the surface, but the heat flux between the surface and the air also influences the TI. When there is a lot of heat flux between the surface and the air above it, the TI is usually higher than average. When the temperature of the surface and the air is almost equal, the TI is usually lower than the average. Hence the TI depends on the daypart and whether that daypart is, for example, in the winter or the summer. Figure 4.10 is analyzed to estimate the influence of the daypart and season on the TI. Note that the data is binned per week and daypart for a total of 913 days of data. Also, note that the vertical dotted lines indicate the meteorological seasons (winter-spring-summer-autumn-winter).

When looking at the mean values of figure 4.10, it can be concluded that the TI difference between the dayparts is negligible. The largest absolute difference between the mean values occurs in the summer between night and evening and is 0.4%. When looking at the yearly mean, the difference is only 0.2%. However, a clear difference in mean values can be observed per season. This might be a consequence of winds frequently occurring from land in that particular season and because the temperature differences between the surface and air are larger in that season. The largest difference is between spring and winter and is around 1%. Therefore, the mapping of the TI on θ_∞ and U_∞ is done per season, resulting in four seasonal TI roses.

The mapping for the TI roses is done using a wind speed bin width of 0.5 m/s and a wind direction bin width of 5°. Two of the roses are shown in figure 4.11. One shows the TI spring rose, and one the TI total rose when all seasons are combined. The TI rose, which considers all data, is relevant for the calibration described in section 6.4. Note that, as expected, the spring TI rose shows, on average, lower values for I_∞ than the other rose. All roses are saved as look-up tables. If there is no data available for a certain combination of parameters, for example, for the combination spring, 200° and 20 m/s, the average I_∞ for 20 m/s is taken. If there are no

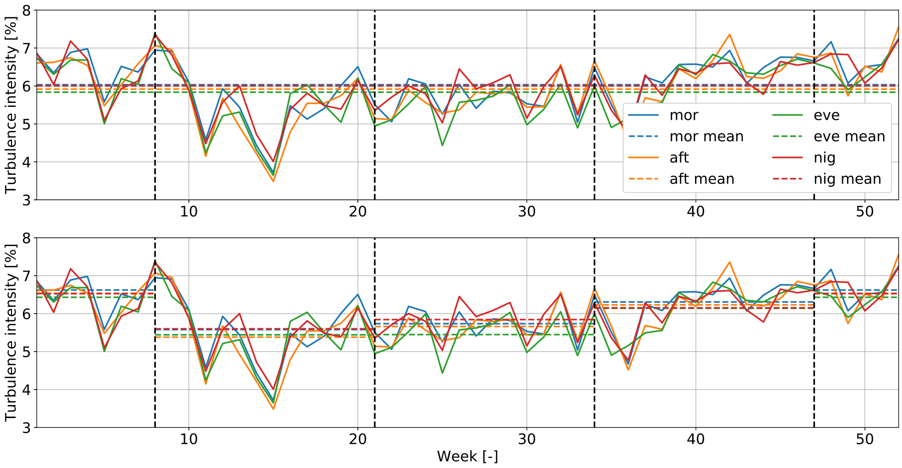

Figure 4.10: Average turbulence intensity binned per week per daypart, including the 1-year average (upper plot) and seasonal average (lower plot). The vertical dotted lines indicate the seasons (winter-spring-summer-autumn-winter).

data points for 20 m/s at all, $I_\infty = 6.0\%$ is used, which is the yearly average. However, as mentioned before, this only happens for high wind speeds, for which the TI is less relevant because all turbines operate at rated power already.

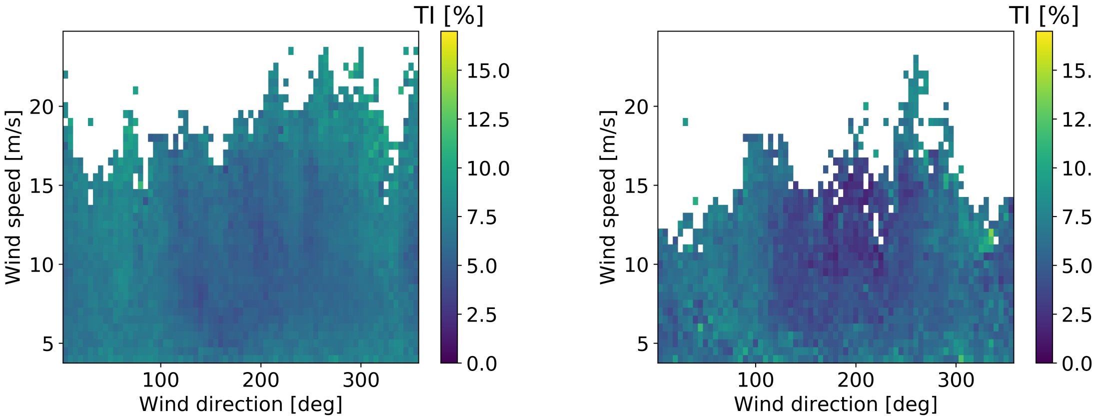

(a) Winter

(b) Spring

Figure 4.11: The turbulence intensity roses for all data and the separate spring data.

5

FLORIS model

This chapter elaborates on the FLORIS model [98], which is used in this thesis to estimate the power and energy gain related to wake steering for the Lillgrund wind farm. The FLORIS model, first presented by Gebraad et al. [47], is a data-driven parametric wind farm model that can predict the steady-state effects of yaw misalignment on the power output for a farm using relatively few model parameters. This chapter first describes the principles upon which the model is based in section 5.1. Then in section 5.2, the framework of the FLORIS model is briefly described and how it can be used in this research.

5.1. Principles

The FLORIS model has several surrogate models available. The original multi-zone wake model is based on the Jensen model [63, 68], augmented with a model for wake deflection through yaw. Further improvements have been included to better model situations with partial wake overlap and to capture wake position effects caused by rotational effects [49]. Through the years, additional improvements and surrogate models have been added to the FLORIS model.

In this thesis, a variant of the Gaussian wake model is used. This model combines a single wake model for wake redirection and turbine de-rating based on the paper by Bastankhah and Porté-Agel [8]. This surrogate model further uses a turbine induced turbulence model [22] and a turbulence summation model [96]. There are several reasons why the Gaussian wake model is selected over the other available models. First of all, it is shown that the Gaussian wake model performs well when the results are compared to experimental data from wind tunnel testing [8] and field testing [6]. Next to that, the Gaussian wake model has fewer tuning parameters compared to, for example, the multi-zone wake model [47]. Recently, the Gaussian wake model can be combined with a model that includes secondary wake steering based on the paper of King et al. [72]. When this sub-model is used, the model is called the Gauss-Curl Hybrid (GCH) model.

Four main principles are fundamental in the GCH wake model: the near-wake region, the 2D Gaussian shape of the velocity deficit in the far-wake region, the deflection of the wake centerline and the secondary wake effects.

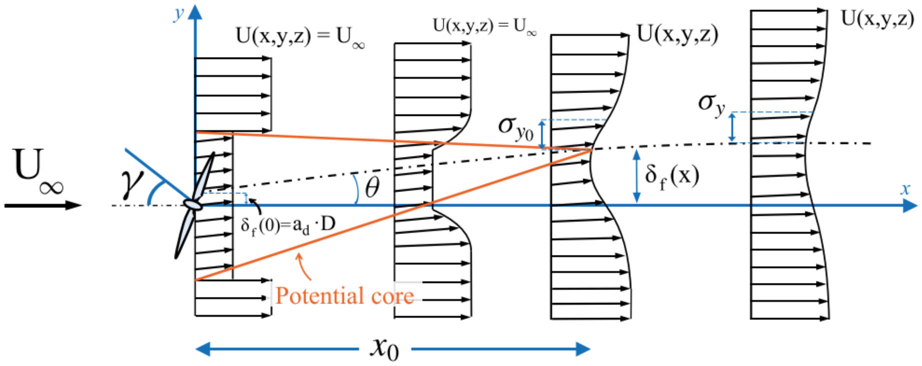

Figure 5.1: Schematic overview of the wake of a yawed turbine (a negative misalignment angle is shown). Originally taken from Bastankhah and Porté-Agel [8] and modified by Doekemeijer et al. [27].

5.1.1. Near wake

In the near-wake region, the velocity deficit does not yet follow a 2D Gaussian shape. The near-wake region is modeled as a linearly converging cone. The base of the cone is located at the rotor, and the tip is located at a distance x_0 downstream. This cone is visualized in figure 5.1 and x_0 can be determined through equation (5.1). In this equation, D is the rotor diameter in m, γ is the yaw misalignment angle with the incoming flow in rad, C_T is the thrust coefficient of the rotor, I_{rotor} is the turbulence intensity at the rotor of the turbine and α and β are the first two tuning parameters.

$$\frac{x_0}{D} = \frac{\cos(\gamma)\left(1+\sqrt{1-C_T}\right)}{\sqrt{2}\left[4\alpha I_{rotor} + 2\beta\left(1-\sqrt{1-C_T}\right)\right]} \tag{5.1}$$

5.1.2. Velocity deficit

The start of the far-wake region is at x_0 distance downstream from the turbine. It is assumed that from this point onwards, the wake velocity deficit has the shape of a 2D Gaussian distribution until the wake is fully recovered. The wake velocity deficit is expressed as in equation (5.2). In this equation, σ is the standard deviation in the specified direction and $(x,\ y,\ z)$ in the Euclidean space. The origin of this space is at the turbine hub, x is aligned with the wind direction, and z is defined positive upwards.

$$\frac{U(x,y,z)}{U_\infty} = 1-\left(1-\sqrt{1-\frac{\sigma_{y0}\sigma_{z0}}{\sigma_y\sigma_z}C_T}\right)\exp\left(\frac{\left(y-\delta_f\right)^2}{2\sigma_y^2}+\frac{z^2}{2\sigma_z^2}\right) \tag{5.2}$$

$$\delta_f = \delta_r(x) + \tan(\theta)x_0 + \frac{\theta}{5.2}\left(C_0^2 - 3e^{1/12}C_0 + 3e^{1/3}\right)\sqrt{\frac{\sigma_{y0}\sigma_{z0}}{k_y k_z C_T}}\ln\left[\frac{\left(1.6+\sqrt{C_T}\right)\left(1.6S_\sigma-\sqrt{C_T}\right)}{\left(1.6-\sqrt{C_T}\right)\left(1.6S_\sigma+\sqrt{C_T}\right)}\right] \tag{5.3}$$

$$\theta \approx \frac{0.3\gamma}{\cos(\gamma)}\left(1-\sqrt{1-C_T\cos(\gamma)}\right) \tag{5.4}$$

5.1.3. Wake deflection

The Gaussian shaped wake velocity deficit is centralized around the centerline of the wake. This centerline is displaced with a distance δ_f in the y-direction from the x-axis as indicated in figure 5.1. This displacement occurs due to yaw misalignment and wake rotation and can be calculated with equation (5.3). In equation (5.3), θ is the initial deflection angle and can be calculated using equation (5.4). Note that γ is in rad and that it is defined in FLORIS as positive counter-clockwise relative to the wind direction (so from top view), hence figure 5.1 illustrates a situation with negative yaw misalignment.

$$C_0 = 1-\sqrt{1-C_T} \tag{5.5}$$

$$S_\sigma = \sqrt{\frac{\sigma_y\sigma_z}{\sigma_{y0}\sigma_{z0}}} \tag{5.6}$$

$$\sigma_y = \sigma_{y0} + (x-x_0)k_y, \quad \text{with} \quad \sigma_{y0} = \frac{D}{2\sqrt{2}}\cos(\gamma) \tag{5.7}$$

$$\sigma_z = \sigma_{z0} + (x-x_0)k_z, \quad \text{with} \quad \sigma_{z0} = \frac{D}{2\sqrt{2}} \tag{5.8}$$

$$k_y = k_z = k_a I_{rotor} + k_b \tag{5.9}$$

$$\delta_r = a_d D + b_d x \tag{5.10}$$

Furthermore, C_0 is the velocity deficit at the start of the far-wake region and is defined as equation (5.5) and S_σ is defined as equation (5.6), while σ_y and σ_z are defined in equation (5.7) and equation (5.8) respectively. The coefficients k_y and k_z are linear wake expansion coefficients, similar to the ones Jensen [63] defines. These linear wake expansion coefficients are both a function of I_{rotor} as can be seen in equation (5.9). This equation contains the next two tuning parameters, namely k_a and k_b. Finally δ_r represents the wake deflection induced by the rotation of the wake [49]. It can be described by a linear function as shown in equation (5.10) and it gives the final two tuning parameters a_d and b_d. So in total the Gaussian wake model can be fitted using the following six tuning parameters: α, β, k_a, k_b, a_d and b_d. More information on the tuning parameters and how they are fitted can be found in section 6.4.

5.1.4. Secondary wake effects

A hybrid model is proposed by King et al. [72] to include the secondary wake effects. This hybrid model combines the Gaussian model described earlier, with analytic approximations made of the curl model presented in Martínez-Tossas et al. [90] and is therefore named the GCH model. The GCH model maintains all the advantages of the Gaussian model while incorporating corrections. These corrections are related to vortices in the wake and address the following discrepancies.

1. Vortices that drive yaw-based wake recovery.
2. The interaction of counter-rotating vortices with the atmospheric boundary layer, resulting in an asymmetry of the wake.
3. The inclusion of vortices makes it possible to model secondary steering and related multi-turbine effects.

$$\Gamma_{wr} = \frac{\pi(a-a^2)U_\infty D}{\lambda} \tag{5.11}$$

$$\Gamma_{cr}(\gamma) = \frac{\pi}{8}\rho D U_\infty C_T \sin(\gamma)\cos(\gamma^2) \tag{5.12}$$

Details of the GCH model can be found in the paper of King et al. [72], but the main principles and equations are described here. There are two types of vortices to be described. One is a result of the wake rotation and hence called the wake rotation vortex. The circulation Γ_{wr} of this vortex is calculated through equation (5.11). This equation uses λ of the turbine and the axial induction a, which is approximated through C_T.

The other type is called the counter-rotating vortex and is a result of yaw misalignment. There is one counter-rotating vortex at the top and one at the bottom of the rotor. The circulation of the counter-rotating vortex Γ_{cr} can be approximated with equation (5.12).

These circulation terms can be used to find the spanwise V and vertical W velocity components based on the wake rotation and yaw misalignment angle anywhere in the space domain. These velocity components influence the streamwise velocity and the wake deflection. The wake recovers more when the turbine operates in yaw misaligned conditions due to large-scale entrainment of flow into the wind farm domain. It is assumed that this is primarily an effect because of the W component. It should be noted that in the calculation of those velocity components, an added parameter is introduced ϵ representing the size of the vortex core. In this thesis, it is assumed that ϵ = 0.3D, as proposed by King et al. [72].

$$u(x,y,z) = u_G(x,y,z) + \frac{W(x,y,z)(x-x_0)(y-y_0)}{\pi\left(\alpha_r(x-x_0)+\frac{D}{2}\right)^2} \tag{5.13}$$

The extra wake recovery is defined as in equation (5.13). Note that u_G is equal to equation (5.2). Also, note that a new tuning parameter α_r is added that dictates how much the entrainment affects the wake recovery. For this thesis, it is assumed to be the same value as in King et al. [72], namely 0.03. Finally, x_0 and y_0 indicate the turbine location that produces the vortex.

On top of the extra wake recovery for yaw misaligned turbines, the vortices also make it possible to predict secondary steering. The vortices described propagate downstream, and when they reach a downstream turbine, the wake of this turbine is influenced by these vortices. This phenomenon is called secondary steering [38]. The V and W components generated by the counter-rotating vortices act like an effective yaw angle at the next turbine. This yaw angle is added on top of the actual yaw misalignment angle of that turbine, see equation (5.14). The presence of γ_{eff} generally has the effect that downstream turbines do not have to yaw as much as upstream turbines to produce energy gains.

$$\gamma = \gamma_{turb} + \gamma_{eff} \tag{5.14}$$

It is important to note something about the two extra tuning parameters introduced in this section: α_r and ϵ. The sub-model describing the secondary wake effects has been added to the FLORIS model in a relatively late stage of this thesis. Therefore, these two parameters are unfortunately not calibrated in section 6.4 due to limited time. So it should be kept in mind that since these two parameters are not calibrated, they introduce additional uncertainties to the output of the FLORIS model.

5.2. Framework

The FLORIS model is an open-source model that can be downloaded from Github for Python [98] or MATLAB [26]. The FLORIS model is designed with three main goals in mind:

- To model the most dominant wake characteristics, such as velocity deficit, position, and turbulence intensity.
- To find the optimal control setpoint for wind turbines to optimize the total wind farm power output.
- To model how wakes interact with each other.

The FLORIS model in this research is used to compare the power and energy production of the Lillgrund wind farm with and without wake steering. In order for FLORIS to estimate these power and energy productions, multiple inputs are required.

First of all, the FLORIS model needs as input for every turbine i its position (X_i, Y_i) in m, D in m, its hub height z_{hub} in m, its power coefficient and thrust coefficient curve and the power-yaw loss coefficient α, which is elaborated upon in sections 6.1 and 6.4.

Next to that, the characteristics of the free stream wind are needed. The FLORIS model needs the following inputs: U_∞ in m/s, θ_∞ in deg, I_∞, $\alpha_{p,\infty}$ and ξ_∞. As mentioned in section 4.1, there is not enough data to get a proper estimate for the inputs $\alpha_{p,\infty}$ and ξ_∞. Therefore ξ_∞ is neglected and set to zero. For $\alpha_{p,\infty} = 0.12$ is used, which is the standard input. It describes a vertical velocity profile calculated using the power-law [10].

Finally, the FLORIS model needs to choose one of the available options per sub-model. The sub-models are: the wake model, the velocity model, and the combination model. Depending on these sub-models, additional inputs are needed. Since the GCH wake model and the added turbulence model are used, the additional inputs are the tuning parameters introduced in section 5.1 and the turbulence intensity parameters modified from Crespo and Hernández [22]. All the inputs earlier are given to the FLORIS model through a .JSON file.

Once all the inputs are correctly in place, the FLORIS model calculates the flow field that includes the wakes based on the free stream wind and the yaw position of the wind turbines. This flow field is then used to estimate the power production per turbine. Further details on the calculations and use of the FLORIS model can be found in the documentation by NREL [98].

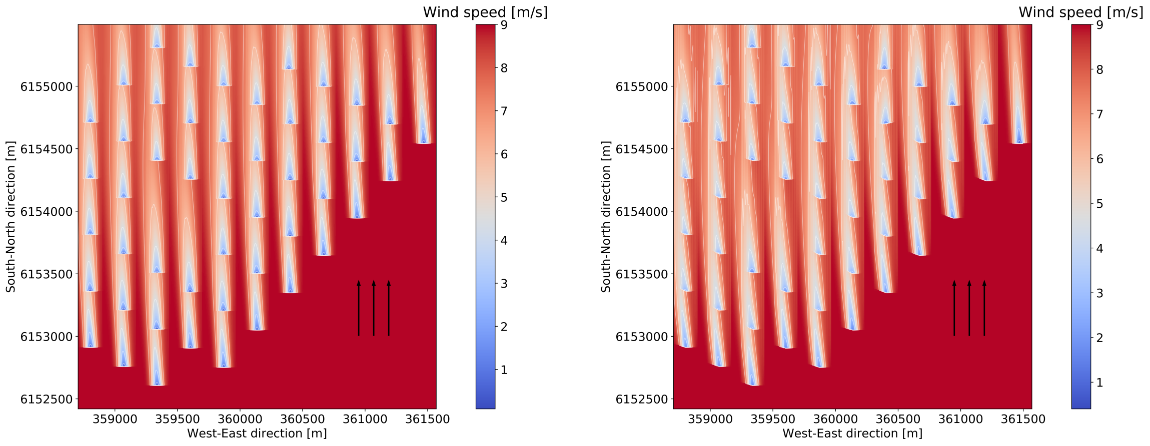

(a) Without yaw optimization

(b) With yaw optimization

Figure 5.2: A visualization of the flow field in the Lillgrund farm for a free stream wind speed of 9 m/s and wind direction of 180° (indicated by black arrows). Both an optimized and normal yaw case are shown.

The FLORIS model comes with additional tools. The most important tool for this research is the tool that calculates the optimal yaw position per turbine to maximize the power output of the entire farm. Other tools include, among others, the visualization of the flow field, the extraction of the flow field, and the visualization of a power rose. A visualization of the flow field in the Lillgrund farm for a $U_\infty = 9$ m/s and $\theta_\infty = 180°$ can be seen in figure 5.2. Figure 5.2a shows the flow field if no yaw optimization is performed and figure 5.2b shows the flow field if the turbines operate at their optimal control set point. Note that the wakes already deflect slightly, even when the turbines are aligned with the wind direction.

It should be emphasized that the FLORIS model assumes steady wakes. This is a significant simplification of reality since the wake of a turbine is actually dynamic and unsteady. This is one of the reasons why the FLORIS model is classified as a low-fidelity model. It cannot capture more complex phenomena in the wake because of the steady-state assumption. The accuracy of the FLORIS output is, therefore, lower than what can be achieved with higher-fidelity models such as, e.g., EllipSys3D or SOWFA, which do take into account the dynamic wake behavior. However, the computation time of the FLORIS model is much lower, making it ideal for the yaw optimization performed in chapter 8.

6
Calibration

The different surrogate models available in FLORIS are not off-the-shelf products that can be used for every wind farm without calibration. FLORIS relies on assumptions and simplifications of the wake physics, for example, that it can be described as a steady-state phenomenon, while in reality, the wake is dynamic and unsteady. Therefore, the surrogate models are topology and turbine dependent and need to be calibrated accordingly.

The calibration process used for the GCH wake model will be explained in this chapter. First, the relevant inputs for the calibration are explained in section 6.1. In section 6.2 is explained how the measured data on which the model is calibrated is obtained and verified. Then in section 6.3 is shown how the first model input, the power curve, is obtained. After that, it is shown how the model is fitted on the measured data by calibrating the tuning parameters in section 6.4. Finally, the obtained power curve and calibrated tuning parameters are validated in section 6.6.

6.1. Inputs

There are multiple relevant inputs in the FLORIS model that change per turbine and wind farm topology and hence need to be calibrated. These model inputs have already been introduced in chapter 4, but are summarized and listed here for convenience.

1. The position per turbine (X_i,Y_i)
2. The turbine characteristics (z_{hub}, D, $C_P - V$ etc.)
3. The free stream wind characteristics (U_∞,θ_∞, I_∞, etc.)
4. The selected surrogate model (Jensen model, Multi-zone wake model, GCH wake model, etc.).
5. The tuning parameters corresponding to the chosen surrogate model (α, β, k_a, etc.)
6. The added turbulence characteristics.

Most of these inputs are given to the FLORIS model through a .JSON file. Note that there are more input options in the .JSON input file than there are listed here, for example, the pitch angle, but these are not used (yet) in the GCH wake model. Two other input parameters that are being used are the generator efficiency and the tilt angle. For this thesis, both parameters are implicitly defined in the power curve. Hence the generator efficiency is set to one, while the tilt angle is set to zero.

Inputs 1, 4, 6, and some parts of 2 are given or chosen and hence do not require any calibration. Input 3 can be found through data analysis, as explained in section 4.2. That leaves input 5 and parts of input 2 that need to be calibrated. To be more specific, the power curve, the tuning parameters, and the power-yaw loss coefficient need to be calibrated.

A power curve is turbine dependent, although the general behavior is always the same (see figure 6.2a). Below rated wind speed, the power curve increases cubed due to the U^3 component (equation (4.2)). Then at rated wind speed, the rated power is reached and kept constant for wind speeds above rated wind speed. The turbine type in the Lillgrund farm is the Siemens SWT-2.3-93, with a rated power of 2.3 MW. For this thesis, the power curve from the manufacturer is provided by DTU.

However, the theoretical manufacturer power curve and the 'real' power curve, the behavior of the power curve in practice, often differ. This mainly happens when the wind shear and TI are different from the design conditions [21, 82]. A poorly calibrated power curve introduces an error and increases the uncertainty of the analysis. This makes this input more complicated then one might expect. Details on the power curve calibration can be found in section 6.3.

The thrust curve is also a relevant input of the FLORIS model that would ideally be calibrated. Unfortunately, a proper method with the data currently available is unknown. Therefore, the theoretical thrust curve provided by the manufacturer through DTU is assumed to be correct. This assumption again introduces additional uncertainty in the model output.

The tuning parameters definitely need to be calibrated. The calibration process of the tuning parameters is one of the main steps and the biggest challenge to make FLORIS usable for this thesis. Calibrating tuning parameters is already challenging when 'true' flow field data is available [27, 29, 49], for example, obtained through SOWFA or EllipSys3D. Let alone when there is no flow field data to calibrate the tuning parameters on. There is no flow field data from high-fidelity models available for the Lillgrund wind farm since it would take too long to produce this data for all the inflow conditions that will be considered in this thesis. The calibration process of the parameters is described in detail in section 6.4.

Also, the turbulence characteristics can be adjusted. These characteristics could be treated as tuning parameters and included in the calibration process. However, it is desired to keep the tuning parameters as low as possible, to avoid over-fitting and too long computation times while calibrating. Hence, it is assumed that the default turbulence characteristics in the FLORIS model describe the added turbulence well. This assumption adds uncertainty to the analysis.

$$P_\gamma = P_0 \cos^\alpha(\gamma) \tag{6.1}$$

The power-yaw loss coefficient α is a parameter that determines how much power is lost when a turbine has a yaw angle with the incoming flow through equation (6.1) [84]. P_0 is the power in W that the turbine would produce if γ=0 rad, hence when it is perfectly aligned with the inflow wind direction. In the paper by Liew et al. [84] is shown that α is different for fully waked turbines compared to free standing turbines. The value of α is the largest for turbines operating in the near-wake region and converges to the free stream value for increasing downstream distance (see figure 6.9). In other words, α depends on the turbine spacing. More details on the calibration of α can be found in section 6.5.

6.2. SCADA data

To be able to calibrate a model, reference data is needed. This data is considered to be the true data, and the goal is to calibrate the FLORIS model such that its output is as close to this data as possible. Hence it is essential that the true data is reliable and adequately derived through data analysis.

As mentioned before, only power measurements can be used to calibrate the model. Two different data sets of reliable power measurements need to be obtained: one containing power measurements of free standing turbines **S1** and one containing power measurements of aligned rows of turbines **S2**.

6.2.1. Free standing turbines

S1 contains power measurements for free standing turbines and is used to obtain the real power curve. This process is described in section 6.3. Three sets of free standing turbine powers are obtained. All three sets are first used to obtain the real power curve and then tested on each set individually in section 6.6 to validate the curve.

S1.1 contains all power measurements for the most western turbine column C1 (see figure 3.1), corresponding to western winds. To be more exact, the data is filtered for θ_∞=270°±20°. To ensure that only valid power measurements are used for the calibration, the data is filtered, and outliers are removed, as shown in figure 3.2. The resulting data set contains 166,149 usable data points for five turbines.

S1.2 is derived similarly, though this time, the most southern turbines D-08 till A-07 are used, and the data is filtered for θ_∞=180°±20°. The data set is filtered, and outliers are removed as always, resulting in 62,871 usable data points for five turbines. This number is significantly lower than for **S1.1**, because southern winds occur less frequent than western winds (see figure 4.1b).

S1.3 contains all power measurements for turbines A-02 till A-06 with winds from around 120°. So the data is filtered for θ_∞=120°±20°. The resulting data set contains 50,015 usable data points for five turbines. This number of data points is again significantly lower than for **S1.1**, because winds from 120° occur less frequent than western winds.

The reason that northern and north-eastern winds are not used is that no consistent data could be obtained for those free standing turbines.

6.2.2. Aligned turbines

S2 contains power measurements for turbines that are aligned with the inflow direction and is used to obtain the tuning parameters. This process is described in section 6.4. In total, six different spacings between the turbines are relevant. In this case, relevant means that changing the tuning parameters for that specific spacing results in additional accuracy. Each spacing has two wind directions with which they are aligned. For example the aligned turbines indicated with C1-C10 in figure 3.1 have a spacing of 4.8D and are aligned with $\theta_{\infty,1} = 0°$ and $\theta_{\infty,2} = 180°$. The six different spacings are visualized in figure 6.1 and listed in table 6.1. Note that the data set names correspond to the underlined alignment directions. Certain alignment directions (e.g., 42° and 90°) are not used because there are not enough data points for these directions to make a calibration or validation possible.

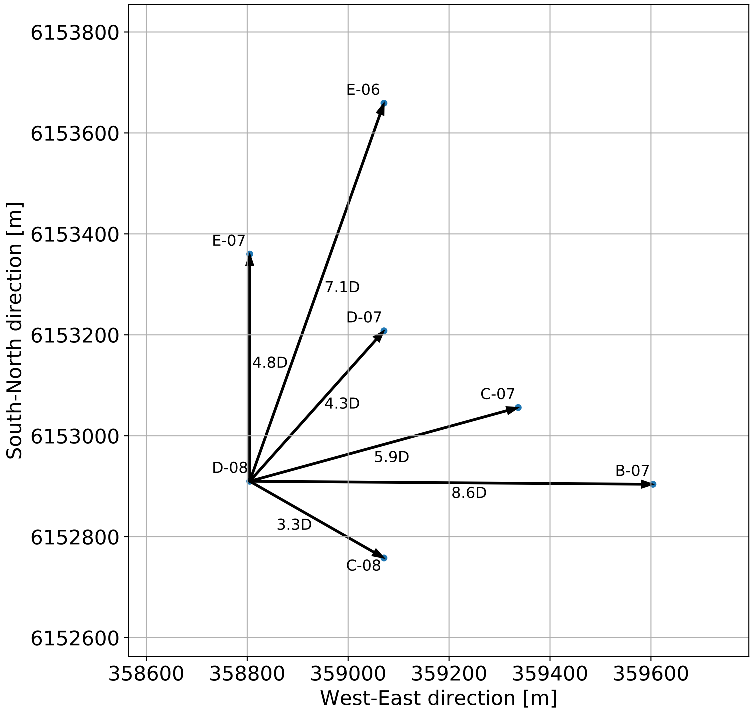

Figure 6.1: Part of the Lillgrund wind farm layout, including the different spacings.

Table 6.1: Data set name per spacing and alignment direction including the number of data points. Note that the underlined directions correspond to the data set name.

Data set [-]	Spacing [D]	$\theta_{\infty,1}$ [deg]	#datapoints [-]	$\theta_{\infty,2}$ [deg]	#datapoints [-]
S2.1 & S2.2	3.3	$\underline{300}$	560,403	$\underline{120}$	448,493
S2.3	4.3	$\underline{222}$	597,033	42	-
S2.4 & S2.5	4.8	$\underline{180}$	311,288	$\underline{0}$	261,067
S2.6 & 2.7	5.9	$\underline{255}$	833,911	$\underline{75}$	326,488
S2.8	7.1	336	-	$\underline{156}$	326,205
S2.9	8.6	$\underline{270}$	927,738	90	-

S2.1-S2.9 are obtained in the same way as **S1**, except that the wind direction bin width is now only 20° instead of 40°. This is because in section 6.4, only power measurements corresponding to a direction close to the alignment direction are relevant.

6.3. Power curve analysis

The power curve provided to FLORIS is crucial in getting a realistic output. Once FLORIS has accurately calculated the local wind speed for every turbine, it should also accurately calculate the power, which is done using the power curve. The power curve of the manufacturer is provided and can be seen in figure 6.2. It is found that this power curve does not match well with the SCADA data. A reason for this can be that the turbines operate under different conditions and at a different site than the design conditions and site used by the manufacturer to obtain the power curve. For example, the turbines at Lillgrund have a hub height of 65 m while the Siemens SWT-2.3-93 is designed for hub heights between 80 and 101 m [64]. An additional reason might be that the manufacturer gives a conservative version of the power curve to make sure not to overstate the turbine's performance potential.

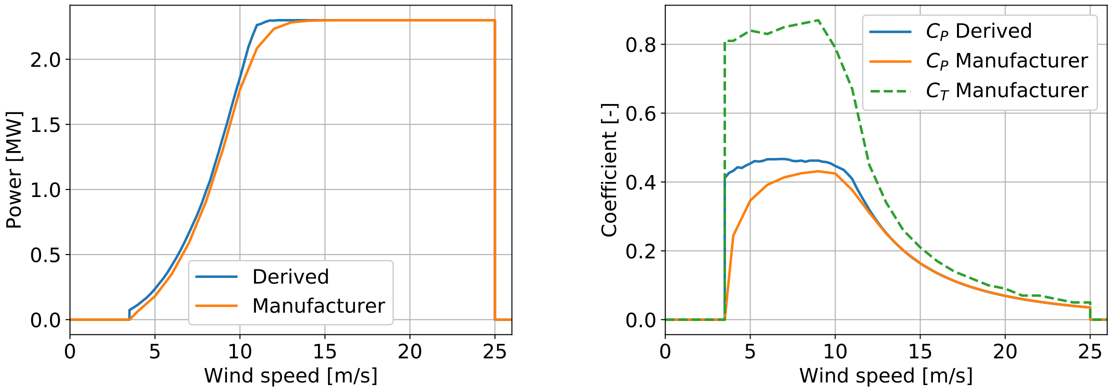

(a) Power curves

(b) Coefficient curves

Figure 6.2: The relevant power curves for the Siemens SWT-2.3-93 turbine and the corresponding power and thrust coefficients.

Since the power curve for the manufacturer is not accurate enough, a new power curve is derived from the SCADA data. Data set **S1** is used to derive the power curve. The following data is known per time step t in the data sets: $P_{t,i}$, $\theta_{t,\infty}$, $U_{t,\infty}$. $P_{t,i}$ is the power per operating turbine in W, so the subscript i indicates the turbine number. Sometimes some $P_{t,i}$ for a certain t are missing. That can happen, for example, when the turbines are not operating because of maintenance, a faulty sensor, or a data point that is removed during the outlier removal.

First is checked if all turbines are operating for a certain t. The remaining data is then filtered for $3.5 \leq U_{t,\infty} < 13.5$ m/s, since below 3.5 m/s the turbines do not operate and above 13.5 m/s all the free stream turbines operate at rated power. After filtering, there are a total of 34,688 t, each having five operating turbines (either turbines in C1, turbines D08 till A-07 or turbines A-02 till A-06). These five $P_{t,i}$ are summed for every time step resulting in 34,688 P_t that can be used to obtain the real power curve.

The method to obtain the real power curve is iterative and performed manually. FLORIS is run for all t and the predicted power $P_{t,F}$ is compared to P_t. Using the power curve of the manufacturer as starting point, the C_P values per wind speed are adjusted in the .JSON input file until a satisfactory fit is achieved. Note that this an inefficient process and that the use of an optimizer could both improve the accuracy of the obtained power curve and potentially decrease the calibration time. However, due to time constraints, an optimization script has not been created.

The C_P values are adjusted based on the Mean Percentage Error (MPE) and Mean Absolute Percentage Error (MAPE) per wind speed, as shown in figure 6.3. The goal is to minimize the MAPE while maintaining the MPE around 0%. Note that the error is calculated between the red output of FLORIS and the blue data points. Also, note that the red output of FLORIS is slightly lower than the actual power curve. This happens because FLORIS uses the REWS to calculate the power, which is slightly lower than the wind speed at hub height due to wind shear effects.

Figure 6.3 shows the final result, and a few observations can be made. First of all, it can be seen that the spread of SCADA data increases with wind speed up until rated wind speed. The observed spread is likely

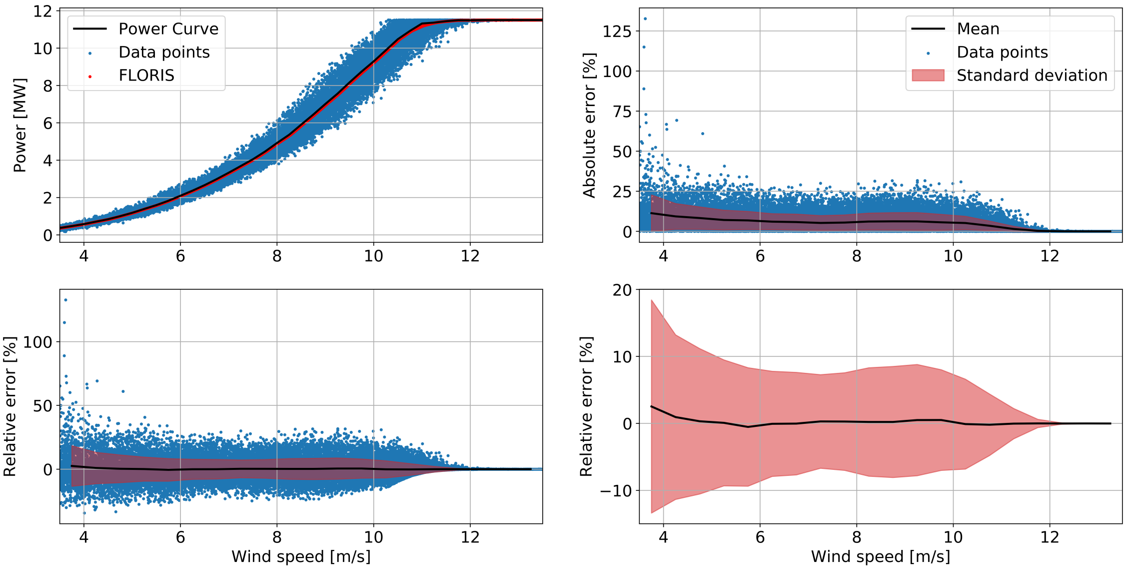

Figure 6.3: FLORIS output for free standing turbines compared to SCADA data (**S1**) including the relative and absolute error. The relative error is plotted once more for clarity.

a result of, amongst others, the TI, wind shear, and atmospheric stability, which are not constant. It can for example be seen in figure 4.8c that in general the TI is lower for $\theta_\infty = 120°$ than for $\theta_\infty = 270°$. This might be explained by how close the land is from those directions and how long the wind has been going over land. Since the roughness length is higher for land than sea, the wind shear is also influenced, which influences the power production of the turbines [10].

This spread inevitably results in errors when FLORIS is compared to the SCADA data. These errors are visualized as well. The mean and the standard deviation of both the relative and absolute errors are shown. The MPE fluctuates around 0% as intended, and the MAPE decreases with increasing wind speed and converges to 0% above rated wind speed. The decrease in MAPE indicates that the increase in power is larger than the increase in the spread.

Furthermore, when looking at the behavior of the relative error per wind speed, it looks like this error is normally distributed. To see whether the relative error is normally distributed, slices at different wind speeds are taken and plotted. An example can be seen in figure 6.4a. Similar plots are made for other wind speeds up until 13 m/s, and similar behavior is observed. Hence, it can be concluded that the relative error is normally distributed.

For $3.5 \leq U_{t,\infty} < 13.5$ m/s the MPE is 0.19% and the MAPE is 5.25%. However, to ensure that the new real power curve performs well, the curve should be validated, which is done in section 6.6.

The final power (coefficient) and thrust coefficient curves can be found in figure 6.2. It can be seen that the power curve of the manufacturer underestimates the power produced per wind speed for all wind speeds below rated. Note that the thrust coefficient curve used is the one provided by the manufacturer.

6.4. Tune parameter calibration

The six tuning parameters that the GCH wake model needs are vital in the calibration process. These tuning parameters are introduced and explained in section 5.1. Note that there are two more parameters, α_r and ϵ, that are not considered in this calibration, due to them being added to the model in a late stage of this thesis.

So for this part of the calibration, the focus is on α, β, k_a, k_b, a_d and b_d. As mentioned, there is no flow data available, hence reproducing a tuning parameter calibration as performed by Doekemeijer et al. [27, 29], Fleming et al. [36] or Gebraad et al. [49] is not possible. Instead, a new method is introduced that tries to minimize the difference between FLORIS output and power measurements by adjusting the tuning parameters. To be more exact, the MAPE between the FLORIS output and the power measurements is minimized. The optimization is defined in equation (6.2), where $\Psi = [\alpha\ \beta\ k_a\ k_b\ a_d\ b_d]$ and $P^i = [P^i_1\ P^i_2 \dots P^i_n]$

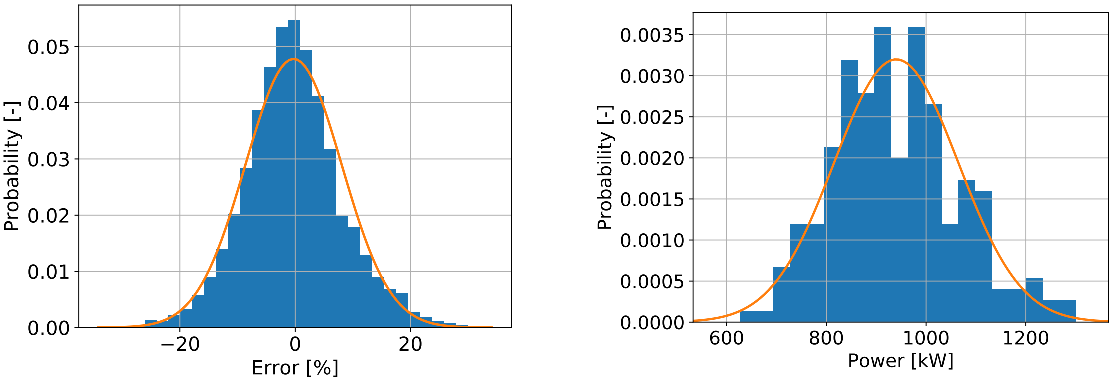

(a) Relative error distribution

(b) Power measurement distribution

Figure 6.4: Histograms of the relative errors occurring in the 6 ± 0.5 m/s bin and power measurements in the 8 ± 0.5 m/s bin for turbine C-08. Both with a fitted normal distribution.

where n is the number of turbines considered and i corresponds to a combination of U_∞ and θ_∞ and hence also turbine spacing. Later in this section is shown that the tuning parameters are indeed wind speed and spacing dependent.

$$\Psi^i_{opt} = \underset{\Psi^i}{\operatorname{argmin}} \frac{1}{n} \sum \left| \frac{P^i_{SCADA} - P^i_{FLORIS}(\Psi)}{P^i_{SCADA}} \right| \tag{6.2}$$

6.4.1. Power measurements

In section 6.2 is explained how the sub data sets in **S2** containing the power measurements for this part of the calibration are derived. The power measurements are connected to free stream wind conditions, as explained in section 4.2. Here, it is also shown that these free stream wind conditions are not always a perfect match with reality. Therefore, the subsets in **S2** need further filtering before it can be used as P^i_{SCADA} in equation (6.2).

A function is written that returns the filtered P^i_{SCADA} based on a wind speed and wind direction range, for example $U_\infty = 8 \pm 0.5$ m/s and $\theta_\infty = 180 \pm 5°$. While using a wind speed bin width of 1 m/s is common practice, the wind direction bin width is not [44]. In section 6.6 is shown that a bin width of 10° is appropriate to use. Another input of the function are the relevant turbine columns/rows. In this example case C2 and C5-C8 are used.

The function filters as follows. First of all, the right sub data set is chosen corresponding to θ_∞, so in this case, **S2.4** is used. After that, the data set is filtered for the wind speed range. Finally, it is checked per time step, whether all turbines in one row/column have a valid power measurement.

The remaining power measurements are averaged to obtain the mean power measurement per turbine $P^i_{SCADA,mean}$. Next to that is the standard deviation of the power measurements obtained to find the 68% Confidence Interval (CI). Note that this is only valid if the SCADA data is normally distributed. In figure 6.4b is shown for one wind speed and turbine that the powers are close to normally distributed, and hence that the standard deviation can be used. It is found that this is valid for all turbines and powers below rated wind speed. An example of $P^i_{SCADA,mean}$ and the corresponding $P^i_{SCADA,68}$ per wind speed can be found in figure 6.5. $P^i_{SCADA,68}$ indicates one standard deviation and hence the range in which 68% of the data points fall within.

6.4.2. Minimizer

Equation (6.2) is applied in Python using an optimization tool. Finding the ideal optimization technique in terms of speed and accuracy for this problem could be a research on its own. It is, therefore, considered to be outside of the scope of this thesis. Two optimizers have been used. Both aim to minimize the cost function defined in equation (6.2) and are part of the commonly used `SciPy` package. The first one that is used during the development of the calibration method uses Sequential Least SQuares Programming (SLSQP) [73]. It solves the optimization problem relatively fast but is more likely to find a local minimum than the other optimization tool. This other optimization tool is called Differential Evolution [121]. It is a lot slower than the SLSQP algorithm, but it is able to find the global minimum of a multivariate function. The Differential Evolution tool is, therefore, eventually used to find the tuning parameters per wind speed and spacing.

Table 6.2: The boundaries per parameter defined for the optimizers.

Parameter	Lower bound	Upper bound
α	0.125	2.5
β	0.015	0.3
k_a	0.05	1.5
k_b	0.0	0.02
a_d	-1.0	1.0
b_d	-0.1	0.1

Both minimizers require boundaries per parameter in between which it should search for the optimum. These boundaries are obtained from Doekemeijer et al. [29], who based them on the values of the parameters found in Bastankhah and Porté-Agel [8]. The boundaries are adjusted to match for the Python version of FLORIS. Only the lower boundary of k_b is changed from -0.01 to 0. The eventually used boundaries can be found in table 6.2. The SLSQP minimizer requires besides defined boundaries, also an initial guess per parameter. For the first wind speed, considered the initial guess is set to the mean value between the upper and lower boundary. For the rest of the wind speeds, the optimal parameters corresponding to the last obtained wind speed are used as the initial guess.

6.4.3. Calibration results

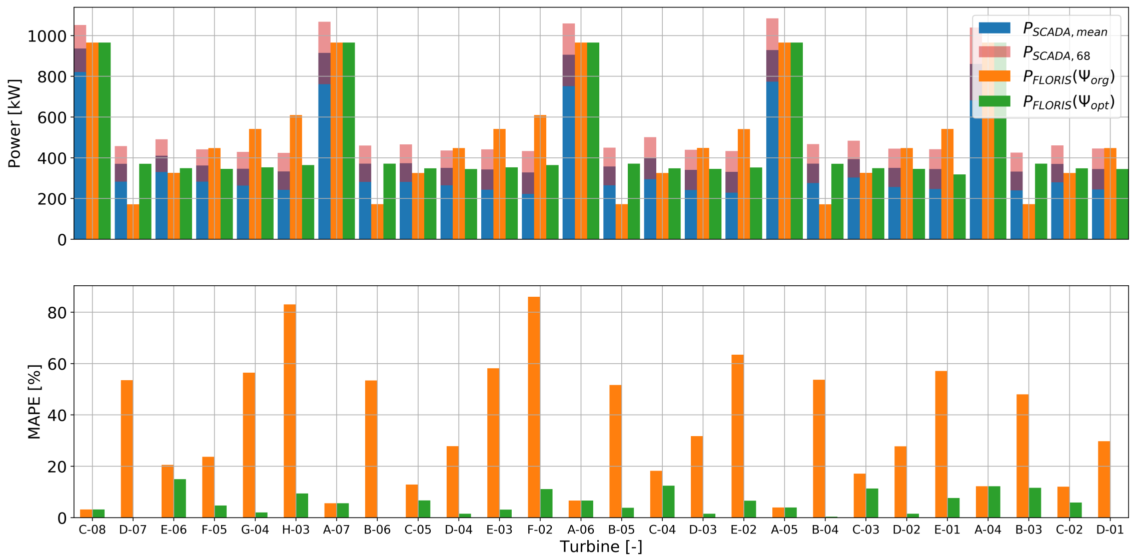

Figure 6.5: Bar plots of P_{SCADA}, $P_{FLORIS}(\Psi_{org})$ and $P_{FLORIS}(\Psi_{opt})$ including the MAPE for $U_\infty = 8$ m/s and $\theta_\infty = 180°$.

With the described calibration method in place, the first results can be obtained. An example of the power performance of FLORIS before and after the optimization compared to the SCADA data for $U_\infty = 8$ m/s and $\theta_\infty = 180°$ can be seen in figure 6.5. This plot shows the output of FLORIS for the original tuning parameters Ψ_{org} and for Ψ_{opt}, including the MAPE per turbine. Note that Ψ_{org} corresponds to the tuning parameters as defined in the example_input.JSON file from NREL [98]. It can be seen that Ψ_{opt} performs much better than Ψ_{org} and that all predictions of FLORIS using Ψ_{opt} fall within one standard deviation of the SCADA data. Note that for all but the last row, the highest MAPE is observed for the second downstream turbine. The average MAPE of all turbines for this wind speed is minimized to 5.7%.

However, when Ψ^A_{opt} is used for another wind speed, the average MAPE is not as low as it can be with a new set of tuning parameters Ψ^B_{opt}. Note that in this case A represents the combination $U_\infty = 8$ m/s and $\theta_\infty = 180°$ and B represents for example $U_\infty = 10$ m/s and $\theta_\infty = 180°$. Hence, the tuning parameters are wind speed dependent. Therefore Ψ^i_{opt} is obtained for every wind speed independently. An example of the trend that the tuning parameters follow per wind speed can be seen in figure 6.6. Note that the tuning parameters are only obtained for wind speeds up to 17.5 m/s, since for higher wind speeds, even turbines in the rows with the smallest turbine spacing, all operate at rated power. Also, note that a_d is manually set to zero. This is done because this parameter showed very chaotic behavior. On top of that is shown in section 7.5 that the impact of a_d on the variability of the output is negligible.

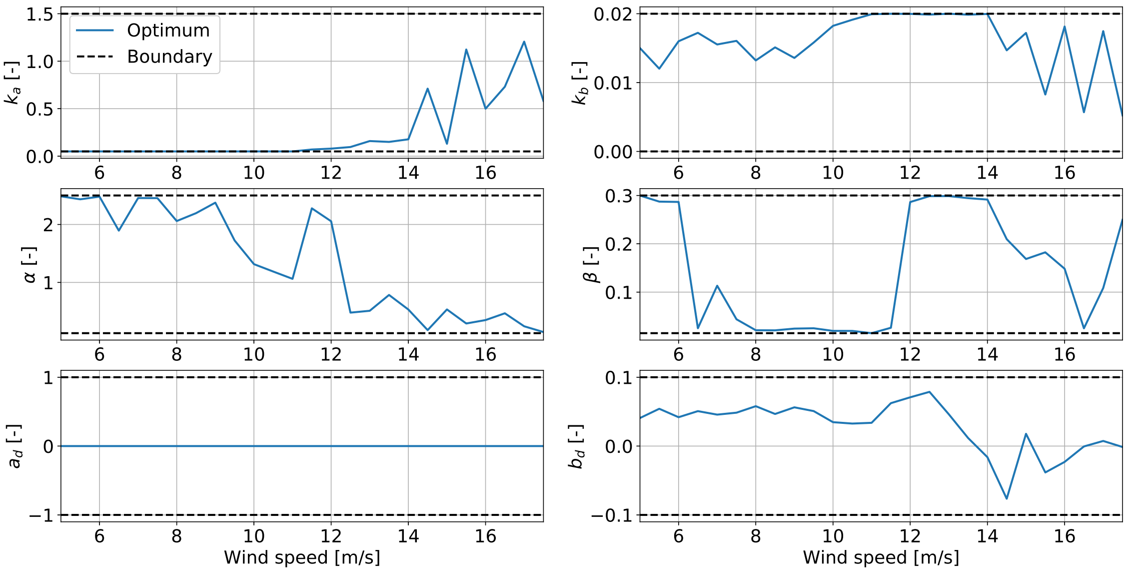

Figure 6.6: This figure shows the tuning parameters per wind speed for a spacing of 4.8D and $\theta_\infty = 180°$.

In figure 6.6 can be seen that for wind speeds up to 15 m/s, almost the full domain is used for every parameter, except for k_a and k_b. It can also be seen that k_a is often close to its lower boundary and that k_b is closer to the upper boundary. However, it should be noted that the sensitivity of k_b is relatively low, as discussed in section 7.5, hence k_a is still the dominating parameter for the wake expansion. Note that in the case shown, all turbines in the farm already operate at rated power from about 15 m/s and higher and that the MAPE, therefore, is almost 0% (see figure 6.7a). Therefore, in this particular case, the optimization of the parameters above 15 m/s is nonsensical, and hence the behavior of this part of the parameters is not analyzed.

The MPE and MAPE per wind speed can be found in figure 6.7a. Note that the wind speed range is from 5-16 m/s. This wind speed range differs per turbine spacing but is chosen to be from the first wind speed that all turbines in the farm are operational till the wind speed that all turbines in the farm are operating at rated power. This is done to get the most relevant average MAPE per spacing.

It can be seen that the MPE fluctuates around 0.5% and that the MAPE decreases with wind speed. This is because the power produced increases with wind speed while the absolute error in power does not change much, hence the error percentage decreases and converges to zero for wind speeds for which all turbines operate at rated power. The average MAPE for the wind speeds considered is minimized to 4.4% with a standard deviation of 2.9%.

In section 6.6 is shown that the tuning parameters for $\theta_\infty = 180°$ and a spacing of 4.8D can not be used to accurately simulate the situation for, e.g., $\theta_\infty = 300°$ and a spacing of 3.3D. Therefore, the tuning parameters are obtained per wind speed for all six turbine spacings that can be found in table 6.1. In section 6.6 is validated that the tuning parameters found using $\theta_{\infty,1}$ also result in a good estimation for $\theta_{\infty,2}$ for the same turbine spacing.

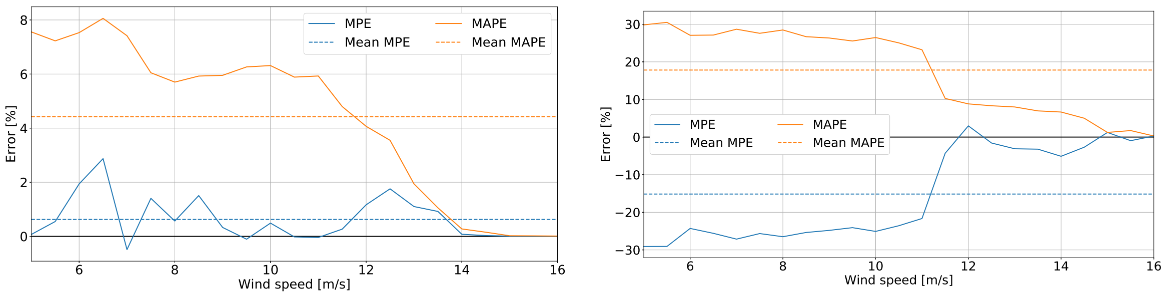

(a) $\theta_\infty = 180°$ on $\theta_\infty = 180°$

(b) 4.8D on 3.3D

Figure 6.7: The errors when tuning parameters corresponding to $\theta_\infty = 180°$ and 4.8D spacing are used on itself and when tuning parameters corresponding to $\theta_\infty = 180°$ and 4.8D spacing are used for $\theta_\infty = 300°$ and 3.3D spacing. Note that the x- and y-axis are different per figure.

So in total, six tuning parameters maps like the one shown in figure 6.6 are obtained. From now on, when FLORIS is used, the tuning parameters are first adjusted based on the spacing between the turbines in a row or column and U_∞. The parameter a_d is set to zero for all speeds and spacings. The average errors and the standard deviation of the errors per wind speed and spacing can be found in figure 6.8. It can be seen that the highest MAPEs occur for low turbine spacings and the lowest MAPEs for higher turbine spacings, this is because the FLORIS model is better able to model the far-wake region than the near-wake region.

Furthermore, it can be seen that the behavior of the MAPE is the same for every turbine spacing. It is generally the highest for low wind speeds, it decreases slightly with increasing wind speed, and it converges to zero as all turbines in the row start to operate at rated power. The standard deviation follows the same trend.

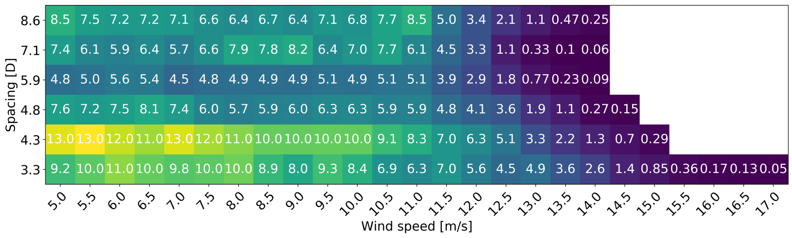

Spacing [D]	5.0	5.5	6.0	6.5	7.0	7.5	8.0	8.5	9.0	9.5	10.0	10.5	11.0	11.5	12.0	12.5	13.0	13.5	14.0	14.5	15.0	15.5	16.0	16.5	17.0
8.6	8.5	7.5	7.2	7.2	7.1	6.6	6.4	6.7	6.4	7.1	6.8	7.7	8.5	5.0	3.4	2.1	1.1	0.47	0.25						
7.1	7.4	6.1	5.9	6.4	5.7	6.6	7.9	7.8	8.2	6.4	7.0	7.7	6.1	4.5	3.3	1.1	0.33	0.1	0.06						
5.9	4.8	5.0	5.6	5.4	4.5	4.8	4.9	4.9	4.9	5.1	4.9	5.1	5.1	3.9	2.9	1.8	0.77	0.23	0.09						
4.8	7.6	7.2	7.5	8.1	7.4	6.0	5.7	5.9	6.0	6.3	6.3	5.9	5.9	4.8	4.1	3.6	1.9	1.1	0.27	0.15					
4.3	13.0	13.0	12.0	11.0	13.0	12.0	11.0	10.0	10.0	10.0	10.0	9.1	8.3	7.0	6.3	5.1	3.3	2.2	1.3	0.7	0.29				
3.3	9.2	10.0	11.0	10.0	9.8	10.0	10.0	8.9	8.0	9.3	8.4	6.9	6.3	7.0	5.6	4.5	4.9	3.6	2.6	1.4	0.85	0.36	0.17	0.13	0.05

(a) Average

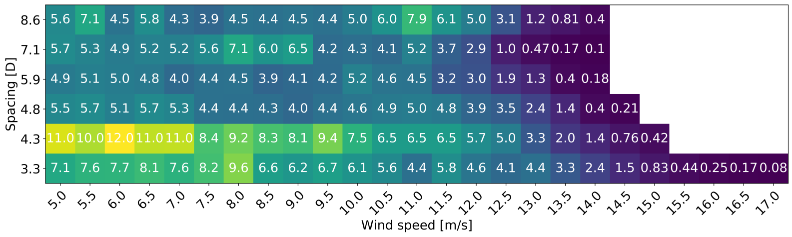

Spacing [D]	5.0	5.5	6.0	6.5	7.0	7.5	8.0	8.5	9.0	9.5	10.0	10.5	11.0	11.5	12.0	12.5	13.0	13.5	14.0	14.5	15.0	15.5	16.0	16.5	17.0
8.6	5.6	7.1	4.5	5.8	4.3	3.9	4.5	4.4	4.5	4.4	5.0	6.0	7.9	6.1	5.0	3.1	1.2	0.81	0.4						
7.1	5.7	5.3	4.9	5.2	5.2	5.6	7.1	6.0	6.5	4.2	4.3	4.1	5.2	3.7	2.9	1.0	0.47	0.17	0.1						
5.9	4.9	5.1	5.0	4.8	4.0	4.4	4.5	3.9	4.1	4.2	5.2	4.6	4.5	3.2	3.0	1.9	1.3	0.4	0.18						
4.8	5.5	5.7	5.1	5.7	5.3	4.4	4.4	4.3	4.0	4.4	4.6	4.9	5.0	4.8	3.9	3.5	2.4	1.4	0.4	0.21					
4.3	11.0	10.0	12.0	11.0	11.0	8.4	9.2	8.3	8.1	9.4	7.5	6.5	6.5	6.5	5.7	5.0	3.3	2.0	1.4	0.76	0.42				
3.3	7.1	7.6	7.7	8.1	7.6	8.2	9.6	6.6	6.2	6.7	6.1	5.6	4.4	5.8	4.6	4.1	4.4	3.3	2.4	1.5	0.83	0.44	0.25	0.17	0.08

(b) Standard deviation

Figure 6.8: Two heatmaps showing the average and standard deviation of the MAPE in percentages for turbine spacing and wind speed combinations.

6.5. Power-yaw loss coefficient

The classical approach of applying the power-yaw loss in the FLORIS model is to give all the turbines in the farm the same value. Different power-yaw loss values are reported. For example, $\alpha = 3$ derived from Blade

Element Momentum (BEM), with which often the power loss due to yawing is overestimated as shown by Madsen et al. [86]. On the other hand, it was found experimentally by Schepers [114] that α = 1.8, while Dahlberg and Montgomerie [24] found a range for α between 1.8 and 5.4. In FLORIS, α = 1.88 is commonly used as obtained by Gebraad et al. [49] through SOWFA. Finally, Medici [91] finds α = 2 from wind tunnel data. All these power-yaw loss coefficients are determined for free standing turbines. However, according to Liew et al. [84], the power-yaw loss coefficient depends on whether a turbine is operating in the wake of another turbine or not.

The power-yaw loss coefficient for waked turbines can be analytically derived [84]. However, this would require results from aero-elastic simulations, which is out of the scope of this thesis. Luckily, Liew et al. [84] performed their research on a turbine with similar turbine characteristics, namely a rated power of 2.3 MW and a diameter of 96.2 m, which is only 3.6 meters more than the Siemens SWT2.3-93 turbine. Therefore, it is assumed that their results can be used to estimate the power-yaw loss coefficient per spacing through figure 6.9. Note that these values are obtained for fully waked turbines, while in this research, they are also used for partly waked turbines. This probably results in an underestimation of the power production of partly waked turbines that operate with a yaw angle.

Table 6.3: The power-yaw loss coefficient per turbine spacing.

Spacing [D]	**α [-]**
3.3	2.20
4.3	2.17
4.8	2.10
5.9	1.98
7.1	1.89
≥8.6	1.80

By using the LES results shown in figure 6.9, which are considered to be the most accurate, α is found for every relevant turbine spacing. The resulting α values can be seen in table 6.3. For all other spacings between turbines that occur, α = 1.8 is used. For free stream turbines, α = 1.7 is used, regardless of the spacing of the row those turbines are in.

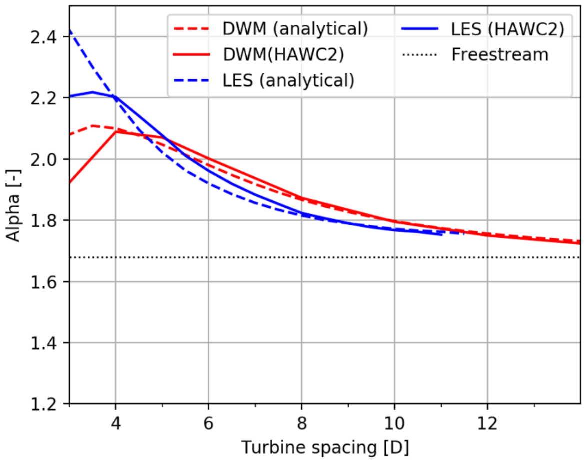

Figure 6.9: This figure shows the power-yaw loss coefficient for waked turbines as a function of turbine spacing for four different power calculation methods. Taken from Liew et al. [84].

6.6. Validation

In this section, both the power curve derivation method and the tuning parameter calibration method are validated. This is done by assessing the model's performance on different parts of the data set.

6.6.1. Power curve

A new power curve is derived in section 6.3 by using data set **S1**. The derived power curve is tested on all three data sets in **S1** separately, to make sure that it is well derived. The result of **S1.3** is visualized in in figure 6.10.

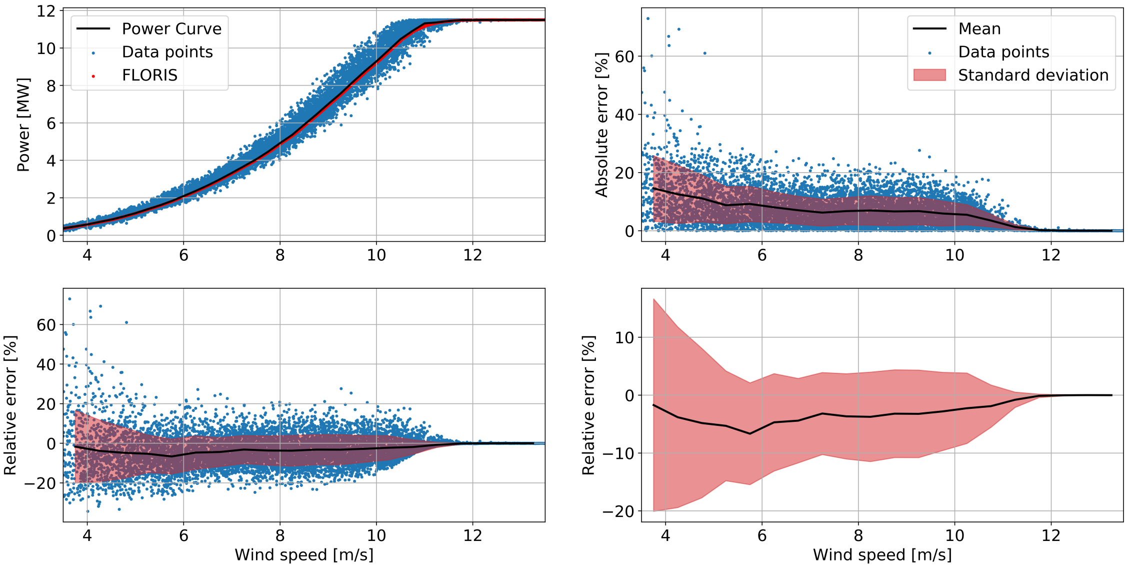

Figure 6.10: FLORIS output for free standing turbines compared to SCADA data (**S1.3**) including the relative and absolute error. The relative error is plotted once more for clarity.

It can be seen that figure 6.10 shows similar behaviour as figure 6.3. However, the MPE is negative, hence FLORIS underestimates on average the actual power for $\theta_\infty = 120°$. The mean and standard deviation of both the relative and absolute error for all three sub data sets can be found in table 6.4.

This table shows interesting results. First of all, the MPE is positive for some wind directions, while it is negative for others. However, apart from this bias, the relative error variability is almost equal for all wind directions considered. The same is true for the absolute error. Note that the smallest MAPE corresponds to the data set with the most data points. This is as expected because **S1.1** contributes to almost 56% of all data that the real power curve is fitted on. Therefore it should be kept in mind that there is a directional bias in the real power curve.

Other explanations for the differences in power production per wind direction might be found in the wind farm's blockage effect. This phenomenon enjoys increasing attention from academics. For example, Forsting and Troldborg [39] show that the power coefficient can be increased due to a combination of local and global blockage effect. In contrast, Bleeg et al. [12] show that wakes-only model types generally overpredict wind farm energy production. Though it is an interesting and relevant phenomenon, the blockage effect is not considered for now since it is outside the scope of this thesis.

Table 6.4: The mean and standard deviation of the relative and absolute error between the fitted power curve and measurements, per data set. Note that the number between the round brackets is the number of data points per data set.

		Relative error		**Absolute error**	
Data set [-]	**θ_∞ [deg]**	**Mean [%]**	**Standard deviation [%]**	**Mean [%]**	**Standard deviation [%]**
S1.1 (19305)	270	1.0	7.0	4.7	5.3
S1.2 (6256)	180	2.8	7.5	5.1	6.1
S1.3 (9127)	120	-3.2	8.3	6.5	6.1

6.6.2. Calibration method

Figure 6.7 shows the errors for θ_∞ = 180° and a spacing of 4.8D corresponding to the tuning parameters for θ_∞ = 180° and a spacing of 4.8D. These errors are low for all wind speeds, as expected since the tuning parameters are optimized for this specific case.

Ideally, the tuning parameters for θ_∞ = 180° and a spacing of 4.8D could be used for all other spacings as well. However, when the tuning parameters for θ_∞ = 180° and a spacing of 4.8D are used on a case with θ_∞ = 300° and a spacing of 3.3D, the error becomes very large for all wind speeds as can be seen in figure 6.7b. The average MAPE is then 17.8%.

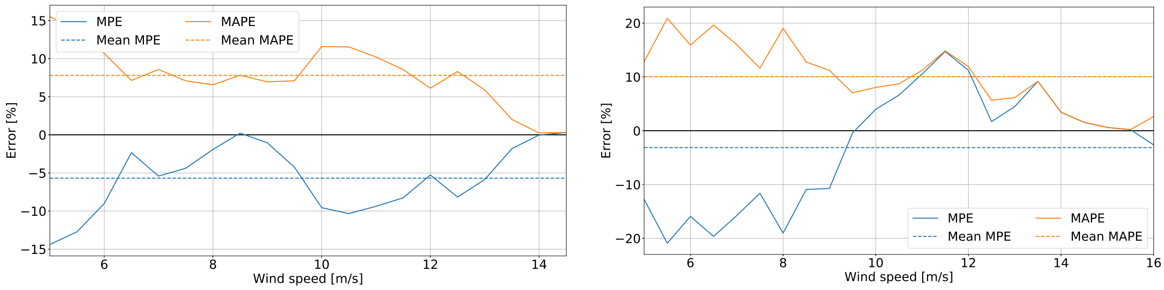

(a) $\theta_{\infty,1} = 180°$ on $\theta_{\infty,2} = 0°$

(b) $\theta_{\infty,1} = 300°$ on $\theta_{\infty,2} = 120°$

Figure 6.11: The errors when tuning parameters corresponding to θ_∞ = 180° and 4.8D spacing are used on θ_∞ = 300° and 3.3D spacing and when tuning parameters corresponding $\theta_{\infty,1}$ = 300° and 3.3D are used for $\theta_{\infty,2}$ = 120°.

However, when the tuning parameters for one spacing are found using $\theta_{\infty,1}$ and then tested on $\theta_{\infty,2}$ corresponding to the same spacing, it is found that the errors are a lot lower. For example when the tuning parameters for $\theta_{\infty,1}$ = 180° and a spacing of 4.8D are used on $\theta_{\infty,2}$ = 0°, the errors are as shown in figure 6.11a. Note that the wind speed range is now from 5.0-14.5 m/s. It is only till 14.5 m/s because of a lack of data for wind speeds higher than 14.5 m/s and θ_∞ = 0°. The MPE is negative, hence FLORIS underestimates the true power production for this wind direction. The average MAPE is 7.8%, and the standard deviation is 3.9%.

In figure 6.11b is shown what the errors are if the tuning parameters for $\theta_{\infty,1}$ = 300° and a spacing of 3.3D are used on $\theta_{\infty,2}$ = 120°. Note that the wind speed range is now 5.0-16.0 m/s, this is again due to a lack of data. First of all can be seen that for this spacing the underestimation (average MPE of -2.5 %) is more significant than for figure 6.11b. Furthermore is the average MAPE 10.1% and the standard deviation is 6.1% , which is also higher.

This increased error can be explained partly by the fact that the GCH wake model performs better in the far-wake region than in the near-wake region and partly by the fact that the derived power curve tends to underestimate the power production for θ_∞ = 120°. Next to these two things should be kept in mind that the FLORIS model relies on a lot of assumptions and simplifications and cannot capture all details in the physics of the flow perfectly. For example, in general, the difference in the power production of the first downstream turbine and the second downstream turbine is turbine spacing dependent. This might be caused by the fact that the first downstream experiences a wake with a relatively high TI, lowering its power performance. Hence, it creates a wake for the second turbine downstream that has a lower velocity deficit and TI. This, in turn, results in a relatively high power production for the second downstream turbine. This effect is especially visible when the turbine spacing is low.

Something that strengthens the hypothesis that the tuning parameters are spacing dependent is that the smaller the difference in spacing, the better the performance is when the tuning parameters of another spacing are used. So when the tuning parameters corresponding to 4.8D are used on the case with 4.3D spacing, the errors are lower than when those tuning parameters are used on the case with 3.3D spacing.

Since it is shown in figure 6.11 that a sufficiently low error is obtained between the FLORIS output and the SCADA data when spacing and wind speed dependent tuning parameters are used, the calibration method is considered to be valid. The tuning parameters can now be used to simulate the Lillgrund wind farm accurately.

One more validation step is performed related to the calibration. In section 6.4.1 is described that a bin width of 10° is used for θ_∞. Discrepancies can be observed between numerical simulations and SCADA data

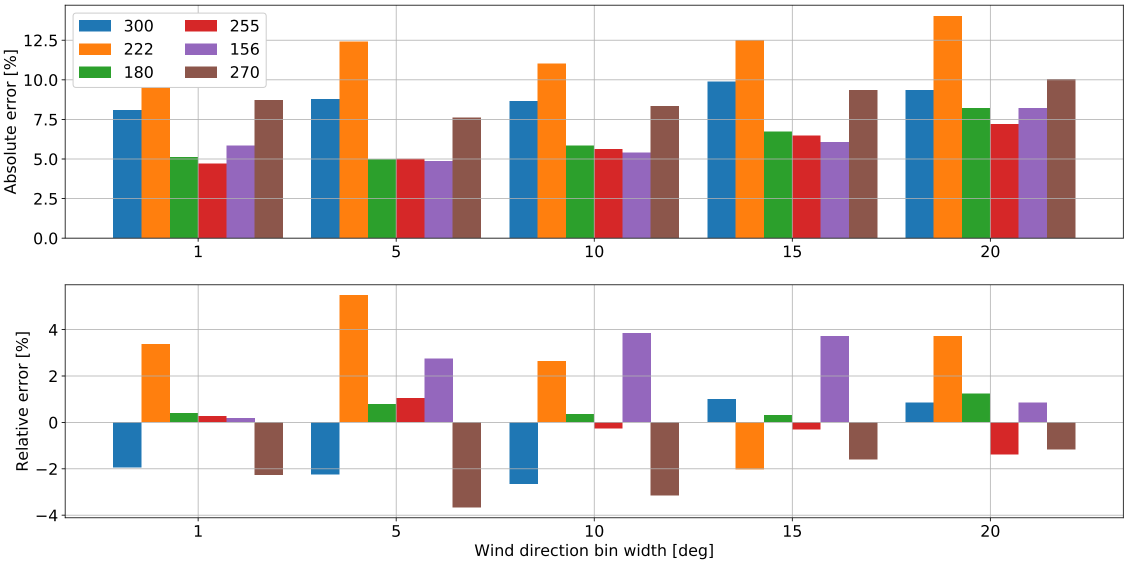

Figure 6.12: The MAPE and MPE per relevant turbine spacing wind direction for different wind direction bin widths.

for narrow sectors (< 10°) due to large wind direction uncertainty [43]. Narrow wind direction sectors may include situations where the turbines operate in conditions that are outside the span of this sector. Therefore it is checked whether the used bin width of 10° is a proper width to use. This is done by comparing the MAPE and MPE for all relevant distances and by using different wind direction bins. This is done for three different wind speeds, namely 6, 9, and 12 m/s. In figure 6.12, the results for 9 m/s are shown, but the results of 6 and 12 m/s show similar behavior. Note that the optimization described in section 6.4 is performed for the different bin widths and turbine spacings.

Figure 6.12 shows, in general, the lowest errors for a bin width of 1°. However, this is an unrealistic bin width to use, since wake effects occur for larger wind direction sectors. Therefore, The comparison between the remaining wind direction bin widths is more interesting. In general, the MAPE is slightly lower for 5° compared to the other bins. However, when also MPE is considered, it turns out that a bin width of 10° is the best performing bin width. Therefore, it can be concluded that the correct bin width is applied during the calibration.

This chapter shows that the FLORIS model is not an off-the-shelf product and that the tuning parameters on which it heavily relies are speed and turbine spacing dependent. A new calibration method is proposed as an alternative to the calibration methods so far used in the literature. The description of the new calibration method in this chapter can be used as a guideline for wind farm owners or other institutions to make the FLORIS model usable for their wind farm without needing to run complex and lengthy simulations with high-fidelity models. An additional advantage is that the model can be calibrated for many different inflow conditions, something which is hard to do when data is to be generated by high-fidelity models. The main requirement for a successful calibration is that there is a significant amount of SCADA and meteorological data available, which is often the case for existing wind farms. This calibration method can not be used for wind farms that do not exist yet, simply because, in that case, there is no SCADA data available. A summary of the new calibration method in the form of a step-by-step guide can be found in appendix B.

7
Uncertainty quantification

This chapter elaborates on the uncertainty analyses performed in this thesis. First, a short general introduction of UQ is given, and previous uncertainty analysis done in this thesis, are summarized in section 7.1. Next, the uncertainty in the input conditions is addressed in section 7.2. Then, a global UQ method for the FLORIS model is described, and the results are analyzed in section 7.3. The effect of the power-yaw loss correction is investigated in section 7.4 and, finally, a global Sensitivity Analysis (SA) is performed in section 7.5.

7.1. Introduction

There are several types of uncertainty present when using a model to simulate reality. These uncertainty types are graphically presented in figure 7.1. The goal of the uncertainty analysis is to estimate and reduce the model prediction error. This model prediction error is an indication of the certainty of the output of the model.

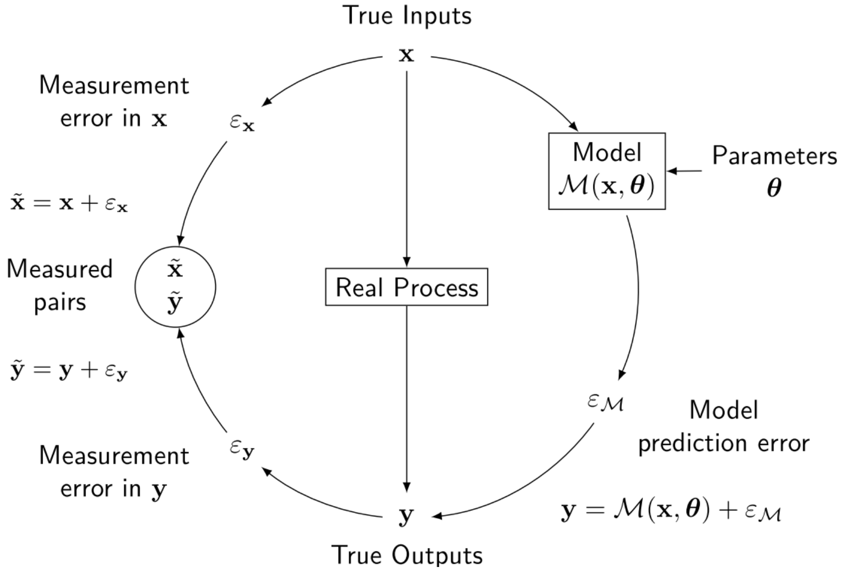

Figure 7.1: Measuring and modeling reality. Taken from Murcia Leon [95].

In this chapter, the uncertainties are divided into four parts: the input uncertainties, model uncertainties, the sensitivity of the tuning parameters, and the uncertainties addressed throughout the thesis. The latter is discussed in this section. The whole thesis is focused on not only analyzing the average values but, if possible, also the variability. This is often done by calculating the standard deviation.

For example, in table 4.1 is shown that the linear fit is the best fit based on the standard deviation of the coefficient of determination. Also, while calibrating, the distributions of the errors are considered in figure 6.4, showing that the standard deviation is a proper way of determining the variability. Afterward is shown in table 6.4 that the spread in error is considerable but in the same order for all the wind directions considered. The same is observed for the MAPE in figure 6.8. Figure 6.5 showed that the prediction of FLORIS is within the 68% confidence interval for all considered turbines. Finally, in table 8.2 is shown what the estimated AEP production would be if the whole wind farm is always operative, again including the 68% interval.

All these small analyses increase the grip on the certainty of the results observed. This happens because by considering the variability, it can be shown that the average value observed is not just a lucky draw, but that the true value is often close to this value (or not). All variabilities observed in these analyses are considered small enough to make the average values connected to them, reliable estimates.

7.2. Input uncertainties

There are several inputs for the FLORIS model. They can be separated into two categories: ones with considerable uncertainty and ones with very low uncertainty that are assumed to be exact. Inputs like the turbine positions, the rotor diameter, and the hub height are, for example, considered to be exact. Inputs with considerable uncertainty are, for example, wind speed, turbulence intensity, and atmospheric stability. Further inputs with uncertainty are the tuning parameters, which are further discussed in section 7.5. The focus in this part will be on the uncertainty in the wind direction and yaw position.

7.2.1. Wind direction and yaw position

Starting with the wind direction. The reference wind direction of the whole farm for the calibration is determined using the nacelle position sensor of upstream turbines (see section 4.2.1). Hansen et al. [58] have shown that this method "results in an uncertainty of more than 7° because the yaw misalignment of the reference turbines also needs to be included". Similar uncertainty can be observed in figure 4.3. Two other sources of wind direction uncertainty are spatial variability of the wind direction within the wind farm and the wind direction averaging period [43]. The spatial variability is relevant for the Lillgrund wind farm since wind direction is not strictly constant and homogeneous due to natural variabilities over a few kilometers. The wind direction averaging period is also relevant. There might be a drift of the wind direction due to large-scale weather phenomena between two averaging periods that is not accounted for [43], hence adding uncertainty. The uncertainty in the yaw positions of turbines is a result of complex phenomena, such as vorticity caused by the blades of the turbines and the speed-up effect around the nacelle [105]. On top of that, added uncertainties for both wind direction and yaw position are added, due to sensor and calibration errors.

Gaumond et al. [44] proposed a way to address the wind direction uncertainty by taking a weighted average of several simulations covering a wide span of directions, rather than one simulation with the average of the same wind direction sector (deterministic approach). This method has since been adopted in several studies, sometimes to address the wind direction uncertainty [27, 29, 109], sometimes the yaw position error [105] and sometimes both [106, 116]. Those studies have shown that the power gains from wake steering are often lower when uncertainty is taken into account. However, the yaw angles found through OUU are, in general, more robust and even outperform yaw angles found through the deterministic approach when noise is added to the input [109]. All these studies, however, only consider small wind farms or short rows and limited inflow conditions. This research aims to show the robustness of the optimized yaw angles for all inflow conditions and hence the robustness of the AEP. Note that, in this thesis, robust means that the yaw angles perform well, even for conditions that they are not optimized for or when there is much variability in the input conditions.

To use the method proposed by Gaumond et al. [44], two Probability Density Function (PDF)s need to be obtained: one of the wind direction and one of the yaw position. Gaumond et al. [44] has shown that those PDFs can be described using a normal distribution. This means that for both PDFs a representative standard deviation needs to be obtained. Unfortunately, with the current data available, it is not possible to do this for Lillgrund specifically. Assuming that the yaw controller and low-frequency wind direction changes used by Simley et al. [116] are representative, opens up the possibility to use the standard deviations that they found. In this paper is shown that the standard deviations are related as in equation (7.1) and that $\sigma_w = 5.25°$. In this paper is also shown that then $\sigma_\theta = 4.95°$ and $\sigma_\gamma = 1.75°$ are the best fit for $\sigma_w = 5.25°$.

$$\sigma_w^2 = \sigma_\theta^2 + \sigma_\gamma^2 \tag{7.1}$$

It should be mentioned that using those standard deviations for all inflow conditions and turbines is a significant simplification. Rott et al. [109] showed that the variability of the wind direction is strongly linked to the atmospheric stability and hence that not just one value can be used for all inflow conditions. Furthermore is shown by Andersen et al. [4] that the variability of the wind direction of the free stream wind is not necessarily the same as the wind that is deeper within the farm. Hence, if possible, σ_θ and σ_γ should be corrected for both of these effects. These corrections are considered outside the scope of this research, but they are interesting to include in future studies.

7.2.2. Stability

An important property of the wind so far neglected in all the analyses is the atmospheric stability. The atmospheric stability is related to the temperature distribution with height, which affects the turbulence, wind shear, and wind veer in the atmosphere due to buoyancy effects [69]. The atmospheric stability is usually classified as stable, neutral, or unstable, but sub-classes are also often used, e.g., very unstable. How the stability is classified depends on the vertical heat flux. The transfer of heat from the surface upwards to the air leads to an unstable Atmospheric Boundary Layer (ABL), while a heat flux directed downwards results in a stable ABL. In other words, when the surface is warmer than the air above it, an unstable ABL will develop, and when the surface is cooler, the ABL will become stable [122].

The effect of atmospheric stability on turbine loads, power production, and wake evolution has been researched extensively. The influence of atmospheric stability on turbine loads and power production is observed using measurements by Barthelmie et al. [7], Hansen et al. [58], Schepers et al. [115] and using simulations by Churchfield et al. [20], Lavely et al. [81], Sathe et al. [113] and is found to be significant. However, these studies do not separate the effect of atmospheric stability from turbulence intensity. Keck et al. [70] does isolate the impact of atmospheric stability using numerical simulations by changing the stability while maintaining the same TI. This paper shows that the atmospheric stability influences the wake dynamics even when TI is kept the same. The effect on the wake dynamics comes from the stability effects on the characteristic turbulence length scale of the ABL. According to the fundamental idea of a split in scales, a change in turbulent length scale should modify the distribution of turbulent energy between the wake meandering and the wake deficit evolution for a given turbulence intensity level [69]. The wind shear is also affected by the stability of the atmosphere, meaning that the velocity profile differs per stability class. This can significantly affect the power production and turbine loads because the wind speed at the lowest point of the rotor can be substantially lower than the highest point of the rotor for a stable atmosphere. This difference in wind speed is usually smaller for a neutral or unstable atmosphere, which generally results in higher wind speeds experienced by the turbine and hence power productions than for a stable atmosphere.

With the available data set of Lillgrund, it is not possible to estimate the atmospheric stability per time step. However, the available data set does contain results from a previous study on a different data set that provides an estimation of the atmospheric stability for a part of the time steps. The stability in this data set is determined by calculating the Monin-Obukhov (M-O) length. The M-O length can be described as the height of a sub-layer where the mechanically driven turbulence due to the ABL shear is more important than the buoyancy-driven turbulence [69]. A large M-O length corresponds to a neutral atmosphere, while for lower lengths, the sign is relevant. Negative for an unstable and positive for a stable atmosphere. For more details on the calculation, the interested reader is referred to the papers of Monin and Obukhov [92], Obukhov [101].

Although the determination of the atmospheric stability is considered outside the scope of this research, the available stability data is still analyzed. Unfortunately, the stability could not be linked to the variability in wind direction σ_θ, which could prove to be useful for the robust yaw angle optimization in section 8.4 [109]. However, the stability can be divided into wind direction bins to see if a correlation can be found. It should be noted that the available stability data is classified using nine sub-classes ranging from extremely unstable to extremely stable. However, in this analysis, the number of classes have been brought back to three: stable, neutral, and unstable.

Table 7.1: Distribution of the atmospheric stability of all time steps.

Stability class	**Count [-]**	**Percentage [%]**
Stable	48,180	31
Neutral	45,926	30
Unstable	60,639	39

The available data set consists of 154,745 time steps for which the stability is known, which sums up to about three years of data. The earliest record is from 26-05-2008 and the last from 20-12-2013. Hence there are time gaps in the data. In figure 7.2 is shown per wind direction bin of 10° what the dominant stability is. The percentage indicates the percentage of occurrence of every stability class in that specific bin, summing up to 100% per bin. The distribution of the full data set is shown in table 7.1.

Figure 7.2 shows clearly that an unstable atmospheric stratification is dominant for north-eastern and eastern winds. In contrast, neutral and stable stratification is often dominant for south-eastern to south-western winds. This suggests that there is a correlation between wind direction and atmospheric stability.

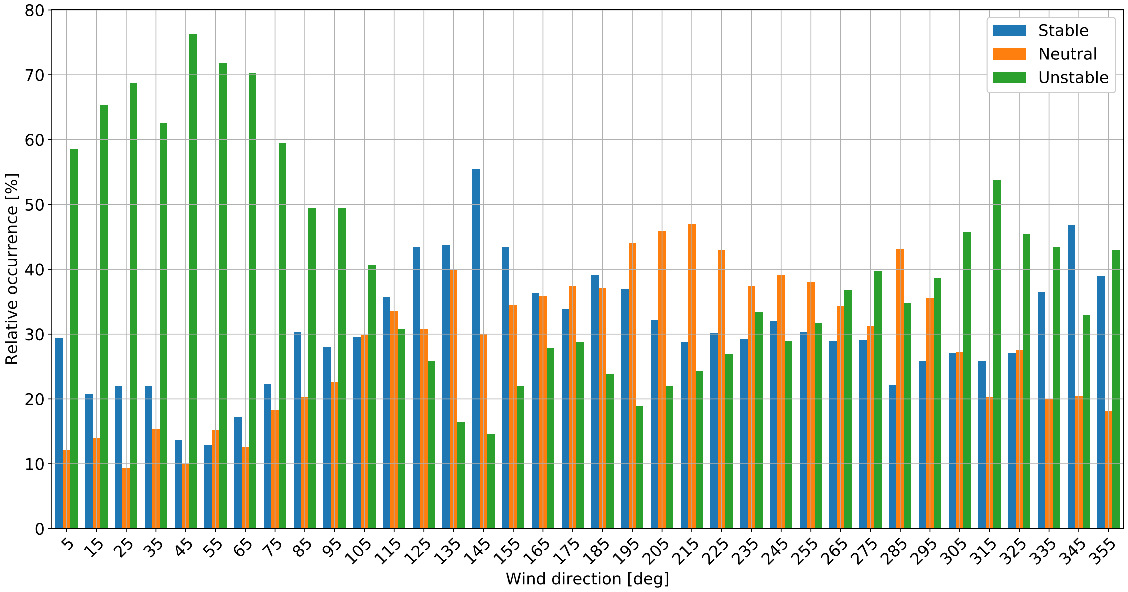

Figure 7.2: The relative occurrence of atmospheric stability per wind direction bin. Each bin width is 10°.

This correlation is probably related to the air temperature and, again, the geographical location of the wind farm (see figure 4.9). Usually, when winds go from a warm land surface to a colder water surface, the ABL becomes stable [128]. However, the opposite can happen when the water surface is warmer than the land surface. For example, north-eastern and eastern winds from land around autumn or winter are often relatively cold. When this cold air first flows over land and then over the water surface that is still cooling down from the summer, an unstable ABL will be formed. Since most land is located from the north-west and the east of the wind farm, this might be one of the reasons why the unstable stratification is so dominant for those wind directions.

The atmospheric stability influences the behavior of the wakes. An unstable atmosphere usually results in more air mixing and hence a faster wake recovery, while the opposite is true for a stable atmosphere. This means that a wake in an unstable atmosphere is usually wider with lower velocity deficits than a wake in a stable atmosphere. This makes wake steering especially effective for stable atmospheres.

There is no option in the FLORIS model to directly implement the atmospheric stability and its effects. However, it is possible to manually adjust the TI, the wind shear, and wind veer and, in that way, approximate the effects of the atmospheric stability. In this thesis, all three properties are not changed based on the atmospheric stability, since the stability is often unknown. By not including the atmospheric stability, uncertainties are added to the analysis. However, since it is found in figure 7.2 that the stability is correlated to wind direction, the TI is indirectly adjusted based on the stability through the TI roses in section 4.2.3. Uncertainties could be brought down further if more stability data is available alongside meteorological mast measurements (including standard deviation statistics). This would make it possible to create wind shear and veer roses similar to the TI roses and get a proper estimate for the uncertainty in wind direction as done by Rott et al. [109]. Including stability through wind veer and shear would be interesting for future research, since especially wind shear is expected to affect the predicted power production significantly.

7.3. Model uncertainty

Besides all the input uncertainties of which parts are addressed in section 7.2, the model itself brings uncertainties as well. One of the main goals of the FLORIS model is to have a short calculation time. This is achieved by making assumptions and simplifications that allow the model to estimate the flow field of the wind farm in a quick but not entirely accurate way. Therefore, in this section, the performance of the FLORIS model is analyzed through an UQ.

7.3.1. Method

There are two main approaches to quantify model uncertainty. The traditional method is called forward uncertainty propagation. This method aims to create a probability distribution of the model output by running the model many times for model inputs and parameters of which the uncertainty is known or estimated. This can be done, for example, using a Monte Carlo simulation [74] or a sensitivity analysis [75]. An alternative method is called inverse UQ, which relies on comparisons between modeled and measured data. Through this method, an estimation of the discrepancy distribution between measured data and the model output is created. The mean of this distribution gives the bias of the model. The width of the distribution gives an estimation of the uncertainty. In this section, the method proposed by Nygaard [100] is followed.

$$L_i = 1 - \frac{\langle P_{net} \rangle_i}{\langle P_{gross} \rangle_i} \tag{7.2}$$

$$\epsilon = \frac{(L_{obs} - L_{model})}{L_{obs}} \tag{7.3}$$

The method is elaborately described in Nygaard [100], but the main principles and equations are explained here. This method relies on the power output as a function of the inflow wind speed and wind direction only, hence making it a desirable method to quantify the uncertainty of the FLORIS model in this thesis. The benchmark of the model is the wake loss defined in equation (7.2). The angle brackets $\langle .. \rangle$ denote the mean value of the time steps considered. The subscript i indicates the case. Either the observed case using measured values L_{obs} or the model case L_{model}. The gross power P_{gross} is the power estimation produced by the farm if no wake effects would be present, and the net power P_{net} is the power estimation of the wind farm, including wake effects. The wake losses can be used to find the relative wake model error, which is defined in equation (7.3). When $\epsilon > 0$, it means that the model underestimates the wake losses, meaning that the estimated AEP is optimistic. When $\epsilon < 0$, the model overestimates the wake losses, meaning that the estimated AEP is conservative.

$P_{gross_{obs}}$ is determined by averaging the power productions of the free stream turbines for every time step and by multiplying it with the number of turbines in the farm. $P_{net_{obs}}$ is determined by summing the power productions of all turbines in the farm per time step. Often a power measurement of one or more turbines in the farm is missing, hence $P_{net_{obs}}$ is scaled to the number of turbines in the farm. $P_{gross_{model}}$ is determined by multiplying the power production obtained from the power curve with the number of turbines in the farm. $P_{net_{model}}$ is determined by summing the predicted power productions of all turbines in the farm.

A validation sample is created that makes it possible to compare the observed output with the model output. This validation sample is based on the available SCADA data. In chapters 3 and 4 is described how the SCADA data is post-processed to get a reliable data set containing power measurements per turbine per time step and the corresponding free stream wind speed and direction. Ideally, only time steps are used for which the power measurements are known for all turbines. However, as mentioned before, some power measurements per time step are often missing due to maintenance, start and stop events, or outlier removal. It is chosen that every time step with a minimum of 45 (out of 48) known power measurements can be used, to have a significant data set, while also maintaining accuracy. As mentioned above, $P_{net_{obs}}$ is scaled according to this minimum amount of known power measurements. It should be mentioned that this relaxation of the quality filter adds uncertainties to the analysis. However, there is a substantial benefit in having more available time steps, and hence these added uncertainties are accepted. The resulting validation sample contains about 37% of all available time steps, which means that in total, there are 70,749 useful time steps equal to about 491 days. The availability of time steps is visualized in figure 7.5b, where it can be seen that especially in the last two years, many time steps are missing.

Table 7.2: The wake losses and relative wake model error for the validation sample.

Parameter	**Value [-]**
L_{obs}	0.28
L_{model}	0.30
ϵ	-0.058

So the validation sample has 70,749 times steps and hence as many $P_{gross_{obs}}$ and $P_{net_{obs}}$ points. The FLORIS model is now run using the free stream wind conditions corresponding to the time steps in the validation sample as input, to obtain $P_{gross_{model}}$ and $P_{net_{obs}}$. When this is done, L_{obs} and L_{model} are obtained through equation (7.2), and eventually ϵ is calculated through equation (7.3). This gives only one value for each wake loss and the relative wake model error, as shown in table 7.2. A statistical re-sampling approach is used to get a distribution of the relative wake model error.

The statistical re-sampling approached used is the bootstrap method [31]. It is a general and versatile method through which arbitrary statistics of a sample can be obtained. The method involves drawing random bootstrap samples with replacement from a sample of data with an unknown distribution. These bootstrap samples can be combined to obtain a distribution of the statistic evaluated, in this case, the relative wake model error. Statistics such as the mean, median, standard deviation, or even more complicated ones can be obtained from this distribution. Hence also the mean value and a confidence interval.

The bootstrap method just described assumes that the data is independent and has identical distributed random variables. This is not the case for the data used here since it is essentially a time series. That means that there is a dependence between the observations. In the case of wind conditions, dependence occurs due to synoptic variations. Therefore, it is better to use an alternative bootstrap method called the circular block bootstrap method [103]. This method divides the validation sample, which essentially is a time series of SCADA data, into overlapping blocks of length k and wraps the time series in a circle to connect the beginning and the end. The created blocks are then randomly sampled with replacement until a new unique time series is created with the same size as the original validation sample. This unique time series is then one bootstrap sample of which L_{obs}, L_{model} and ϵ can be derived. This process is repeated N times to get a distribution of ϵ. The procedure followed is visualized in figure 7.3.

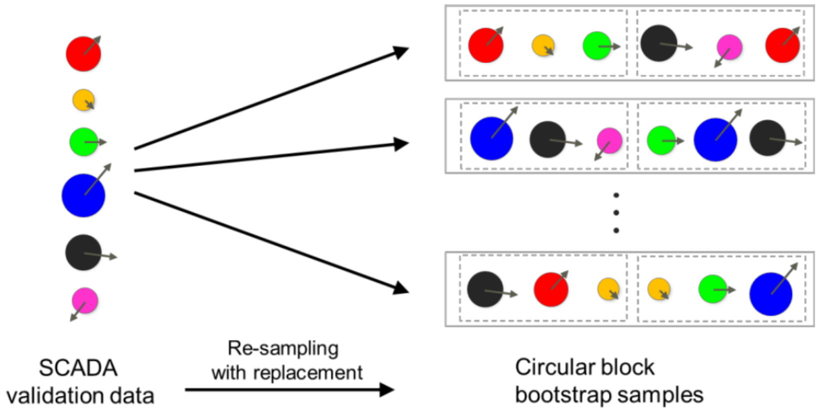

Figure 7.3: Illustration of the bootstrapping procedure. Left: the validation time series of filtered SCADA data. Each time step is a separate colour and indicates the wind speed and direction with the arrow and the wind farm power through size. Right: the validation data is sampled with replacement, generating bootstrap samples of the same length as the validation sample. The block length here is three, but in reality, it is longer. Taken from: Nygaard [100].

Each bootstrap sample obtained is a part of the original sample. However, it has unique wind conditions and a distinct distribution of variables such as the TI and atmospheric stability that influence the wake losses. Therefore, when ϵ is calculated for every bootstrap sample, a distribution of values is obtained. The bias and confidence interval of this distribution can be obtained. The bias is simply the mean value of the distribution, and it indicates how much the model generally over or underestimates the wake loss. The confidence interval is obtained by finding the 16^{th} and 84^{th} percentile. This gives the 68% confidence interval. The uncertainty of the model is defined as the half-width of this 68% confidence interval. The reason that the 16^{th} and 84^{th} percentiles are used is that it is assumed that the distribution is not normally distributed.

The circular block bootstrap method can easily be applied in Python using the open-source `arch` package. There are two inputs to the circular block bootstrap method that need to be determined, the block length k and the number of bootstrap samples N. It is important to have a block length that is large enough to capture correlations in the wind climate, such as synoptic variations. However, the larger the block length, the fewer blocks are used to create the validation sample. So the larger the block length, the 'coarser' the validation sample. The influence of the block length is further analyzed in section 7.3.2.

According to Nygaard [100], the block length should be about two days. However, later this has been adjusted to about five days (N. G. Nygaard, personal communication, June 6, 2020). A build-in function of `arch` also gives an optimal block length of $k = 5.5$ days or $k = 798$ time steps. This function finds the optimal block length by applying the algorithm proposed by Politis and White [104] and its correction by Patton et al. [102]. The minimum number of bootstrap samples N is found by performing a convergence analysis where N is increased in steps. It is found that the bias and uncertainty with a fixed block length are converged for $N \geq 1500$.

Finally, it should be noted that the experimental uncertainty is not separated from the model uncertainty for this approach. That means that the uncertainty in the inflow wind speed and wind direction are implicitly included in the model uncertainty. This is a pragmatic and conservative approach.

7.3.2. Results

When the method described in section 7.3.1 is followed, and the circular block bootstrap method is applied with $k = 798$ time steps and $N = 5000$, the distribution looks like figure 7.4. Note that a high N value is chosen to obtain a smooth distribution; however, it does not affect the observed bias and uncertainty.

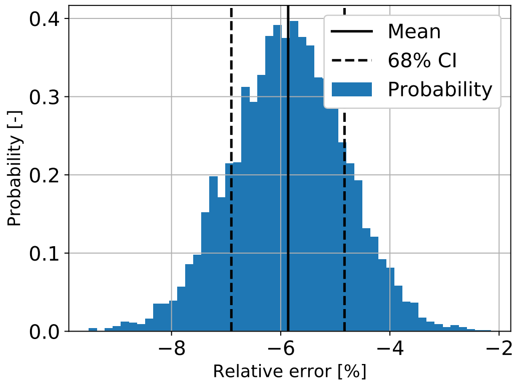

Figure 7.4: The distribution of the relative wake model error ϵ, including the mean value and 68% confidence interval.

It is observed that the mean value, hence the bias, is -5.8%, similar to what was seen from the original validation sample (table 7.2). The uncertainty, as mentioned before, is defined as the half-width of the 68% confidence interval. In this case, the model uncertainty is 1.1%. "It is common that the bias far exceeds the uncertainty for a specific wind farm, but when results of multiple wind farms are combined, the opposite is often true" - (N. G. Nygaard, personal communication, June 6, 2020). This model uncertainty means that if a wake loss is obtained of 30% that the predicted wake loss has a 68% confidence interval of [29.7%,30.3%]. However, this is without taking into account the bias that is found. When the bias is taken into account, the predicted wake loss with a 68% confidence interval becomes [27.9%,28.6%].

7.3.3. Block length and sample size

It is also investigated how the block length influences the bias and uncertainty. The UQ method is applied multiple times with different block lengths, and the bias and uncertainty are saved. The behavior can be seen in figure 7.5a. While the bias is barely affected by the block length, the uncertainty first increases with block length and eventually converges. This might be explained due to time gaps being present in the validation sample as a result of the data filtering. When large block lengths are used, the blocks might include some of these gaps, which increases the uncertainty. The time gaps in the validation sample are visualized in figure 7.5b. A value of 1 means that the time step is present in the validation sample, and a value of 0 means that the time step is not present. It can be observed that there are multiple significant gaps, and hence can be concluded that the increasing uncertainty with block length is partly a result of blocks that have these time gaps encapsulated. As mentioned before, another influence of using large block lengths is that fewer blocks per validation sample are used. This probably increases the uncertainty as well. Eventually, the optimized block length of $k = 798$ time steps is used.

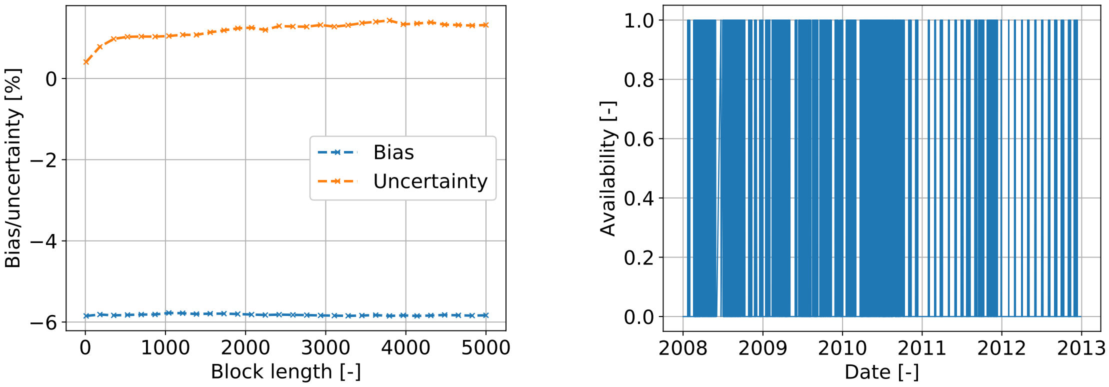

(a) Block length

(b) Availability

Figure 7.5: The bias and uncertainty as a function of block length and a visualization of the time gaps in the validation sample.

7.3.4. Bias and stability analysis

The bias and uncertainty found in section 7.3.2 show that although there is a significant bias, the FLORIS model is able to estimate the wake loss with relatively low uncertainty. When exact values of power or AEP, estimated in section 8.2, are needed, the bias can be used to correct the model output. It should be noted that it is an assumption that the bias is the same for the baseline and optimized case. It is tested what happens with the baseline and optimized AEP and energy gain when the FLORIS output power of each turbine is increased by 5.8%. The values found can be seen in table 7.3.

Table 7.3: The annual energy production for the baseline and optimized case and the energy gain of the optimized case with respect to the baseline case both with and without the bias applied.

Case	**AEP original [GWh]**	**Gain original[%]**	**AEP bias [GWh]**	**Gain bias[%]**
Baseline	379.0	-	401.2	-
Optimized	401.9	6.1	425.7	6.1

The baseline AEP, including the bias, is 5.9% larger than the original baseline AEP, which is a little more than 5.8% due to rounding errors. The optimized AEP, including the bias, is also 5.9% larger than the original, optimized AEP. Since both the baseline and optimized AEP go up with about 5.9%, the energy gain is equal to what is found in section 8.3, namely 6.1%. So with the assumption that the bias is the same for the baseline and optimized case, the effect of adding the bias in terms of energy gain is negligible.

$$\eta_{f,d} = \frac{\langle P_{net,d} \rangle}{\langle P_{gross,d} \rangle} = 1 - \langle L_d \rangle \tag{7.4}$$

Finally, the wind farm efficiency is investigated for different wind direction bins. The wind farm efficiency is calculated through equation (7.4), where the subscript d denotes the wind direction bin and the angle brackets ⟨...⟩ indicate that the mean value of all powers in that wind direction bin are used. Note that the wind farm efficiency is the same as one minus the wake loss. The result can be seen in figure 7.6. First of all can be seen that FLORIS follows the general trends and matches most peaks and troughs, except for two turbine spacings, namely 5.9D (75° & 255°) and 8.6D (90° & 270°). The wake loss is especially overestimated by FLORIS in the 10-100° sector.

The stability data derived and analyzed in section 7.2 is also shown in figure 7.6. This is done to visualize a possible correlation between the atmospheric stability and the ability of FLORIS to predict the wake loss accurately. As just concluded in this section, the overestimation of the wake loss by FLORIS seems to be mainly happening in the 10-100° wind direction sector. The stability distribution shows clearly that the atmosphere is dominantly unstable in that same wind direction sector. This analysis suggests that FLORIS is,

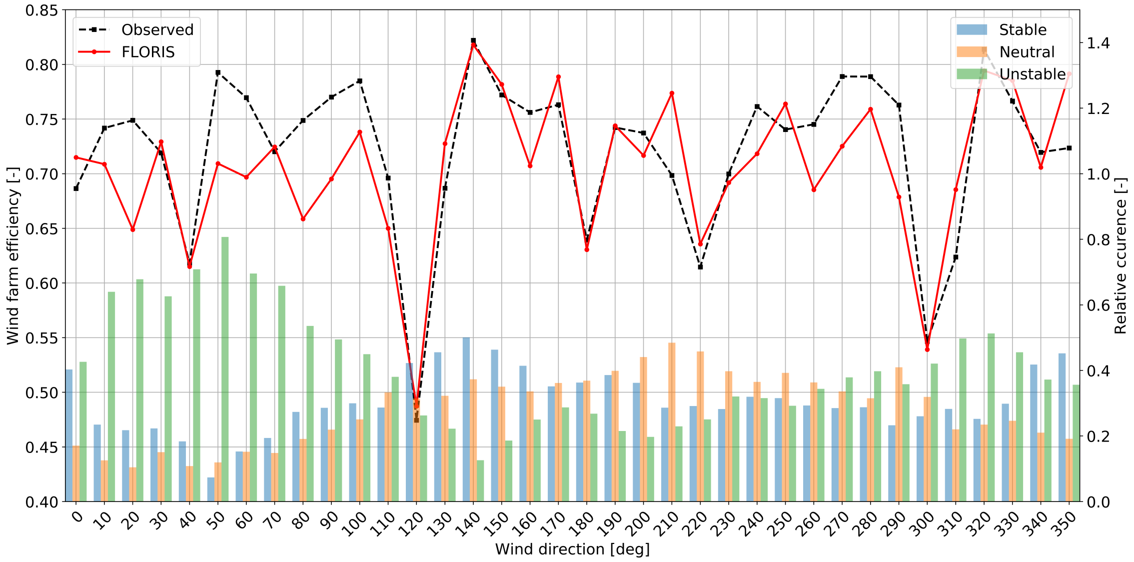

Figure 7.6: The wind farm efficiency per wind direction bin, including the stability distribution per wind direction bin.

in general, better able to model a stable and neutral atmosphere (with the current settings) than an unstable atmosphere. This again highlights the importance of taking into account the atmospheric stability and that it has a separate effect from TI [69].

Note that the stability data shown does not necessarily correspond to the same time steps as the power measurements. Furthermore, it should be mentioned that there might be another reason why FLORIS is less precise in predicting the wind farm efficiency in the 10-100° sector. In this thesis, it is assumed that the tuning parameters are dependent on the turbine spacing. Most of those tuning parameters are found with data from 146-310° sector since most wind occurs from those directions, and hence more data is available. On the other hand, a close match is found at 30°, 40°, and 70°.

This section shows that the FLORIS model with the current tuning parameters can mimic the real output with a low uncertainty but a considerable bias. The relative wake error distributions are similar to what is observed in the paper by Nygaard [100], while, of course, another model is considered here. This section also shows that although there is a considerable bias, that this bias barely influences the observed energy gain. The bias can be partly explained by the fact that in this research, the FLORIS model is not adjusted based on atmospheric stability. All in all is found that FLORIS can predict the wake losses with low uncertainty, albeit with a bias, and hence that the FLORIS model gives a reliable representation of the wake losses for non-yawed cases. Assuming that this means that FLORIS can also predict wake losses well for yawed cases, gives confidence in the gains observed in chapter 8.

7.4. Power-yaw loss correction

In section 6.5 is described how the power-yaw loss coefficient is applied per turbine depending on the turbine spacing relevant to certain free stream inflow conditions. This is done because Liew et al. [84] show that the power-yaw loss is more significant for turbines operating in the wake of another turbine. As mentioned before, it is assumed in this research that the same power-yaw loss coefficient should be used for partly waked turbines as fully waked turbines. The power output found through FLORIS is, therefore, probably an underestimation. In section 8.2 and section 8.3 is explained how the AEPs, the wind farm efficiency, and the energy gains found in this section are determined.

In figure 7.7 is shown what the difference is between running FLORIS with and without the power-yaw loss correction. It can be seen that, as expected, the wind farm efficiency is generally slightly higher when $\alpha = 1.7$ is used for all turbines. These differences can only be observed close to alignment wind directions since the turbines in the farm are yawed in that case. The most significant differences occur for the smallest

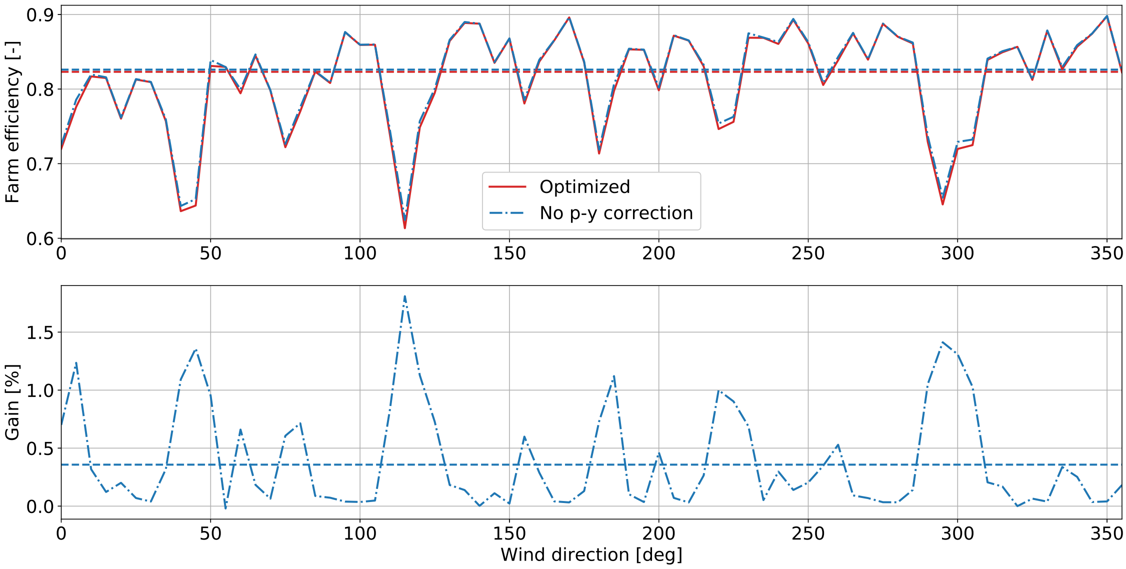

Figure 7.7: The wind farm efficiency and deterministic gain using optimized yaw angles with and without the power-yaw loss correction.

turbine spacings, likely because the largest yaw angles are used in that case, amplifying the effect of using a different power-yaw loss coefficient.

Table 7.4: The optimized annual energy production with and without applying the power-yaw loss correction.

Parameter	**Value**
AEP_{pyc}	401.9 GWh
AEP_{nopyc}	403.4 GWh
Difference	0.4 %

The difference in the optimized AEPs can be seen in table 7.4. AEP_{pyc} is the original optimized AEP with power-yaw loss correction and AEP_{nopyc} is the optimized AEP with all turbines having $\alpha = 1.7$. As can be seen, AEP_{nopyc} is slightly larger than AEP_{pyc}, which is as expected since the predicted power-yaw loss is lower. The difference is about 0.4%. Though this does not seem like a big difference, it is considered to be more accurate, and the power-yaw correction becomes more relevant when yaw and wind direction uncertainty is taken into account [84].

The baseline AEP is not shown because it is the same in both cases. That is because, for the baseline case, all yaw angles are set to zero, meaning that the cosine part in equation (6.1) is always one, no matter what α is used.

This section shows what happens to the AEPs and the optimization when the power-yaw loss correction is not applied. This means that instead of applying the power-yaw loss coefficient per turbine per spacing, the power-yaw loss coefficient is set to the free stream value of $\alpha = 1.7$ for all turbines. Apart from that, the whole model is the same, and the yaw angles are once more optimized. A small increase is observed when the power-yaw loss correction is not applied. Applying the correction factor is expected to be more significant when yaw and wind direction uncertainties are taken into account. However, the effect in this section is minimal, and the decrease in AEP is therefore small.

7.5. Global sensitivity analysis

Another uncertainty of the FLORIS model is present in the tuning parameters of the model. The FLORIS model is a complex model with multiple parameters and non-linear relations. Determining the exact uncertainty related to those parameters is, therefore, a complex and computationally expensive task. However, the influence of the uncertainty can be estimated by estimating the sensitivity of the parameters on the model output. In other words, a method can be used to rank the parameters based on their sensitivity to the model's output using a SA.

The method used is the variance-based Sobol method. The Sobol method, first described by Sobol [117] (in Russian), is a global and model-independent sensitivity analysis method based on variance decomposition. It can handle non-linear and non-monotonic functions and models. The method decomposes the variance of the model output into fractions that can be connected to parameters or sets of parameters in the model. Hence, the method makes it possible to distinguish highly influential parameters from non-influential parameters. Furthermore, background information about the interaction between pairs of parameters can be obtained.

For example, imagine a simple model that has two inputs and one output. Using the Sobol method, it can be determined that in this hypothetical case, 60% of the output variance can be attributed to input one, 25% to input two, and 15% to the interaction between input one and two. These percentages give a direct estimation of the sensitivity of the parameters. These percentages are so-called Sobol sensitivity indices and generally given in fractions rather than percentages. The first two percentages are called first-order Sobol sensitivity indices, while the last percentage describing the interaction between the set of inputs, is a second-order Sobol sensitivity index. The method to find these indices is described in section 7.5.1. The input cases considered are introduced in section 7.5.2, and finally, the results of the analysis are discussed in section 7.5.3.

7.5.1. Method

Full details on the Sobol method can be found in Sobol [117] with examples on how to apply the Sobol method by Nossent et al. [97], Zhang et al. [130]. In this subsection, only the main principles and equations are shown and explained. Let a function describe the model as in equation (7.5), where $\vec{X}$ contains p amount of uncertain parameters, and Y is a chosen univariate model output. The function $f(\vec{X})$ can be decomposed into summands of increasing dimensionality as shown in equation (7.6), where $\vec{X} = [X_1, ..., X_p]$ is the parameter set.

$$Y = f(\vec{X}) = f(X_1, ..., X_p) \tag{7.5}$$

$$Y = f_0 + \sum_{i=1}^{p} f_i(X_i) + \sum_{i=1}^{p} \sum_{j=i+1}^{p} f_{ij}(X_i, X_j) + ... + f_{1,2,...,p}(X_1, ..., X_p) \tag{7.6}$$

The parameter f_0 in equation (7.6) is a constant when the input factors are independent, and each term in the equation is chosen to have a zero average and is square-integrable. The value of f_0 is, in that case, equal to the expectation value of the output. Besides, the summands are mutually orthogonal, and the decomposition is unique.

$$V(Y) = \int_{\Omega^p} f^2(\vec{X}) d\vec{X} - f_0^2 \tag{7.7}$$

$$V_{i_1...,i_s} = \int_0^1 ... \int_0^1 f^2_{i_1...i_s}(X_{i1}, ..X_{i_s}) dX_{i_1} ... dX_{i_s} \tag{7.8}$$

$$V(Y) = \sum_{i=1}^{p} V_1 + \sum_{i=1}^{p-1} \sum_{j=i+1}^{p} V_{ij} + ... + V_{1,...,p} \tag{7.9}$$

The total unconditional variance is defined as in equation (7.7) where Ω^p represents the p-dimensional unit hyperspace where the parameter ranges are scaled between 0 and 1. The partial variances, that are components of the total variance decomposition, are computed in equation (7.8) from each of the terms in equation (7.6). In equation (7.8) $1 \leq i_1 \leq ... \leq i_s \leq p$ and $s = 1, ..., p$. Given the assumption that the parameters are mutually orthogonal, the variance decomposition can be written as equation (7.9).

$$S_i = \frac{V_i}{V} \tag{7.10}$$

$$S_{ij} = \frac{V_{ij}}{V} \tag{7.11}$$

$$S_{Ti} = S_i + \sum_{j \neq i} S_{ij} + ... = 1 - \frac{V_{\sim i}}{V} \tag{7.12}$$

So the variance is now decomposed in parts that contribute to the total variance. These variance contributions can be connected to individual parameters or parameter interactions. This makes it possible to find the Sobol sensitivity indices. The Sobol sensitivity indices are characterized by the ratio of the partial variance to the total variance. In this thesis, in total, three Sobol sensitivity indices are investigated. The first-order sensitivity S_i found through equation (7.10), the second-order sensitivity S_{ij} found through equation (7.11) and the total-order sensitivity S_{Ti} found through equation (7.12). Note that $V_{\sim i}$ means the variance as a result of all parameters except X_i [110].

S_i is a measure of the contribution of the individual parameter X_i to the total model variance. The partial variance V_i is the variance of the conditional expectation $V_i = [E(Y|X_i)]$. S_{ij} is a measure of the contribution of the interaction between parameters X_i and X_j on the total model variance. Finally, S_{Ti} describes the main effect of X_i and all its interactions with other parameters in all orders, so up to the p^{th}-order.

When it is assumed that the input factors are orthogonal and that the model is additive, S_i is equal to S_{Ti}, and the sum of S_i is 1. Note that this also means that the sum of S_{Ti} is 1. When a model is non-additive, it means that there is an interaction between the input factors. This has as a result that $S_{Ti} \geq S_i$, that the sum of S_i is now smaller than 1 and that the sum of S_{Ti} is larger than 1. In other words, one can determine the impact of the interaction between X_i and other parameters by analyzing the difference between S_{Ti} and S_i.

To find S_i and S_{Ti}, a total of $N(p+2)$ evaluations of the model are needed, where N is the sample size. To find S_{ij} an additional $N \cdot p$ evaluations are needed, resulting into a total of $N(2p+2)$ evaluations. This means that the number of model evaluations quickly increases with increasing sample size. The sample size needed for a reliable Sobol method depends on the number of parameters p and on how complex the model is [130]. The reliability of the Sobol indices can be tested by using confidence intervals. Since it is unrealistic to perform the SA with $N(p+2)$ model evaluations multiple times, bootstrapping with resampling is used instead [32]. The N samples used to get all the evaluations are sampled 100 times with replacement to get distributions of S_i, S_{Ti}, and in some cases also S_{ij}. From these distributions, the 95% confidence intervals are determined. In general, one wants the confidence intervals of the most sensitive parameters to be as small as possible, which should be at least smaller than 10% of the sensitivity indices for the most sensitive parameters [130].

7.5.2. Cases

The global sensitivity of the input and tuning parameters are analyzed for several cases shown in table 7.5. In this table is indicated per case number what the model output and the wind input is, the number of parameters p considered, the Sobol sensitivity indices considered, the corresponding turbine spacing, and the sample size N. The cases are determined for the same reasons and similarly as described in section 8.4, so in general, the wind directions and speeds that result in the highest AEP gains (section 8.2) are investigated further.

Sometimes only first and total-order sensitivity indices are considered because then only $N(p+2)$ evaluations are needed as opposed to $N(2p+2)$ when also the second-order sensitivity indices are considered. So when only S_i and S_{Ti} are investigated, the sample size N, and hence the reliability of the Sobol indices, can be increased.

Furthermore, sometimes the farm power P_f is used as model output and sometimes the farm AEP AEP_f. The model complexity is different for these two outputs. Calculating the AEP_f is more complex and computationally expensive than calculating P_f. Both outputs can only be considered for one wind direction per case. For AEP_f all operational wind speeds are considered and for P_f only one wind speed, making the computation time for P_f significantly lower.

The boundaries for Ω^p for the tuning parameters are the same as in table 6.2. Note that α_r and ϵ are not considered due to time constraints. The boundaries for the wind input parameters can be found in table 7.5 for the relevant cases. I_∞ is determined through a look-up table $I(U_\infty, \theta_\infty)$ shown in figure 4.11a. Finally, the sensitivity analysis is done for both the baseline and the optimized case.

Table 7.5: Cases for the global sensitivity analysis using the Sobol method.

Case	Output	Input	p [-]	Indices	Spacing [D]	N [-]
1.1	P_f	$U_\infty = 9 \pm 0.25$ m/s $\theta = 305 \pm 2.5°$ $I_\infty = I(U_\infty, \theta_\infty) \pm 0.02$	9	S_i, S_{Ti}, S_{ij}	3.3	50,000
1.2	P_f	$U_\infty = 9 \pm 0.25$ m/s $\theta = 305 \pm 5°$ $I_\infty = I(U_\infty, \theta_\infty) \pm 0.02$	9	S_i, S_{Ti}, S_{ij}	3.3	50,000
2.1	AEP$_f$	$U_\infty = 3.5 - 25$ m/s $\theta = 305°$ $I_\infty = I(U_\infty, \theta_\infty)$	6	S_i, S_{Ti}	3.3	12,000
2.2	AEP$_f$	$U_\infty = 3.5 - 25$ m/s $\theta = 45°$ $I_\infty = I(U_\infty, \theta_\infty)$	6	S_i, S_{Ti}	4.3	12,000
2.3	AEP$_f$	$U_\infty = 3.5 - 25$ m/s $\theta = 185°$ $I_\infty = I(U_\infty, \theta_\infty)$	6	S_i, S_{Ti}	4.8	12,000
2.4	AEP$_f$	$U_\infty = 3.5 - 25$ m/s $\theta = 80°$ $I_\infty = I(U_\infty, \theta_\infty)$	6	S_i, S_{Ti}	5.9	12,000
2.5	AEP$_f$	$U_\infty = 3.5 - 25$ m/s $\theta = 160°$ $I_\infty = I(U_\infty, \theta_\infty)$	6	S_i, S_{Ti}	7.1	12,000
2.6	AEP$_f$	$U_\infty = 3.5 - 25$ m/s $\theta = 95°$ $I_\infty = I(U_\infty, \theta_\infty)$	6	S_i, S_{Ti}	8.6	12,000

7.5.3. Results

The results of case 1.1 are first analyzed in detail and then compared to what happens when a larger bin width is used for θ_∞ in case 1.2. After that, the results of case 2.1 are analyzed in detail. Finally, the results of cases 2.1 to 2.6 are compared with each other.

Cases 1.1 and 1.2

The first and total-order sensitivity indices for case 1.1 can be seen in figure 7.8. The confidence intervals are shown in black and the 10% intervals in red. As mentioned in section 7.5.1, for reliable estimates of the indices, the confidence intervals should be smaller than 10% of the most sensitive parameters. Both the values of the indices for the baseline and optimized model are shown.

First of all, it can be seen in figure 7.8 that the confidence intervals of the most sensitive parameters are well within the 10% deviation. Furthermore, it can be seen that, when only the tuning parameters are considered, only three parameters have a significant effect on the model output. Those tuning parameters are k_a, α and b_d. The effect of a_d is negligible, hence this parameter is set to zero for all speed and spacing combinations in section 6.4, and consequently for every time the model is used. The most sensitive wind input parameter is U_∞.

It should be noted that the Sobol sensitivity indices of these wind input parameters are strongly dependent on the used boundaries. For example, when the wind direction boundaries are set to $305 \pm 5°$ instead of $305 \pm 2.5°$, the influence of θ_∞ will increase significantly, as is shown in figure 7.8c. It can be seen that at the same time, the influences of the tuning parameters decrease.

It is also interesting to note the difference in the indices' values and the ranking of the parameters when the baseline model results in figure 7.8a, and the optimized model results in figure 7.8b, are compared. Especially the influence of k_a and α decrease for the optimized model, while the influence of b_d and U_∞ increase. These changes suggest that the optimal tuning parameters determined in section 6.4 for the baseline model

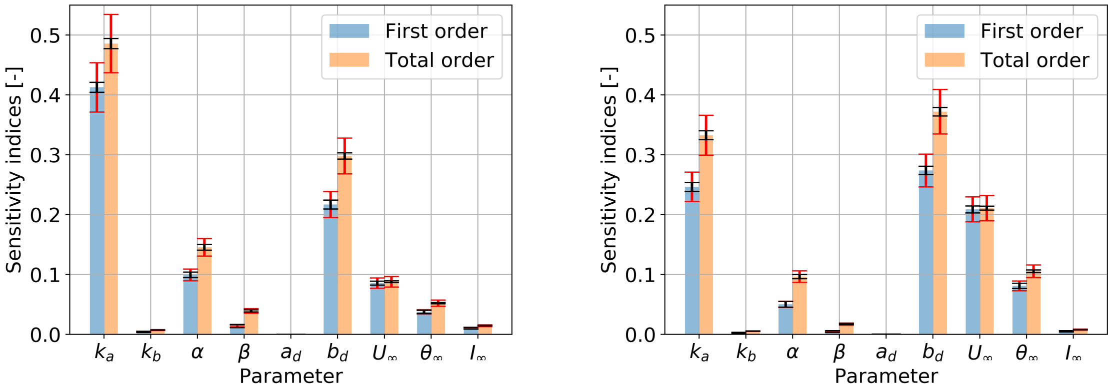

(a) Case 1.1 Baseline

(b) Case 1.1 optimized

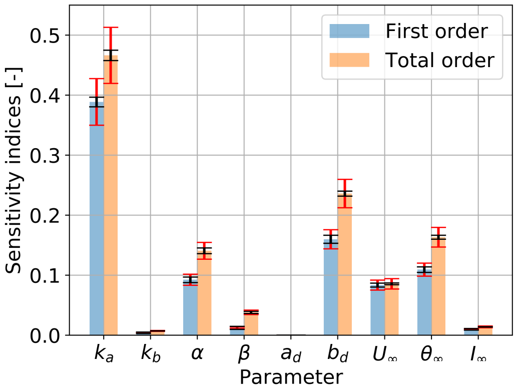

(c) Case 1.2 Baseline

Figure 7.8: The Sobol sensitivity indices of case 1.1 for both the baseline and optimized model and of case 1.2 for the baseline model only including the 95% confidence interval error bar (black) and the 10% deviation (red).

are not necessarily optimal for the optimized model. This also suggests that the assumption that the model bias and uncertainty determined in section 7.3 is the same for the optimized model might not be valid.

Finally, it can be seen in figure 7.8 that for all parameters but U_∞, the total-order sensitivity is significantly higher than the first-order sensitivity. This indicates that there are some significant second-order (and probably higher-order) sensitivities. Note that it makes sense that S_{Ti} is higher than S_i since theoretically S_i is a part of S_{Ti}. However, since the indices' values are estimates due to numerical integration, it should be checked that S_{Ti} is indeed higher than S_i. The most relevant S_{ij} are shown in table 7.6. Note that the S_{ij} values are low compared to the S_i values and that $S_{ij,CI}$ is relatively large compared to S_{ij}. Also, note that only two relevant parameter interactions are shown. The reason others are not relevant to show is either because the index value is really low, or the confidence interval is larger than the actual value. The confidence intervals decrease when larger sample sizes are used. However, the computing time, in that case, becomes unreasonably high.

Case 2.1 to 2.6

Case 2.1 is investigated in the same way as case 1.1. The first and total-order sensitivity indices are shown in figure 7.9. Note that this time only six parameters are taken into account and that the model output is AEP. Hence, all operational wind speeds are taken into account.

The first thing that is important to note when analyzing figure 7.9 is that the 95% confidence intervals are

Table 7.6: The relevant second-order sensitivity indices with the 95% confidence interval of case 1.1.

Parameters	**Model**	S_{ij} [-]	$S_{ij,CI}$ [-]
$k_a - b_d$	Baseline	5.1E^{-2}	1.4E^{-2}
	Optimized	5.6E^{-2}	1.3E^{-2}
$\alpha - b_d$	Baseline	1.2E^{-2}	7.7E^{-3}
	Optimized	1.9E^{-2}	8.3E^{-3}

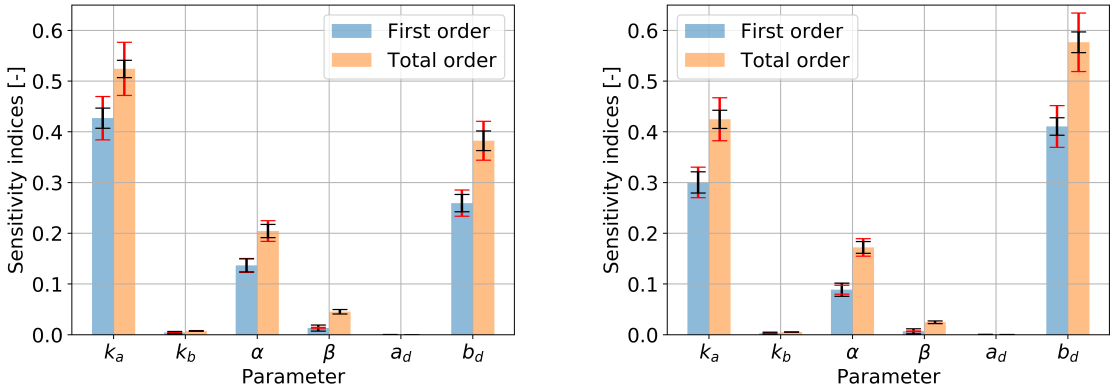

(a) Baseline

(b) Optimized

Figure 7.9: The Sobol sensitivity indices of case 2.1 for both the baseline and optimized model, including the 95% confidence interval error bar (black) and the 10% deviation (red).

visibly larger than the ones observed in figure 7.8. This is a result of N being lower than for cases 1.1 and 1.2, even though p went from nine to six. So although the 95% confidence intervals are generally within the 10% deviation, the reliability of these sensitivity estimations is lower than those of cases 1.1 and 1.2.

That being said, it can be seen that the ranking of the parameters in figure 7.9 is identical to the one observed in figure 7.8. Also, the same behavior can be observed when looking at the difference between the baseline and the optimized model. The sensitivity of the parameters k_a and α are lower for the optimized model, while the sensitivity of b_d is higher. It is interesting that both figures 7.8 and 7.9 show that for yawed turbines, the sensitivity of the tuning parameter b_d, describing the lateral deflection of the wake due to the rotation of the wake (equation (5.10)), increases. In contrast, the tuning parameters influencing the lateral deflection of the wake due to a yaw angle (equation (5.3)) are decreased. So it seems that the introduction of the yaw angle lowers the influence of the parameters directly related to this yaw angle through equation (5.3), which in turn increases the influence of b_d. An explanation can be found in the fact that equation (5.10) does not take into account the deformation and alterations to the wake as a result of the yaw angle. This results in the absolute influence of b_d being unaltered, while the relative influence increases since the other parameters do take parts of the wake deformations and alterations into account, as explained in section 5.1.

Until now, the tuning parameters have been investigated for one wind direction or a small wind direction bin only. Cases 2.1 to 2.6 are chosen to investigate whether the Sobol indices change with different wind direction or, rather, with different turbine spacings. The wind directions are determined in such a way that all six relevant turbine spacings, shown in table 6.1, are investigated. The wind direction for each spacing is chosen based on the energy gains observed in section 8.2. Although not shown here to reduce the number of figures, it has been found that the Sobol indices behave similarly for wind directions corresponding to the same spacing. The Sobol indices per spacing can be seen in figure 7.10.

Figure 7.10 shows first of all that for α the 95% confidence intervals are outside the 10% deviation, as well as for b_d for the largest two spacings. This means that the estimations of the Sobol indices are not fully converged and reliable yet. However, despite this, it still tells something about the parameters' behavior in relation to the turbine spacing. The changing sensitivity of the tuning parameters per turbine spacing

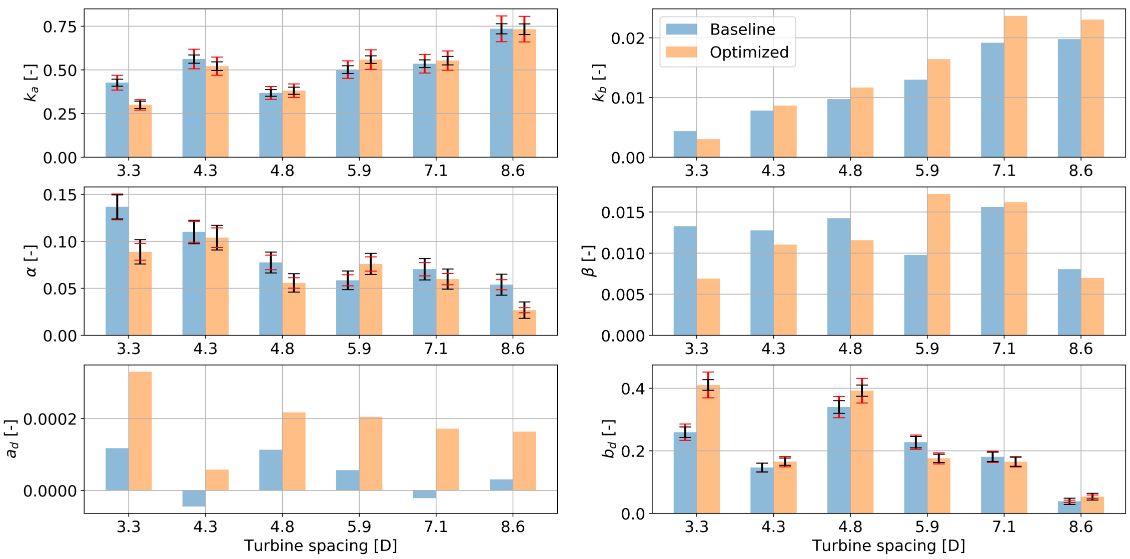

Figure 7.10: The Sobol indices per turbine spacing corresponding to cases 2.1 to 2.6 for both the baseline and optimized model. The 95% confidence interval error bar (black) and the 10% deviation (red) are shown for the most sensitive parameters.

explains in part why tuning parameter maps per turbine spacing needed to be made during the calibration.

Focusing on the three most sensitive parameters, it can be seen that the sensitivity of k_a is lower for distances closer to the near-wake region and increases when they are deeper in the far-wake region. The sensitivity of α and b_d follow an opposite trend; they are higher for distances closer to the near-wake region and lower for distances deeper in the far-wake region. These changes in sensitivity make sense because α influences the near-wake, which is further away for larger spacings. The lateral deflection influenced by b_d is relatively small compared to the lateral deflection described by the other parameters for larger spacings. No clear trends between the baseline and optimized model can be observed, except that the effect of an increasing b_d described before, only seems to happen for distances close to the near-wake region.

This section shows the sensitivity of both the tuning parameters and some input parameters. It can be concluded that the parameters' ranking is the same for both P_f and AEP_f as output. Also can be concluded that the accuracy of the model would not change much when only the tuning parameters α, k_a, and b_d are considered, while the other parameters are set to a fixed value of zero. Next to that is shown that the sensitivity of the parameters are different for the baseline and optimized model, making the assumption that the bias and uncertainty of the baseline model can be used for the optimized model, less valid. Finally can be concluded that the sensitivity of the parameters is spacing dependent, which explains in part why the tuning parameters need to be calibrated per turbine spacing.

8

Optimization

This chapter describes the optimizations performed on the Lillgrund wind farm. First, in section 8.1, it is described how the FLORIS tool is used to optimize the power output of the farm by finding the optimal yaw angles deterministically. Then, in section 8.2, it is explained how the AEP is determined for the baseline and the optimized case, and the potential energy gain is analyzed in section 8.3. Next, in section 8.4, it is shown how more robust yaw angles can be obtained by taking into account uncertainties. Finally, the findings in this chapter are discussed elaborately in section 8.5.

8.1. Yaw optimization

The FLORIS model can be used to find the optimal yaw angles. These optimal yaw angles are expected to increase the power output of the farm and hence the energy output. The goal is to increase the AEP as much as possible using yaw-based wake steering. The idea is to use the data from which the wind rose of Mast-p1 is derived as input, and see how much gain in AEP can be obtained when optimized yaw angles are used. That means that the optimization must be performed for about 100,000 time steps, which is only possible if the yaw angle optimization can be performed in a reasonable amount of time.

There is a tool available in the FLORIS model that makes it possible to obtain the optimal yaw angles for a specific case of free stream conditions. However, this requires optimization of, in the case of the Lillgrund wind farm, 48 turbines. So the tool has 48 variables to optimize. The same SLSQP minimizer used to calibrate the tuning parameters in section 6.4, is used to find the optimal yaw angles. The cost function is the negative farm power. It is negative because the optimization tool is a minimizer. An initial guess of the yaw angles can be given to speed up the optimization, but if it is not available, then the initial guess is simply a zero yaw angle for all turbines. Though it is not part of this thesis, other optimization methods than the SLSQP have been considered, and it is found that those methods are not significantly faster or more accurate in finding the global minimum. It should be noted that this has only briefly been researched for a few inflow cases.

It is tested how long it takes for a single processing core (Intel(R) Core(TM) i7-8750H CPU @ 2.20GHz) to perform one single optimization for one specific case of free stream conditions. For this test, the wind speed is set to 8 m/s, the wind direction to 185°, and the TI to the corresponding value found from the TI rose. The reason that 185° is chosen is because it is found that wake steering is especially effective when the downstream turbine is partially waked (section 8.2). The standard options for the optimization are used in this case. The computing time is in the order of ~700 s. Note that this measured time is for the optimization part only and not the initialization of the model or adjustment of tuning parameters.

This computing time is rather high, and therefore it is unfortunately impossible to perform the optimization for 100,000 time steps. Instead, it is now chosen not to optimize all 100,000 time steps but optimize for combinations of U_∞ and θ_∞ in steps over the full domain of the wind rose. Both the baseline power, meaning all yaw angles are zero, and the optimized power, can then be calculated for all these U_∞ and θ_∞ combinations. These can, in turn, be used to estimate the AEP for the baseline and optimized case, as described in section 8.2. The goal is to perform the optimization for $U_\infty = 3.5 - 25$ m/s with a step size $\Delta U_\infty = 0.5$ m/s and for $\theta_\infty = 0 - 360°$ with a step size of $\Delta\theta_\infty = 5°$. This means that 3,168 optimizations need to be performed. There is no point in optimizing for $U_\infty > 17.5$ m/s, since the farm produces rated power for all wind directions regardless of the yaw angles. That leaves the total number of optimizations needed at 2,088. Still, this would mean almost ~400 hours of computing for a single core.

So before this optimization is performed, a few things should be altered to decrease the computing time. First of all, the options of the optimization can be altered. It is found that when the function tolerance is altered from the standard $1e^{-7}$ to $1e^{-2}$, only a third of the computing time is needed, while the eventual outcome is near equal. Next to the adjustment of this tolerance, an initial guess can be given to the minimizer, which is as close as possible to the eventual outcome. At first, no initial guess is known, but it opens the possibility of a warm start. This means that first larger step sizes, e.g., $\Delta U_\infty = 1$ m/s, $\Delta\theta_\infty = 15°$, are used. The outcome of these optimizations can then be used as the initial guess for the next optimization with smaller step size, e.g., $\Delta U_\infty = 1$ m/s, $\Delta\theta_\infty = 10°$. So, for example, the optimal yaw angles are known for $U_\infty = 8.5$ m/s and $\theta_\infty = 15°$. These can then be used for the not yet determined combination of $U_\infty = 8$ m/s and $\theta_\infty = 10°$. This step can be repeated until the desired step sizes are reached.

Although the suggested alterations do decrease the computing time, it would still take a long time on a single core. Hence, a cluster of cores available through DTU is used. To properly use this cluster, the script should be written in such a way that it supports parallel computing. This is done using the package called `joblib`, which is another open-source package for Python. The optimization is eventually performed using 64 cores (Dual 12-core 2.6GHz Intel Xeon processors). Even after all the warm starts, the computing time was about half a day. Note that this includes the initialization of FLORIS, so setting the right tuning parameters, power-yaw loss coefficients per turbine, and turbulence intensity, as well as the calculation of the baseline and the optimized farm power.

Eventually, one more iteration is done. The optimal yaw angles found with step sizes $\Delta U_\infty = 0.5$ m/s and $\Delta\theta_\infty = 5°$ and a function tolerance of $1e^{-2}$ are used as an initial guess to find the optimal yaw angles with the same step size, but a function tolerance of $1e^{-3}$. This is done to make sure that a converged optimum is found.

8.2. Annual energy production estimation

Although it is interesting to know the power gain for a specific wind speed and wind direction combination, eventually, what is relevant for the owner of the wind farm is whether a gain in AEP is possible and how significant this gain is. The AEP is estimated using the powers and yaw angles derived in section 8.1 and the wind rose of Mast-p1. The equation for AEP can be found in equation (8.1) and has as unit Wh. T is the number of hours in a year, $p(U_{\infty,i}, \theta_{\infty,j})$ is the probability that a specific combination of wind speed and direction occurs and $P_k(U_{\infty,i}, \theta_{\infty,j})$ is the power corresponding to that combination of wind speed and direction. Note that the subscript k indicates whether it is the baseline or the optimized case.

$$AEP_k = T\left[\sum_i \sum_j \left(p(U_{\infty,i}, \theta_{\infty,j}) P_k(U_{\infty,i}, \theta_{\infty,j})\right)\right] \tag{8.1}$$

As mentioned before, the probability that a specific combination of wind speed and direction occurs is derived from the data of Mast-p1. A heatmap of the probability can be seen in figure 8.1. Note that this is just a more readable version of figure 4.8a.

These probabilities are combined with the calculated baseline and optimized power to estimate the AEP. The total AEP of the farm for three different cases can be found in table 8.1. This table also shows the wake losses, which is calculated with respect to the case that assumes that there are no wake effects. Finally, the energy gain of the optimized case with respect to the baseline case is shown. So this table shows that a potential gain of 6.1% is achievable when wake steering is applied. It should be emphasized that this energy gain is found using a deterministic approach, so without taking into account input uncertainties and hence the gain is likely an upper limit. This is further investigated in section 8.4.

Table 8.1: The annual energy production for the three different cases together with the wake losses for the baseline and optimized case and the mean energy gain of the optimized case with respect to the baseline case.

Case	**AEP [GWh]**	**Wake loss [%]**	**Gain [%]**
Baseline	379.0	22.4	0
Optimized	401.9	17.7	6.1
No-wake	488.3	-	-

The baseline AEP value in table 8.1 is higher than the 330 GWh AEP found in the literature [124]. Therefore the baseline AEP is verified by comparing it with the SCADA data available. Per year an estimation is made of the total AEP that the Lillgrund wind farm has produced. This is done by calculating the fractional AEP of every 10-minute time step AEP$_t$ available and sum it over a full year. Note that the AEPs in table 8.1 are

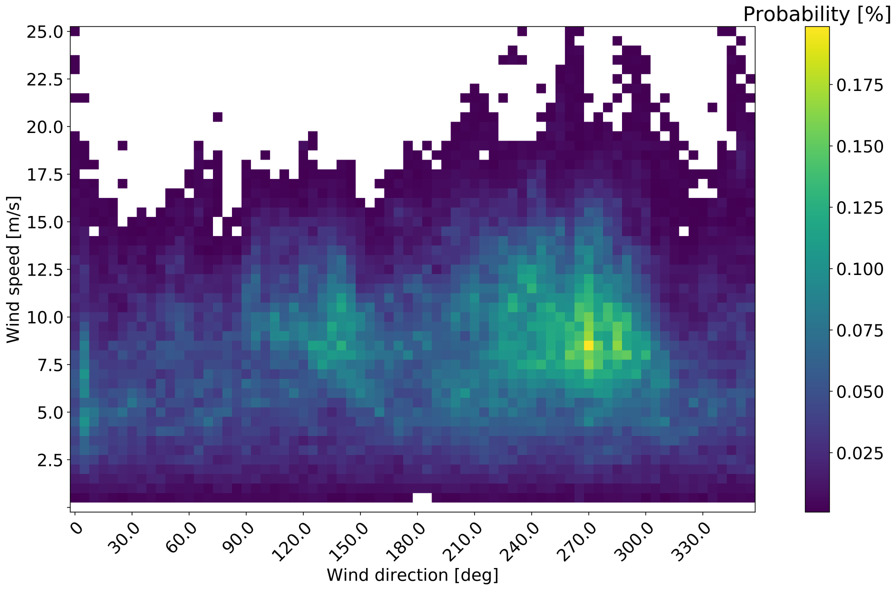

Figure 8.1: Probability of combinations of U_∞ and θ_∞ in percentages.

calculated with the assumption that all turbines in the Lillgrund wind farm are fully operative the entire year. In reality, this is often not the case. On top of that, is it possible that measurements are missing due to one of the reasons mentioned in section 3.3. Therefore when AEP_t is summed, it is afterward corrected for the number of operative turbines and the number of known time steps. The AEP after the operative turbines correction is indicated as AEP_{oc}, and the final AEP after the correction for known time steps is indicated as AEP_{tc}. The resulting AEP found per year can be seen in table 8.2. It can be seen that the AEP found through FLORIS is well within the standard deviation of the corrected AEP found through SCADA data. The discrepancy between AEP_{tc} and the 330 GWh found in literature can partly be attributed to turbines being inoperative due to either maintenance or a decrease in energy demand.

Table 8.2: The annual energy production per year obtained from SCADA data corrected for the number of operative turbines AEP_{oc} and the number of known time steps AEP_{tc}. The last row shows the average and when relevant one standard deviation.

Year [-]	**Known time steps [%]**	**AEP_{oc} [GWh]**	**AEP_{tc} [GWh]**
2008	88.4	353.6	400.2
2009	93.1	330.0	354.6
2010	89.5	306.9	343.0
2011	92.5	357.4	386.5
2012	93.8	341.0	363.6
Average	91.4	337.8±18.2	369.6±20.9

A remark should be made about this section. It is found in section 7.3 that the FLORIS model overestimates the wake loss of this wind farm on average by 5.8%. When it is assumed that this wake loss underestimation is the same for the optimized as for the baseline case, both the baseline and optimized AEP increase with 5.8%. The baseline AEP would then be 400.9 GWh, which is outside the one standard deviation range of AEP_{tc}, but almost equal to the AEP_{tc} in 2008. This can mean that the years on which figure 8.1 is based were relatively windy. Another explanation can once again be sought in the wind farm blockage effect earlier discussed in section 6.6.1. The wind rose is based on data from Mast-p1, which was present at the Lillgrund site before the wind farm was constructed. Hence wind speeds were measured that are probably higher since there was no wind farm present that could affect the wind speed due to blockage effects. This means that higher wind speeds might have been measured, and higher wind speeds mean a higher AEP. In general, it can be concluded that the AEPs found in this section are reasonable and can be further investigated in section 8.3.

8.3. Energy gain

A value of AEP on itself does not say much, hence to compare the baseline and optimized case, the quantities of interest are the normalized farm energy κ_f, wind farm efficiency η_f and the energy gain G_f. The normalized energy indicates how much energy is produced per wind direction. The wind farm efficiency describes how much AEP is produced compared to the AEP that could have been produced if there would be no wake effects at all, which is again calculated per wind direction. The energy gain shows how much the AEP can be increased per wind direction when yaw-based wake steering is applied.

The AEP is indicated as E in the following equations. The definition of κ_f can be found in equation (8.2). Note that the k in E_k again indicates either the baseline or the optimized case. E_k is normalized by the highest optimized AEP. The definition of η_f can be seen in equation (8.3). E_{nw} indicates the AEP if there would be no wakes. The energy gain G_f is calculated using E_{base} and E_{opt} through equation (8.4) and is expressed in percentage.

$$\kappa_f = \frac{E_k}{E_{opt,max}} \tag{8.2}$$

$$\eta_{f,k} = \frac{E_k}{E_{nw}} \tag{8.3}$$

$$G_f = 100\frac{\left(E_{opt} - E_{base}\right)}{E_{base}} \tag{8.4}$$

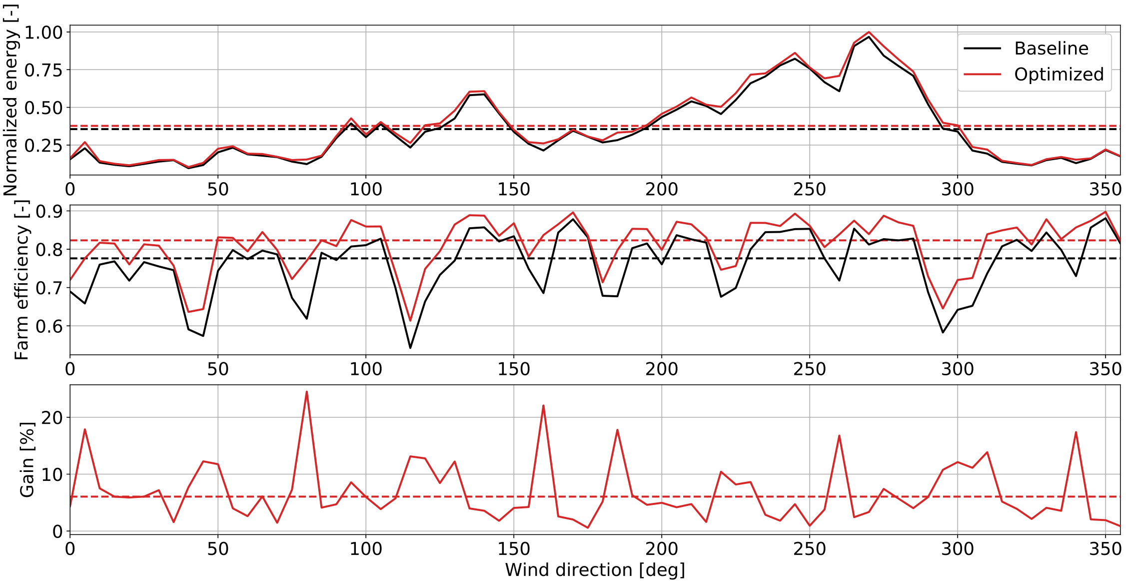

Figure 8.2: The normalized energy, wind farm efficiency and deterministic energy gain per wind direction for the baseline and optimized case when wake steering is applied.

Figure 8.2 shows the normalized energy, wind farm efficiency, and energy gain per wind direction. These values are obtained by summing the AEP of all wind speeds per wind direction. From the normalized energy plot can be seen that the most energy is produced with western and south-western winds. This is not surprising since this is the most frequently occurring wind direction (figure 4.1a). Also, winds from around 135° contribute significantly, again, because this is a wind direction that occurs relatively often.

The more interesting plots are the lower two plots showing a comparison of the wind farm efficiency and energy gain per wind direction for the baseline and optimized case. The wind farm efficiency plot shows some deep troughs for the wind directions for which the turbines are aligned. At the same time, high gains are achieved for wind directions close to these alignment directions. Although it is hard to see in the figure, the gain peaks generally occur for an alignment wind direction plus 5° (e.g., 80°, 160°, 185°, etc.). This suggests that wake steering is especially effective when downstream turbines are partly waked instead of fully waked.

Interestingly enough, the highest peaks occur for 80° and 160° corresponding to 5.9D and 7.1D turbine spacing, respectively. It should be noted that it is shown in section 7.3.4 that there is especially an increased uncertainty for those wind directions. However, the gains observed are between two model outputs; hence it might still very well be valid gains. The gains around wind directions for the lowest turbine spacing 3.3D, so 120° and 300°, do not show sharp peaks, but rather relatively high gains for a wide range of wind directions. The same can be observed for 4.3D, hence around 42° and 222°. This is because the turbine spacings are so small that the wake losses are relevant for these wider ranges of wind directions. Therefore, the gain achieved at these low turbine spacings are very relevant for the total AEP gain.

The dashed lines indicate the values when the wind farm efficiency and the energy gain are averaged over the wind directions. Note that since the probability of the wind is already an integral part of the AEP, these mean values indicate absolute values. So this graph shows that when wake steering is applied, the wake loss decreases from 22.4% for the baseline case to 17.7% for the optimized case. The graph also shows that the mean energy gain in total AEP is 6.1%, which is in the same order as the 3.7% found by Gebraad et al. [50].

Although 6.1% is an exciting increase, it should be noted that this gain is a deterministic value and is likely an upper boundary. It is expected that this gain will decrease when uncertainty in wind direction and yaw angle is taken into account. Furthermore should be noted that Lillgrund is a very densely spaced wind farm. That means that although the gain is relatively large for this wind farm, the values for more widely spread wind farms is probably lower. This can also explain why the energy gain for Lillgrund is higher than the one found by Gebraad et al. [50]. Finally, in section 7.3 is shown that the FLORIS model predicts the wake losses with a bias of -5.8% and uncertainty of 1.1%. In section 7.3.4 is shown that the influence of the bias on the energy gain is negligible. However, the uncertainty of 1.1% is still present and influences the wake losses. When it is assumed that the model uncertainty for the optimized case is the same as for the baseline case, then the 68% CI of the energy gain can be given as [5.5%,6.6%]. Note that the energy gain from now on is given in the form x.x% [x.x%,x.x%], where [x.x%,x.x%] indicates the 68% CI.

8.4. Optimization under uncertainty

This section will address the influence of input uncertainty on the optimization process, and the eventually found energy gain. So far, in some sections, the uncertainty is quantified in terms of variability to estimate the reliability of the results found, but the variability is not used further. In this section, the variability of the inputs is used to decrease the model output uncertainty. The origin and chosen quantification of the input uncertainties are introduced earlier in section 7.2.

Two methods to use these uncertainties are described in section 8.4.1, where also a convergence study is done for the number of evaluations needed per PDF. The results of the first method are analyzed and compared to the deterministic results in section 8.4.2. The second method results in an additional yaw set of which the performance is compared to the deterministic yaw set in section 8.4.3.

8.4.1. Method

The origin and estimation of the uncertainty for the wind direction and yaw position are described in section 7.2. In this part of the thesis, the uncertainties in wind direction and yaw position expressed as σ_θ and σ_γ, respectively, are used. The values used are $\sigma_\theta = 4.95°$ and $\sigma_\gamma = 1.75°$, which are taken from Simley et al. [116]. It is assumed that two representative normal distributed PDFs can be obtained using those standard deviations. Both can be found in figure 8.3a. Note that the PDFs are normalized such that the sum of the discrete points is 1.0.

$$P_w(\theta_\infty,\gamma) = \int_{-\pi}^{\pi}\int_{-\pi}^{\pi} \rho\left(\Delta_\theta,\Delta_\gamma\right)\sum_{i}^{N_t} P^i_{FLORIS}(\theta+\Delta_\theta,\gamma+\Delta_\gamma)\,d\Delta_\theta\, d\Delta_\gamma \tag{8.5}$$

$$N = N_\theta \cdot N_\gamma \tag{8.6}$$

The PDFs are used to obtain a weighted average of the wind farm power for combinations of wind direction and yaw position. The weighted average of power is defined as equation (8.5) [116], where the superscript

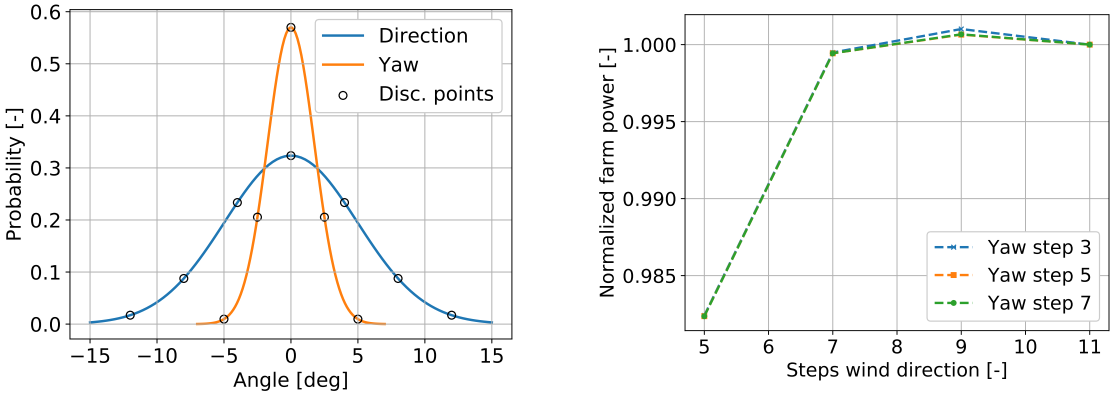

(a) PDFs

(b) Convergence

Figure 8.3: The probability density functions of the wind direction and yaw position, including the used discretization points for the weighted average. Also, the weighted average of the farm power for different combinations of discretizations is shown.

i denotes the turbine number. Note that $\rho\left(\Delta_\theta, \Delta_\gamma\right)$ represents the joint probability distribution of wind direction and yaw position and that Δ_θ and Δ_γ represent deviations of the wind direction and yaw positions from their mean values. Equation (8.5) is implemented by approximating the integration as a double summation and discretizing the wind direction and yaw position.

Determining the weighted average of the farm power takes longer than when the farm power is calculated using the deterministic approach. This is because the model needs to be evaluated multiple times with different combinations of wind directions and yaw positions. The number of evaluations needed per P_w, N, depends on the number of discretization points of the wind direction PDF N_θ and yaw position PDF N_γ, through equation (8.6). A convergence study is done to determine the number of steps needed per PDF in figure 8.3a, to make sure that the weighted average calculation is as fast as possible, without losing too much accuracy.

The convergence study is done by evaluating P_w for different combinations of N_θ and N_γ, which range from 5 to 11 and 3 to 7 respectively. The step size for the number of steps is 2 to check odd numbers only because, in that way, the peak of the PDF is always evaluated. The convergence of P_w for $U_\infty = 8$ m/s and $\theta_\infty = 80°$, corresponding to the largest gain in figure 8.2, can be seen in figure 8.3b. Note that P_w is normalized with the value corresponding to the largest N_θ and N_γ combination. From this graph can be seen that for $N_\gamma > 5$, barely anything changes. For $N_\theta > 7$, some fluctuation can be observed, but in general, the power is almost converged. The same behavior is obtained for different combinations of wind speeds and directions. Therefore, from now on, $N_\theta = 7$ and $N_\gamma = 5$ are used when the weighted average is calculated. The corresponding discretization points are indicated in figure 8.3a. Note that the numbers of discretization points are similar to what is used by Doekemeijer et al. [27].

Ideally, the weighted average is used during the yaw angle optimization, to find more robust yaw angles. As mentioned before, robust means that the yaw angles perform well, even for conditions that they are not optimized for or when there is much variability in the input conditions. To perform this optimization in a similar way as described in section 8.1 is computationally expensive. Therefore, first, a different method is chosen and explained here. Instead of taking the weighted average into account during the optimization, the performance of the baseline and optimized yaw angles are tested with the weighted average approach. This way, the robustness of the yaw angles is tested, rather than increased.

So the farm powers using the baseline and optimized yaw angles are calculated, but now with the weighted average approach. This is done for all U_∞ and θ_∞ combinations in figure 8.1. This way, the weighted average AEP can be compared to the deterministic AEP, determined in section 8.2. The results are analyzed and discussed in section 8.4.2. To test the sensitivity of σ_θ and σ_γ, the results obtained are also compared to the results of the same method, but then $2\sigma_\theta = 9.9°$ and $2\sigma_\gamma = 3.5°$ are used. Note that a similar convergence study is done for these standard deviations and that now $N_\theta = 9$ and $N_\gamma = 7$ are found. From now on the following

Table 8.3: The three different input condition cases considered.

Case	Standard deviation [deg]
$\sigma_{w,0}$	0.0
$\sigma_{w,1}$	5.25
$\sigma_{w,2}$	10.5

two symbols are used when talking about the weighted average case for one standard deviation $\sigma_{w,1} = [\sigma_\theta \ \sigma_\gamma]$ and for the two standard deviations case $\sigma_{w,2} = [2\sigma_\theta \ 2\sigma_\gamma]$. The input condition cases are shown once more in table 8.3 for clarity. Note that a normally distributed PDF with a standard deviation of 0° essentially becomes a Dirac delta, hence $\sigma_{w,0}$ represents the deterministic case.

The second method is focused on increasing the robustness of the optimized yaw angles by actually performing the yaw optimization using the weighted average. Such an optimization is called OUU. This is first done for the combinations of U_∞ and θ_∞ per turbine spacing corresponding to the highest energy gains observed in figure 8.2 to evaluate its potential. This means that OUU is performed for the combinations shown in table 8.5. To test the sensitivity of $\sigma_{w,1}$, the results obtained are compared to OUU with $\sigma_{w,2}$. So in total four sets of yaw angles are produced, one is the baseline case where all turbines have $\gamma = 0°$ indicated as $\vec{\gamma}_{base}$, one with the deterministic yaw optimization indicated as $\vec{\gamma}_{w,0}$, one with the yaw OUU using $\sigma_{w,1}$ indicated as $\vec{\gamma}_{\sigma_{w,1}}$ and one with the yaw OUU using $\sigma_{w,2}$ indicated as $\vec{\gamma}_{\sigma_{w,2}}$. The gains corresponding to three different input uncertainties are calculated for each yaw set, to compare their robustness. Afterward, OUU is performed for every combination in the windrose, similar to the deterministic case, to find a more robust yaw angle set for every combination and improve the energy gain with input uncertainty.

8.4.2. Weighted average results

The comparison between the deterministic and weighted average AEPs is not trivial, since not only the optimized powers change using the weighted average, but also the baseline power and the power when there are no wake effects. The easiest way to compare the three cases is to compare the three total AEPs and energy gains, which are shown in table 8.4. Also, a graph similar to figure 8.2 is shown in figure 8.4, although they can not directly be compared. The best direct comparison per wind direction for the deterministic and weighted average cases can be obtained by comparing the energy gain graphs, which is done in figure 8.5.

Table 8.4: The annual energy productions for three different methods and their corresponding wake losses and mean energy gains.

Method	Case	AEP [GWh]	Wake loss [%]	Gain [%]
$\sigma_{w,0}$	Baseline	379.0	22.4	0
	Optimized	401.9	17.7	6.1
	No-wake	488.3	-	-
$\sigma_{w,1}$	Baseline	378.6	22.4	0
	Optimized	388.0	20.5	2.5
	No-wake	488.1	-	-
$\sigma_{w,2}$	Baseline	376.9	22.7	0
	Optimized	381.2	21.8	1.1
	No-wake	487.5	-	-

Starting with the results shown in table 8.4. It can, first of all, be seen that indeed the AEPs of the baseline and no-wake case are not the same, though the differences in especially the no-wake case are small. This is as expected since only the yaw uncertainty influences the AEP in that case. Note that the power-yaw loss coefficient has the free stream value $\alpha = 1.7$ for all turbines for the no-wake case. More importantly, it can be seen that the optimized AEP is much lower for the weighted averages cases. This has as a logical consequence that also the mean energy gains are much lower. They decrease from 6.1% [5.5%,6.6%] for the deterministic case, to 2.5% [1.8%,3.1%] when $\sigma_{w,1}$ is used for the weighted average and 1.1% [0.5%,1.8%] when $\sigma_{w,2}$ is used. The differences in gains clearly show that these gains are very sensitive to the certainty of the input conditions.

On the other hand, even with large uncertainty in the input conditions, the lowest energy gain in the CI is still positive. From these observations can be concluded that the optimized yaw angles are robust enough to have a positive impact on the total AEP, but there are probably more robust yaw angles that result in an even higher energy gain. That is an exciting conclusion since it strongly increases the feasibility of using yaw-based wake steering to obtain energy gains.

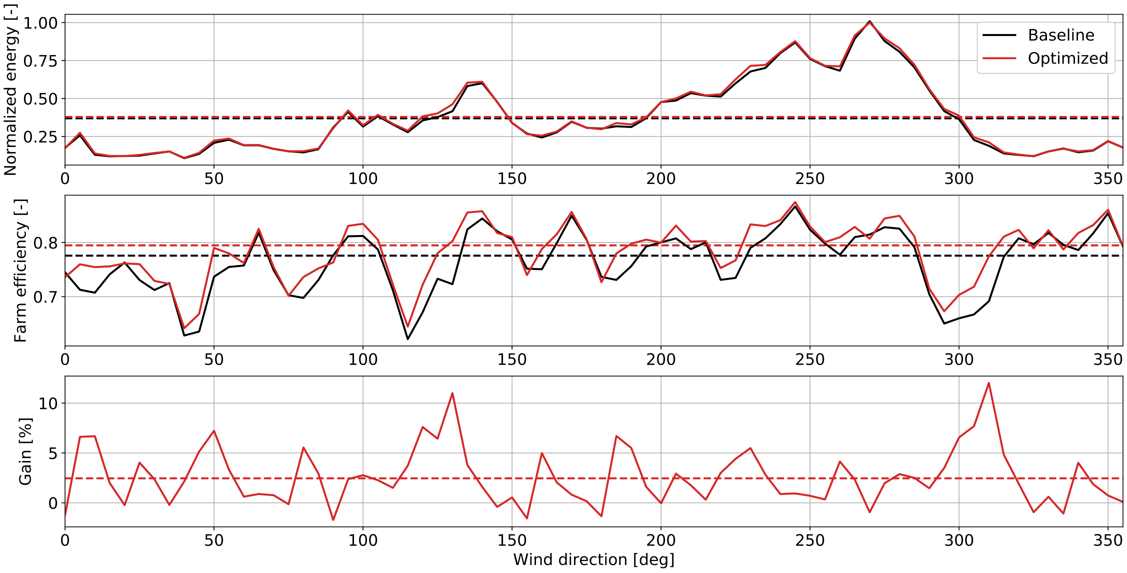

Figure 8.4: The normalized energy, wind farm efficiency and energy gain per wind direction when wake steering is applied using the weighted average.

The results corresponding to $\sigma_{w,1}$ are further analyzed using the graph shown in figure 8.4. It shows similar behavior compared to what is observed in the deterministic case in figure 8.2. However, the highest peaks in energy gain are observed for different wind directions; the peaks are generally smaller and sometimes even negative. Interestingly, the negative energy gains occur for (close to) alignment directions, e.g. 90°, 155° and 180°. This indicates that for some alignment wind directions, it is better to use zero yaw angles than optimized yaw angles and that locally not all optimums are found yet. These negative gains also show that the use of yaw-based wake steering can sometimes fail its objective, which is also found by Rott et al. [109]. Both these things have as a result that the total AEP gain is lower than for the deterministic case, although a positive energy gain is still achieved.

It is shown in table 8.4 that not only the optimized AEP changes but also the baseline and no-wake AEP. The direct comparison between the deterministic and weighted average cases is, therefore, done by comparing the energy gains per wind direction of all three cases. This is done in figure 8.5. Note that the differences in the lower plot are calculated by subtracting the energy gains of the deterministic case from the weighted average cases. Hence the values are absolute values expressed in percentages rather than the percentage of the deterministic case.

The upper plot of figure 8.5 shows once more that the weighted average cases have lower and more smoothed peaks than the deterministic case. It can also be seen that for $\sigma_{w,2}$ negative gains occur for more wind directions and are more negative than for $\sigma_{w,1}$. Furthermore, the lower plot shows that the difference is generally more negative for $\sigma_{w,2}$ than for $\sigma_{w,1}$. However, for some wind direction bins, the difference of $\sigma_{w,1}$ is actually more negative. This might have to do with the fact that for $\sigma_{w,2}$ the deviations Δ_θ and Δ_γ sometimes become so large, that the optimized yaw positions might already become beneficial for the next alignment direction. In general, $\sigma_{w,2}$ comes with much more negative gains and shows that wake steering can often fail its objective. However, in total, the yaw angles are robust enough to achieve a positive energy gain.

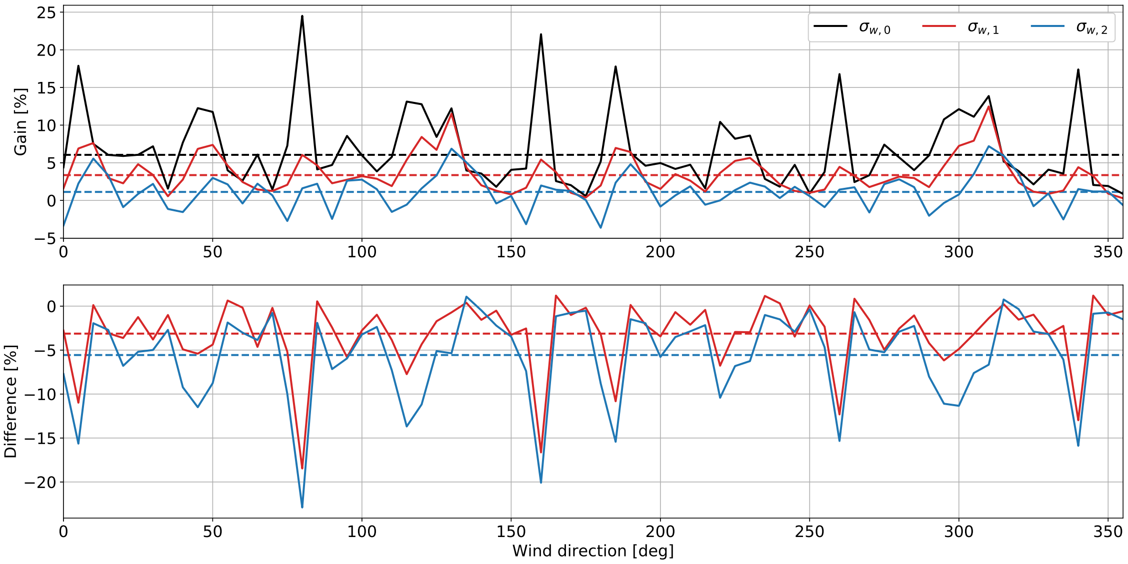

Figure 8.5: The energy gain per wind direction for the deterministic and weighted averages cases. The difference between the weighted average cases and the deterministic case is also shown.

8.4.3. Optimization under uncertainty

OUU takes much time when all possible combinations of U_∞ and θ_∞ in the wind rose are considered. Therefore, at first only cases for which the highest energy gains in the deterministic case are observed are chosen to estimate the potential of OUU. In total, six cases are considered, one for each relevant turbine spacing. The cases can be found in table 8.5. Note that in this section, the farm power is calculated for these cases, rather than the AEP. However, the gain is identical, since both the baseline and optimized power are multiplied by the same occurrence fraction obtained from the wind rose.

Table 8.5: The cases which will be optimized under uncertainty. The gain corresponds to the deterministic case.

Case [-]	**Distance [D]**	**Wind direction [deg]**	**Wind speed [m/s]**	**Deterministic gain [%]**
1	3.3	310	9.0	23
2	4.3	45	9.0	19
3	4.8	5	8.0	30
4	5.9	80	8.0	44
5	7.1	160	9.0	39
6	8.6	95	9.0	18

Figure 8.6a shows how the powers change when they are optimized under uncertainty. First of all, it can be observed that the baseline power increases with increasing uncertainty. That makes sense since they are calculated with input uncertainty. This means that also wind directions are considered for which turbines are partly waked instead of fully, resulting in a higher weighted average baseline power than the deterministic power. It is surprising to see that, for this case, the optimized power is the lowest for $\sigma_{w,1}$ rather than for $\sigma_{w,2}$.

The most interesting are, of course, the power gains. In figure 8.6a can be seen that they decrease significantly. Also can be seen that the difference between $\sigma_{w,1}$ and $\sigma_{w,2}$ is small compared to the difference between them and $\sigma_{w,0}$. So in total, there are now four sets of yaw angles per case: $\vec{\gamma}_{base}$, $\vec{\gamma}_{\sigma_{w,0}}$, $\vec{\gamma}_{\sigma_{w,1}}$, $\vec{\gamma}_{\sigma_{w,2}}$. To avoid a chaotic plot, only $\vec{\gamma}_{base}$, $\vec{\gamma}_{\sigma_{w,0}}$ and $\vec{\gamma}_{\sigma_{w,1}}$ are shown for case 4 in figure 8.7. The upper value per turbine corresponds to $\vec{\gamma}_{\sigma_{w,0}}$ and the lower to $\vec{\gamma}_{\sigma_{w,1}}$. In general, it can be seen that the yaw angles of $\vec{\gamma}_{\sigma_{w,1}}$ are lower than $\vec{\gamma}_{\sigma_{w,0}}$. The most downstream turbines generally have a slightly larger yaw angle for $\vec{\gamma}_{\sigma_{w,1}}$. Similar behaviour is observed by Rott et al. [109]. In this paper is also observed that more robust yaw angles result in

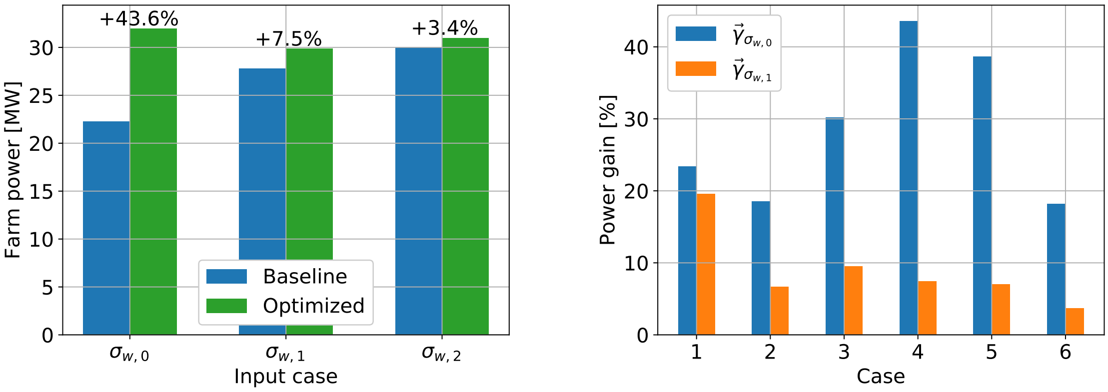

(a) Farm powers

(b) Gain comparison

Figure 8.6: The baseline and optimized power and the corresponding gain are shown per method for case 4. Also, the gains for all cases corresponding to $\vec{\gamma}_{\sigma_{w,0}}$ are compared with the gains corresponding to $\vec{\gamma}_{\sigma_{w,1}}$, each calculated with the input conditions they are optimized for.

a lower yaw activity, which is beneficial for especially the control of the turbine.

However, the powers observed in figure 8.6a correspond to the input uncertainty for which they are optimized. Hence, they do not tell much about the actual robustness of the yaw angles. To test the robustness of the different yaw angle sets, it is also investigated how they perform under two different input cases. These two cases are $\sigma_{w,1.5} = 7.88°$ and $\sigma_{w,2.5} = 13.13°$. The gain is given per yaw set. The results for case 4 are shown in table 8.6. Unsurprisingly, the highest gain is found for $\vec{\gamma}_{\sigma_{w,0}}$ with the input case $\sigma_{w,0}$. However, $\vec{\gamma}_{\sigma_{w,0}}$ shows the lowest gains for the other two input cases. Although $\vec{\gamma}_{\sigma_{w,2}}$ has a slightly higher gain for the $\sigma_{w,1.5}$ input case, in general the highest gains are achieved with $\vec{\gamma}_{\sigma_{w,1}}$. This observation strengthens the motivation to apply OUU to find better and more robust optimal yaw angles. It also shows that the choice of $\sigma_{w,1}$ used for OUU is non-trivial and requires further study.

Table 8.6: Table showing the mean gains per yaw set for different input conditions for case 4.

Input case	**Gain $\vec{\gamma}_{\sigma_{w,0}}$ [%]**	**Gain $\vec{\gamma}_{\sigma_{w,1}}$ [%]**	**Gain $\vec{\gamma}_{\sigma_{w,2}}$ [%]**
$\sigma_{w,0}$	43.6	39.7	29.0
$\sigma_{w,1.5}$	1.9	3.3	3.4
$\sigma_{w,2.5}$	6.5	7.5	6.7

For case 4, it is found that OUU with $\sigma_{w,1}$ results in the most robust yaw angles for the cases considered. The same analysis is done for all other cases, and it consistently showed that the yaw angle set obtained through OUU with $\sigma_{w,1}$ outperforms the other two when all three input cases are considered. Assuming that the uncertainty in input conditions is correctly described when $\sigma_{w,1}$ is used, it means that the gains observed for $\vec{\gamma}_{\sigma_{w,1}}$ with input conditions $\sigma_{w,1}$ give the most reliable gains. Therefore, these gains are compared to the deterministic gains for all cases in figure 8.6b.

Figure 8.6b shows the gains for $\vec{\gamma}_{\sigma_{w,0}}$ with input conditions $\sigma_{w,0}$ and $\vec{\gamma}_{\sigma_{w,1}}$ with input conditions $\sigma_{w,1}$ per case. Note that each case corresponds to a unique turbine spacing, where case 1 has the smallest turbine spacing and case 6 the largest. It can be seen that the gains of $\vec{\gamma}_{\sigma_{w,1}}$ are lower than for $\vec{\gamma}_{w,0}$, ranging from 16% to 83% lower. These are significant reductions that strongly suggest that the total energy gain is indeed much lower than the 6.1% observed in section 8.3. However, it is probably possible to obtain gains that are slightly higher than the 2.5% observed in table 8.4 by finding the optimal $\vec{\gamma}_{\sigma_{w,1}}$ for all combinations of U_∞ and θ_∞ in the wind rose.

A script is written to find the optimal $\vec{\gamma}_{\sigma_{w,1}}$ for all combinations of U_∞ and θ_∞ in the wind rose. It calculates all yaw sets per wind speed for one wind direction at the time. Using parallel computing and 24 cores (Dual 12-core 2.6GHz Intel Xeon processors) on the DTU cluster, it takes between 12-24 hours per wind di-

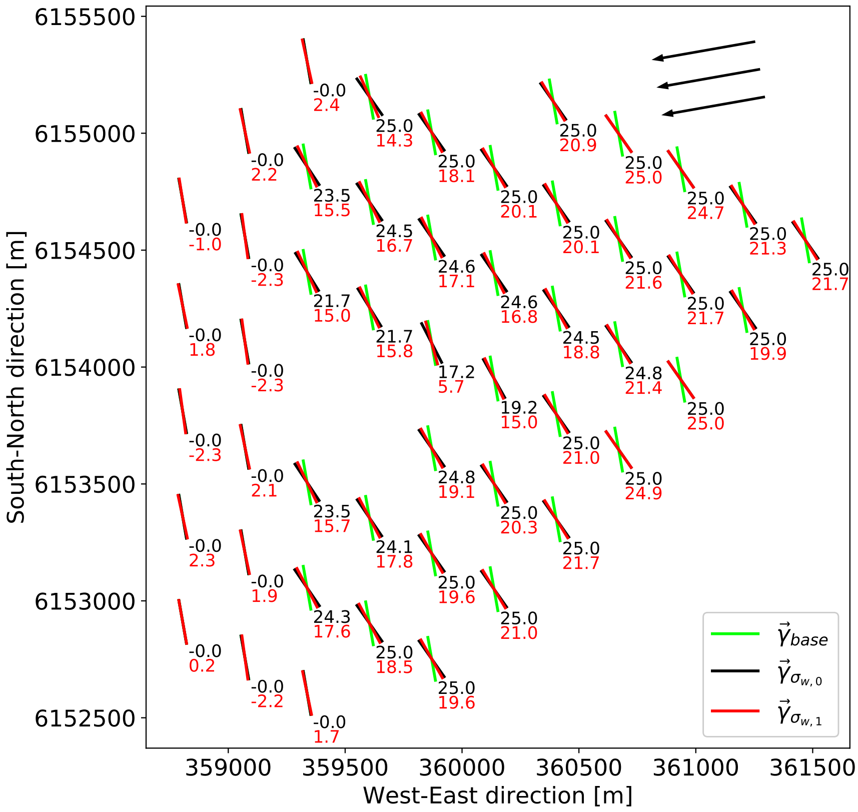

Figure 8.7: Three sets of yaw angles are shown. The upper value per turbine indicates the deterministic yaw angle and lower value the robust yaw angle. The wind direction is indicated with black arrows.

rection. The calculation is done for a total of 72 wind directions, and it was possible to do between five to seven wind directions at a time, depending on the available nodes in the cluster. The total calculation time was, therefore, about ten days. Note that when there are more nodes available, the number of days can be reduced. This is mentioned to emphasize that even with a fast low-fidelity model, it is very computationally expensive to obtain robust yaw angles.

Table 8.7: Gains per input condition case for the deterministic yaw angles set and the robust yaw angle set.

Input case	**Gain $\vec{\gamma}_{\sigma_{w,0}}$ [%]**	**Gain $\vec{\gamma}_{\sigma_{w,1}}$ [%]**
$\sigma_{w,0}$	6.1	4.3
$\sigma_{w,1}$	2.5	3.4
$\sigma_{w,2}$	1.1	2.2

Although it takes a lot of computational effort, the results are worth it. The gains for $\vec{\gamma}_{\sigma_{w,0}}$ and $\vec{\gamma}_{\sigma_{w,1}}$ with input case $\sigma_{w,1}$ are compared in figure 8.8. It can be seen that where $\vec{\gamma}_{\sigma_{w,0}}$ has locally negative gains, $\vec{\gamma}_{\sigma_{w,1}}$ shows positive gains throughout. When looking at the differences in the lower plot it can be seen that the highest differences are observed for alignment directions (90°, 155°, 180°, etc.). The total AEP gains for $\vec{\gamma}_{\sigma_{w,0}}$ and $\vec{\gamma}_{\sigma_{w,1}}$ for three different input cases are shown in table 8.7. It can be seen that with $\vec{\gamma}_{\sigma_{w,1}}$ higher gains are observed for $\sigma_{w,1}$ and $\sigma_{w,2}$, while it is lower for $\sigma_{w,0}$. The hypotheses that the gain for $\sigma_{w,1}$ could be higher than 2.5% is found to be correct. The final estimation of the possible energy gain for the Lillgrund wind farm is 3.4% [2.8%,4.0%] using $\vec{\gamma}_{\sigma_{w,1}}$ and $\sigma_{w,1}$ as input uncertainty.

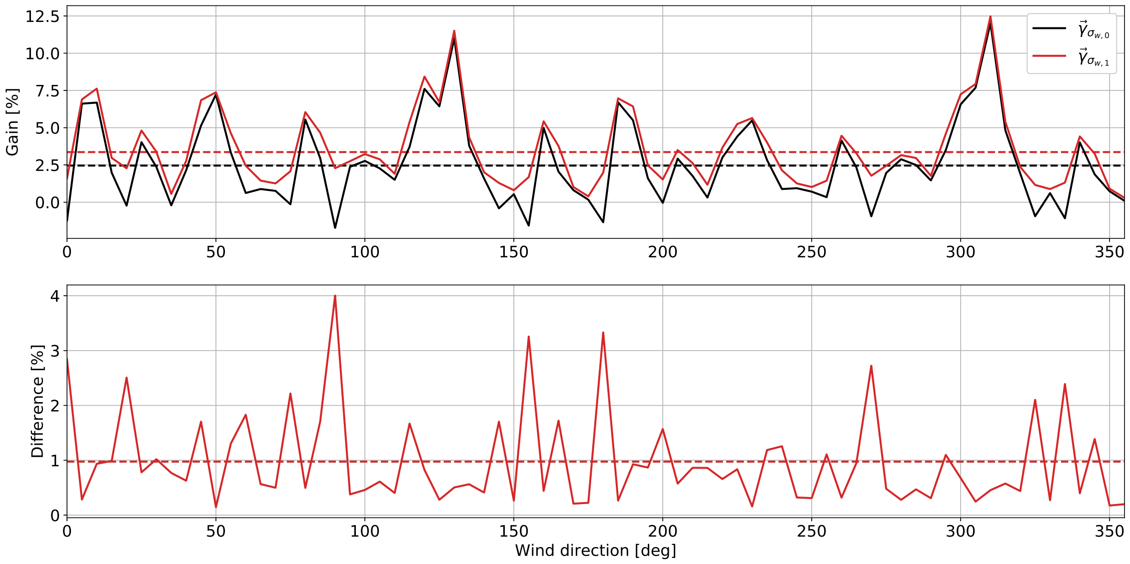

Figure 8.8: The energy gain per wind direction for $\vec{\gamma}_{\sigma_{w,0}}$ and $\vec{\gamma}_{\sigma_{w,1}}$ with input case $\sigma_{w,1}$. The difference between the gains is also shown.

8.5. Discussion

This section couples and discusses the findings of this chapter in section 8.5.1. These findings are then critically reviewed in section 8.5.2. Next, it is analyzed how the gains are divided over the turbines in the wind farm in section 8.5.3. Finally, a short discussion is presented in section 8.5.4, about the effect on the turbine loads when a turbine operates with a yaw misalignment angle.

8.5.1. Findings

First of all, it is found that the baseline AEP predicted by the model is close to what is obtained through SCADA data, although when the bias is applied, it is on the high side. Finding the optimal yaw angles for many combinations of the wind rose takes significant computation time because there are 48 turbines and hence 48 variables to optimize. It is found that a warm start helps to decrease the computation time. With the optimal yaw angles in place, a total AEP gain of 6.1% [5.5%,6.6%] can be achieved using the input case $\sigma_{w,0}$. The most significant contributions to this gain are from wind directions related to small spacings. However, the highest peaks correspond to other turbine spacings and wind directions, which are the alignment direction plus 5°. This indicates that yaw-based wake steering is especially effective for partly waked situations.

It is found that the energy gain corresponding to a yaw set depends heavily on the input conditions. The more uncertainty is present in the input conditions, the lower the energy gain found. This is both a result of the baseline AEP going up and the optimized AEP going down. This, in turn, is a result of the weighted average method that not only takes fully/partly waked situations into account but also partly/non-waked situations, for wind directions that, with the deterministic method, would only result in fully/partly waked situations. Hence, the gains go down and can even become negative for certain wind directions. This undesirable behavior and should be avoided as much as possible. It should be noted that the weighted average method decreases uncertainties in the analysis and thus increases the reliability of the gains observed.

Since $\vec{\gamma}_{\sigma_{w,0}}$ is found using optimization with the deterministic method, it is believed that it is not the optimal yaw set for uncertain input conditions. A more robust yaw set can be obtained when OUU is applied, although it takes significant computation time. It is shown in table 8.6 for one specific case that indeed higher gains can be achieved for yaw sets found through OUU, even for uncertainties for which they are not optimized. These yaw sets generally have lower yaw angles than the deterministic yaw set, and from the literature, it is known that this results in less yaw activity [109]. Based on these findings, a robust yaw set $\vec{\gamma}_{\sigma_{w,1}}$ is found for every combination of U_∞ and θ_∞ in the wind rose. In figure 8.8 is shown that while $\vec{\gamma}_{\sigma_{w,0}}$ shows negative gains for certain wind directions, $\vec{\gamma}_{\sigma_{w,1}}$ shows positive gains for all wind directions. Next to that can be observed in table 8.7 that $\vec{\gamma}_{\sigma_{w,1}}$ results in higher gains for the uncertain input cases. Based on this analysis can be concluded that the potential energy gain for the Lillgrund wind farm when applying yaw-based wake steering is likely to be around 3.4% [2.8%,4.0%]. Hence, achieving an energy gain using yaw-based wake steering is feasible for the Lillgrund wind farm.

8.5.2. Validity of findings

While the entire thesis is focused on quantifying and, if possible, mitigating the uncertainties in the analyses as much as possible, there is still significant uncertainty in the observed energy gain of 3.4% [2.8%,4.0%]. First of all, the FLORIS model is a low-fidelity model that relies on many simplifications (e.g. the steady-state assumption, near-wake modeled as a cone, velocity deficit modeled as 2D Gaussian shape) to calculate the flow field and wake deflection. The advantage of those simplifications is that the model can calculate the flow field quickly; the downside is that it is not as accurate as higher-fidelity models and that there are tuning parameters that need to be calibrated. This calibration adds uncertainty when turbines are operating outside the operating regions used for the calibration. Besides, two tuning parameters are not calibrated as they were added to the model in a late stage of the thesis.

Secondly, it has been shown in section 7.3 that the model output for the baseline case has a significant bias, while the uncertainty is low. The influence of the bias on the energy gain is negligible according to the model; however, this only holds under the assumption that the same bias and uncertainty can be used for the optimized case. On top of that, the bias and uncertainty are calculated with SCADA data as a benchmark, which comes with uncertainty in itself due to miscalibrated sensors and noise in the measurements.

Thirdly, uncertainties in the input conditions are used to find more robust yaw angles in section 8.4.3. However, this uncertainty is assumed to be well described by $\sigma_{w,1}$, which is found by Simley et al. [116] to describe the wind direction and yaw error uncertainty accurately. For the Lillgrund wind farm, $\sigma_{w,1}$ is assumed to be correct as well, but this is not necessarily the case. Next to that, only the uncertainty in wind direction and yaw error is taken into account, while there is also uncertainty in other inputs such as wind shear, wind veer, and C_T.

Finally, it is assumed that the tuning parameters obtained for the baseline case can also be used for the optimized case. That adds significant uncertainty to the outcome as well. When this added uncertainty is accepted, the FLORIS model proves to be very useful to find robust, optimized yaw angles. However, when a higher certainty of the output is desired, high-fidelity models might be more useful, e.g., PALM, SOWFA, or EllipSys3D (see section 2.1.3). A way to test the reliability of the found gains can be to run one of these high-fidelity models for a few test cases with the yaw angle sets obtained through the FLORIS model and compare the obtained gains. However, this would be a significant study on itself, since these high-fidelity models are very computationally expensive and have their own uncertainties as well.

This subsection highlights the fact that this research relies on multiple significant assumptions and that, therefore, the observed results should be used with caution.

8.5.3. Turbine gains

One part, which is not investigated yet, is how the gains are per individual turbine. This is analyzed by plotting the flow field for case 4 and indicate per turbine the gain achieved compared to the baseline yaw set and underneath the used yaw angle to achieve this gain. The plot can be seen in figure 8.9. The light green bars indicate the baseline yaw position per turbine, perpendicular to the inflow direction, and the black bars indicate the optimized yaw position. As expected, the large yaw misalignment angles of the upstream turbines result in a negative gain ranging from -6 to -16%.

However, the first downstream turbines, although having yaw angles as well, immediately show positive gains ranging from 8% in the middle of the farm to 33% at the edge of the farm. The turbine at the edge of the farm does not have to steer the wake because there is no downstream turbine; hence a significant gain is achieved. Similar gains are observed for the next downstream turbine, up to the last downstream turbine per row. These turbines at the end of the row show only small yaw angles, probably aligning themselves with the slightly deviated local wind direction. They have substantial energy gains, ranging from 11 to 38%, except for the turbines behind the gap in the farm. Note that these results correspond to only one inflow case with relatively significant gains. The behavior of both the gains and yaw angles are different for other inflow conditions. However, the general trend is that negative gains are observed for the upstream turbines and positive gains for all other turbines.

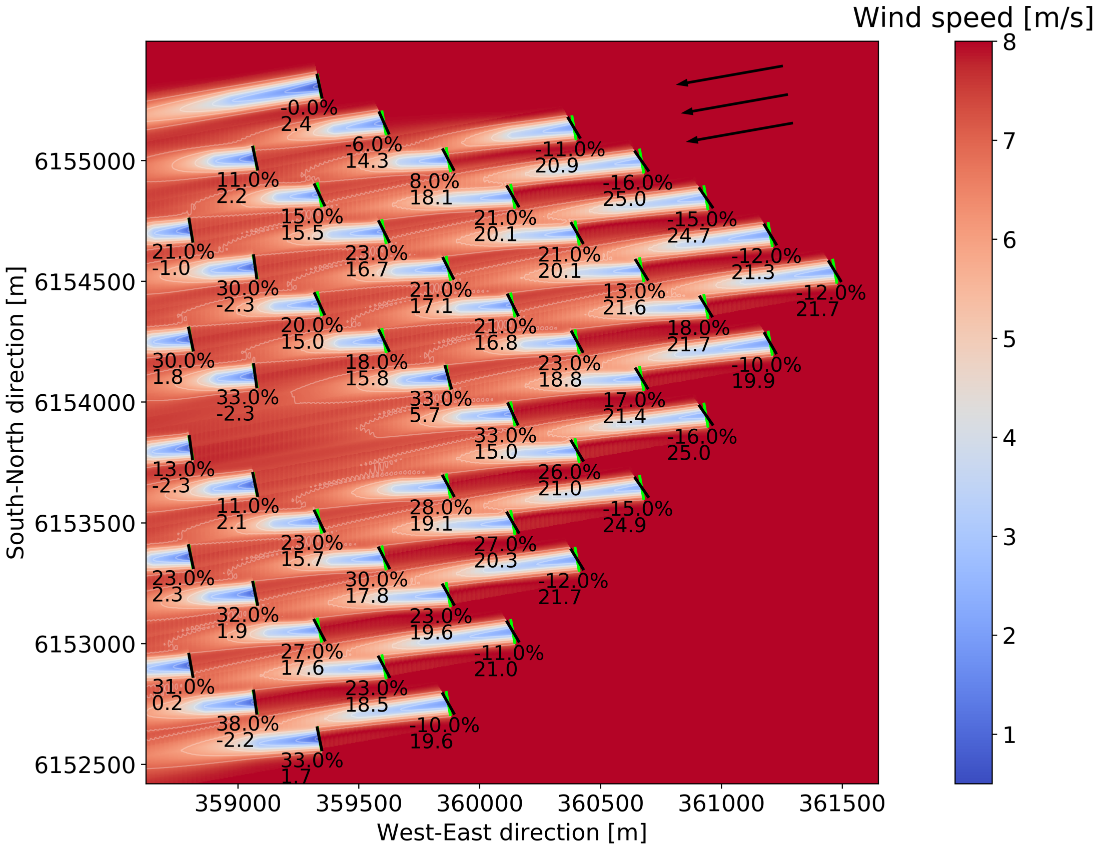

Figure 8.9: The flow field, the gains (upper value) and yaw angles (lower value) corresponding to $\vec{\gamma}_{\sigma_{w,1}}$ per turbine are shown for case 4, $\vec{\gamma}_{base}$ are shown in light green. The wind direction is indicated with black arrows and the turbine diameter is doubled in size for clarity.

8.5.4. Turbine loads

Another factor so far left out in this thesis is the turbine loads. A wind turbine is generally well able to handle extreme loads [119], but it is considered to be a fatigue-critical machine [123]. Therefore the focus will be on fatigue loads only. The fatigue loads on a turbine are different when it operates in a wake compared to when it operates in free stream air. Also, when the turbine operates with a yaw misalignment angle, the loads change. That means that the energy gain obtained through wake steering might come with the cost of a reduced turbine lifetime.

Starting with a turbine operating in the wake of another upstream turbine. Two effects of the wake influence the turbine loads. First of all, the flow in a wake has increased turbulence, which means that the fatigue loads on a turbine increase [16, 83]. This is, in turn, shortens the lifetime of the turbine. The second effect happens when the center of the wake does not align perfectly with the downstream turbine, e.g., due to wake meandering, resulting in the downstream turbine operating in a partly waked situation. That, in turn, results in more thrust being produced at one side of the rotor than the other side of the rotor. This results in large cyclic variations because the blades pass in an out of the wake [126, 129]. The large cyclic variations can accelerate the structural degradation of waked turbines.

Turbines operating with a yaw misalignment angle experience different fatigue loads in their components than aligned turbines. The influence of the yaw angle is non-trivial, and the studies so far are not conclusive. Some studies show increased blade-root fatigue loads for negative yaw angles and decreased fatigue loads for positive yaw angles [15, 34]. The reduction can be a result of the yaw misalignment counteracting the effect that wind shear has on the cyclic loads of the blades for specific wind conditions [35].

However, Damiani et al. [25] found through aero-elastic simulations that not all fatigue load components necessarily decrease with increasing yaw angles. Although they too observed that fatigue loads of the blade-root decreased for positive yaw angles, large variations were found as well, making generalizations difficult. Also, more complex trends are observed for other components such as shaft torque and tower-top bending. Whether the fatigue loads increased or decreased for the shaft torque seemed to depend on wind speed and control dynamics. On top of that is found that the aero-elastic simulation results did not always match field-derived data. For example, according to the aero-elastic simulation, the tower-top bending is not influenced by yaw misalignment. However, field-derived data showed a sharp increase in fatigue loads for yaw misalignment.

The general trend between the different studies is that the fatigue loads of wind turbine components change, but whether they increase or decrease is so far inconclusive and highly dependent on specific conditions such as the wind and control dynamics. This means that although yaw-based wake steering shows potential in terms of energy, the technique should be used with caution when loads are considered. For future research, a field study with turbines operating with intentional misalignment would be preferable to get a better grip on the effects of yaw misalignment on the loads.

9

Conclusion and recommendations

This chapter aggregates the conclusions drawn throughout the report to arrive at one general conclusion in section 9.1. After that, using the insights obtained from this thesis, recommendations for further studies are given in section 9.2.

9.1. Conclusion

When clustering turbines, one of the main disadvantages is that aerodynamic interaction can occur between the turbines, resulting in so-called wake losses. One possible control technique to reduce wake losses is called yaw-based wake steering. The main objective of this thesis is to estimate the feasibility of achieving an energy gain by optimizing the use of yaw-based wake steering on an existing wind farm, by calibrating a parametric wind farm model on measured data and taking input and model uncertainty into account.

The wind farm model can be calibrated by adjusting six tuning parameters. These tuning parameters influence three typical properties of the wake: expansion, deflection, and velocity deficit. A global sensitivity analysis revealed that the influences of the tuning parameters differ significantly. It also showed that one tuning parameter, a_d, can be left out in future calibrations, as its influence on the model output is found to be negligible. This removal brings the number of tuning parameters down to five, decreasing the optimization time and the risk of over-fitting during the calibration.

The common way of calibrating the wind farm model is by using flow field data or power predictions found through computationally expensive, high-fidelity models. This thesis proposes a new calibration method using conventionally measured data of a wind farm only, with a focus on the power measurements. It should be noted that the new method is only useful for existing wind farms and only if enough data is available. During the development of this new calibration method, it turned out that the tuning parameters depend on both the turbine spacing and wind speed. This is not previously shown in the literature because in those calibrations, the inflow conditions, and hence turbine spacings considered, are limited. These findings highlight the fact that the wind farm model is not an off-the-shelf product.

After calibration, the model's performance is tested by comparing wake losses and annual energy productions with measured data as a benchmark. It is found that the model can predict the wake losses well for most wind directions. However, a negative model bias of -5.8% is present, meaning that the model generally overestimates the wake losses. This bias can potentially be reduced if the effect of the atmospheric stability on the wind shear and veer is taken into account. Finally, although the model has a significant bias, the related model uncertainty is only 1.1%, making the model evaluation a reliable estimate.

With the assumption that the bias and uncertainty found are valid for both the yawed and non-yawed cases, it is possible to optimize the yaw angles and assess the potential energy gain. The 68% confidence interval of the energy gain potential using the deterministic approach is 5.5-6.6%. It has been shown that applying a correction for the model bias has a negligible effect on the energy gain. The deterministic approach assumes perfect knowledge of the wind direction and control of the yaw position. From the literature, it is known that both these input conditions come with uncertainty. The robustness of the deterministic optimal yaw angles is tested by taking into account these input uncertainties. It is found that the energy gain potentially decreases to a minimum of 0.5-1.8% for the uncertainty cases considered and that for some wind directions, the optimal yaw angles even result in a negative gain. This means that the technique sometimes fails its objective and indicates low robustness of the deterministic yaw angles.

It is possible to increase the robustness of the yaw angles when the input uncertainty is taken into account during the yaw angle optimization. It is a computationally expensive task to do this for a full wind rose. Hence, first, a case study is done to show that the robustness of the yaw angles can indeed be increased by applying optimization under uncertainty. However, the level of uncertainty influences the robustness of the yaw angles, and its determination is non-trivial. A proper estimate of the uncertainty in input conditions will increase the robustness and potential energy gain. It can be concluded that an energy gain between 1.8-3.1% can be achieved for the deterministic yaw angle set and with the most realistic level of input uncertainty considered. This can be further increased by finding more robust yaw angles using optimization under uncertainty. The more robust yaw angles set results in an energy gain between 2.8-4.0%, which is an increase of 36% compared to the 1.8-3.1% for the deterministic yaw angle set, emphasizing the importance of optimization under uncertainty. The gains are in the same order as found in the literature.

The observed energy gain of 2.8-4.0% shows an exciting potential for wake steering, though the effect of yaw misalignment on the turbine loads should be kept in mind. The effect of wake steering on the turbine loads is an ongoing field of research, and the first results indicate that yaw misalignment decreases the lifetime of most components in the turbine; hence wake steering comes with a cost in loads. Although the lifetime of the wind farm might decrease, the main conclusion of this thesis is that achieving an energy gain using yaw-based wake steering is feasible for the Lillgrund wind farm.

9.2. Recommendations

The feasibility of wake steering for the Lillgrund wind farm does not mean that it is feasible for all wind farms. Lillgrund is a relatively old and dense wind farm. The control technique's potential decreases when turbine spacings are larger, which is often the case for modern wind farms. Hence, this thesis does not prove the feasibility of yaw-based wake steering for all wind farms. Instead, it provides a guideline to calibrate the model and assess the feasibility for existing wind farms, using Lillgrund as a test case.

However, when this control technique proves to be feasible in reality, the layout of future wind farms can be optimized while taking into account the control technique. These layouts would likely have smaller turbine spacings, resulting in more turbines per area and thus a higher power production per area. Nevertheless, first, the feasibility of wake steering should be further assessed. Therefore, additional analyses must be done to assess the feasibility of wake steering on existing wind farms, similarly as done for the Lillgrund test case here. If it is consistently found that an energy gain through wake steering is feasible in theory, then the next step is to test it in practice using a field study on a wind farm scale.

Next to testing the feasibility using other wind farms with, for example, larger turbine spacings, the accuracy of the proposed calibration method and the model overall can be increased. Throughout the thesis, recommendations for improvements have been mentioned. The reason that these are not applied in this thesis is often due to time constraints. The recommendations are aggregated here.

First of all, only one power curve is used for all turbines, while a large spread in power can be observed per wind speed. Further research can be done into this spread, and power curves can be derived, for example, per atmospheric stability or turbulence intensity class. Next to that, the thrust curve of the manufacturer is used in this thesis. The thrust curve has a significant impact on the wake calculations. It would be interesting to see if the uncertainty related to the thrust curve can be brought down by using turbine load measurements to verify the thrust curve.

Furthermore, the model takes secondary wake effects into account, which are caused by a wake rotation vortex and two counter-rotating vortices. The latter only occur when there is yaw misalignment. These affect the wake recovery and the perceived yaw misalignment of downstream waked turbines. There are two tuning parameters related to secondary wake effects that are not calibrated in the proposed method. Hence, they should be included in the calibration method in the future.

Another issue that should be investigated to verify its use in this thesis is the power-yaw loss correction. It would be valuable to know whether this correction changes significantly for partly waked turbines compared to fully waked turbines. In this thesis, the correction for fully waked situations is also used for partly waked situations. Also, the wind farm blockage effect is mentioned multiple times. It could be relevant to know the significance of this effect on the power output and whether it can be accounted for in the model, to improve the accuracy of the output.

An important wind condition that is not accounted for is the atmospheric stability. The literature shows that this stability affects the turbulence intensity, wind shear, and wind veer. It is believed that the overall accuracy of the model will increase if these three properties are adjusted based on the stability class, especially by accurately estimating the wind shear. Taking into account the stability can also improve the estimation of wind direction variability. The better this estimation is, the more the uncertainty can be decreased. On top of that, the evolution of wind direction variability through the park is assumed to be negligible. However, the significance of this assumption should be investigated in future studies and possibly incorporated to improve the model.

Also, an assumption made is that the model bias and uncertainty is equal for both the baseline and optimized case. It should be researched whether this assumption is valid. The ideal way to do this is through a field study with yawed turbines. When such a study would be done, it is also the ideal opportunity to analyze the effects of a yaw misalignment angle on the turbine loads.

Finally, since the FLORIS model is based on many simplifications, it would be interesting to verify the obtained results using higher-fidelity models. A way to test the reliability of the found gains can be to run a high-fidelity model for a few test cases with the yaw angle sets obtained through the FLORIS model and compare the obtained gains. When similar results are found, it would strongly increase the reliability of the results and conclusion of this thesis.

It has already been extensively discussed in section 8.5. However, it should be noted that the model used has its limitations and that many significant assumptions are made during this research that influence the results. So, although those results have been calculated as accurately as possible within the framework and available time, the observed findings should be used with caution. However, it is believed that despite this, the results of this thesis have shown that it is feasible to achieve energy gains using yaw-based wake steering for the Lillgrund wind farm.

A

Newton-Raphson method

The Newton-Raphson method is a powerful technique to solve equations numerically. It is a technique based on the simple idea of linear approximation. Starting from an initial guess x_0, the root-finding algorithm produces successively better estimates to the roots using equation (A.1).

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)} \tag{A.1}$$

In equation (A.1), n is the iteration number, $f(x)$ is the equation to solve and $f'(x)$ is the corresponding derivative. The equation to be solved is first introduced in section 4.2.2 and is equation (A.2) where P is the turbine's power in W, ρ the air density in kg/m^3, R the rotor radius in m and U_{REWS} the REWS in m/s. This equation needs to be solved iteratively because $C_P(\lambda,\theta)$ (equation (A.3)) depends on λ which in turn depends on U_{REWS} (equation (A.5)). Note that in this case the variable to be found x is U_{REWS}, which is done for every available time step. Also note that the coefficients $c_1 - c_9$ can be found in table A.1.

Table A.1: The coefficients as defined in Göçmen et al. [52, 53].

Coefficient	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9
Value [-]	0.47	101	0.4	0.01	1.95	5	16.5	0.089	0.02

$$f(U_{REWS}) = P - 0.5\rho C_P(\lambda,\theta)\pi R^2 U_{REWS}^3 \tag{A.2}$$

$$C_P(\lambda,\theta) = c_1\left(\frac{c_2}{\lambda_i} - c_3\theta - c_4\theta^{c_5} - c_6\right)\exp\left(-\frac{c_7}{\lambda_i}\right) \tag{A.3}$$

$$\lambda_i = \left[\left(\frac{1}{\lambda + c_8\theta}\right) - \left(\frac{c_9}{\theta^3+1}\right)\right]^{-1} \tag{A.4}$$

$$\lambda = \frac{\omega R}{U_{REWS}} \tag{A.5}$$

The derivative $f'(U_{REWS})$ can be derived through equation (A.6) where the capital letters are defined as in equations (A.7) to (A.10) to make the equation readable. Implementation of the method can be found in van Beek [125].

$$f'(U_{REWS}) = \frac{\mathrm{d}f(U_{REWS})}{\mathrm{d}U_{REWS}} = \frac{0.5\rho ARc_1\omega E(c_7D - c_2)U_{REWS}}{B^2} - 1.5\rho Ac_1EDU_{REWS}^2 \tag{A.6}$$

$$A = \pi R^2 \tag{A.7}$$

$$B = \left(\frac{\omega R}{U_{REWS}} + c_8\theta\right) \tag{A.8}$$

$$D = c_2\left(\frac{1}{\frac{\omega R}{U_{REWS}} + c_8\theta} - \frac{c_9}{\theta^3+1}\right) - c_3\theta - c_4\theta^{c_5} - c_6 \tag{A.9}$$

$$E = \exp\left(-c_7\left(\frac{1}{\frac{\omega R}{U_{REWS}} + c_8\theta} - \frac{c_9}{\theta^3+1}\right)\right) \tag{A.10}$$

B

Step-by-step calibration

This appendix presents a summary of the proposed calibration method in the form of a step-by-step guide. All the steps performed for the Lillgrund test case are mentioned, but it might be that additional steps are needed for other data sets or that some can be skipped. Bold steps are linked to a Python script; normal steps are not. When the step is italic, it means that it is part of another data set. The data set for steps 1-4 contains 10-minute average measurements and is the main data set used. The data set is confidential, meaning that all scripts except 3c, 3d and 4d run on dummy data sets (100 random time steps). The Python scripts for steps 1-4 are available in Gitlab [125] under the folder 'Relevant Scripts'.

The data set for step 5 contains 1 Hz data and is used to derive the REWS. This data is confidential as well, meaning that all scripts are run on dummy data sets (2 out of 223 days available). The Python scripts for step 5 are available in Gitlab as well.

1. Data analysis data set 1
    (a) **Import data set per channel from the SQL database of DTU.**
    (b) Analyze data by creating a scatter plot. Delete unrealistic values.
    (c) **When relevant: apply LOF method to detect outliers.**
    (d) Obtain the free stream wind direction per time step.
    (e) *Obtain correction factor to convert nacelle wind speeds to REWS. (see setp 5a)*
    (f) **Use the correction factor to convert nacelle wind speeds to REWS.**
    (g) **Obtain the free stream wind speed per time step based on the REWS of free stream turbines. (parallel)**
    (h) **Analyze wind conditions such as a wind rose or a TI rose.**
    (i) When relevant: determine the tip-speed-ratio of the turbines.
2. Derive power curve
    (a) Obtain power curve of the manufacturer and if needed convert power to power coefficient.
    (b) Create a complete .JSON file from the example.JSON as reference and add the power and thrust curve data. The layout can be added in Python. Choose sub-models but leave tuning parameters as they are.
    (c) **Compare the predicted power of FLORIS with power measurements of free stream turbines. (parallel)**
    (d) Manually adjust the power coefficients in the .JSON file based on the discrepancy between the predicted powers and measured powers until satisfied with the power curve. (iterate with step 2c)
    (e) **Validate obtained power curve by analyzing the error for (new) parts of the data set.**
3. Derive tuning parameters
    (a) Determine the relevant turbine spacings for the wind farm.
    (b) Determine turbine rows (preferably in the middle of the park) on which the model will be calibrated per turbine spacing.
    (c) **Calibrate the tuning parameters by minimizing the average MAPE of all turbines considered.**

   (d) **Validate (when possible) the performance of the tuning parameters on the second alignment direction corresponding to that turbine spacing.**

4. Set the power-yaw loss coefficient and tuning parameters per turbine spacing
   (a) Obtain the power-yaw loss coefficient for different turbine spacing using aero-elastic software following the method of Liew et al. [84].
      i. Alternative: use figure 7.7 obtained from Liew et al. [84]. Note that the coefficient is turbine dependent.
   (b) **Create a function/script to adjust the tuning parameters, and the power-yaw loss coefficient per turbine depending on the turbine spacings and whether they are free stream or not.**

5. Obtain correction factor to convert nacelle wind speeds to REWS.
   (a) **When relevant: Convert the stored data from .mat (MATLAB) files to pkl files.**
   (b) **Investigate data and delete unrealistic or erroneous data points.**
   (c) **Calculate REWS based on power, pitch angle and rotational speed. (parallel)**
   (d) **Determine the best model fit, for example with the k-fold cross validation technique.**
   (e) **Choose the best model fit and save the correction factor parameters.**

Bibliography

[1] OpenFOAM, 2020. URL `https://openfoam.org/`.

[2] J. F. Ainslie. Calculating the flowfield in the wake of wind turbines. *Journal of Wind Engineering*, 27 (1-3):213–224, 1988.

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

[4] S. J. Andersen, B. Witha, S-P. Breton, J. N. Sørensen, R. F. Mikkelsen, and S. Ivanell. Quantifying variability of Large Eddy Simulations of very large wind farms. *Journal of Physics: Conference Series*, 625(1): 0–12, 2015. doi: 10.1088/1742-6596/625/1/012027.

[5] J. Annoni, P. M. O. Gebraad, A. K. Scholbrock, P. A. Fleming, and J-W. van Wingerden. Analysis of axial-induction-based wind plant control using an engineering and a high-order wind plant model. *Wind Energy*, (19):1135–1150, 2016. doi: 10.1002/we1891.

[6] J. Annoni, P. A. Fleming, A. K. Scholbrock, J. Roadman, S. Dana, C. Adcock, F. Porté-Agel, S. Raach, F. Haizmann, and D. Schlipf. Analysis of control-oriented wake modeling tools using lidar field results. *Wind Energy Science*, 3(2):819–831, 2018. doi: 10.5194/wes-3-819-2018.

[7] R. J. Barthelmie, K. S. Hansen, S. T. Frandsen, O. Rathmann, J. G. Schepers, W. Schlez, J. Phillips, K. Rados, A. Zervos, E. S. Politis, and P. K. Chaviaropoulos. Modelling and measuring flow and wind turbine wakes in large wind farms offshore. *Wind Energy*, 12(5):431–444, 2009. doi: 10.1002/we.348.

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

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

[10] J. Berg, J. Mann, M. Kelly, and M. Nielsen. *Micro meteorology for Wind Energy*. 2017.

[11] C. Bergmeir, R. J. Hyndman, and B. Koo. A note on the validity of cross-validation for evaluating autoregressive time series prediction. *Computational Statistics and Data Analysis*, 120:70–83, 2018. doi: 10.1016/j.csda.2017.11.003.

[12] J. Bleeg, M. Purcell, R. Ruisi, and E. Traiger. Wind farm blockage and the consequences of neglecting its impact on energy production. *Energies*, 11(6), 2018. doi: 10.3390/en11061609.

[13] S. Boersma, B. M. Doekemeijer, P. M. O. Gebraad, P. A. Fleming, J. Annoni, A. K. Scholbrock, J. A. Frederik, and J-W. van Wingerden. A tutorial on control-oriented modeling and control of wind farms. *Proceedings of the American Control Conference*, pages 13–120, 2017. doi: 10.23919/ACC.2017.7962923.

[14] S. Boersma, B. M. Doekemeijer, M. Vali, J. Meyers, and J-W. van Wingerden. A control-oriented dynamic wind farm model: WFSim. *Wind Energy Science*, 3(1):75–95, 2018. doi: 10.5194/wes-3-75-2018.

[15] K. Boorsma. Power and loads for wind turbines in yawed conditions. (November), 2012.

[16] J. Bossuyt, M. F. Howland, C. Meneveau, and J. Meyers. Measurement of unsteady loading and power output variability in a micro wind farm model in a wind tunnel. *Experiments in Fluids*, 58(1):1–17, 2017. doi: 10.1007/s00348-016-2278-6.

[17] M. M. Breuniq, H. P. Kriegel, R. T. Ng, and J. Sander. LOF: Identifying density-based local outliers. *SIGMOD Record (ACM Special Interest Group on Management of Data)*, 29(2):93–104, 2000. doi: 10.1145/335191.335388.

[18] F. Campagnolo, V. Petrović, C. L. Bottasso, and A. Croce. Wind tunnel testing of wake control strategies. *Proceedings of the American Control Conference*, 2016-July:513–518, 2016. doi: 10.1109/ACC.2016.7524965.

[19] D. Cavar, P.-E. Réthoré, A. Bechmann, N. N. Sørensen, B. Martinez, F. Zahle, J. Berg, and M. C. Kelly. Comparison of OpenFOAM and EllipSys3D for neutral atmospheric flow over complex terrain. *Wind Energy Science*, 1(1):55–70, 2016. doi: 10.5194/wes-1-55-2016.

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

[21] A. Clifton and R. Wagner. Accounting for the effect of turbulence on wind turbine power curves. *Journal of Physics: Conference Series*, 524(1), 2014. doi: 10.1088/1742-6596/524/1/012109.

[22] A. Crespo and J. Hernández. Turbulence characteristics in wind-turbine wakes. *Journal of Wind Engineering and Industrial Aerodynamics*, 61(1):71–85, 1996. doi: 10.1016/0167-6105(95)00033-X.

[23] J-A. Dahlberg. Assessment of the Lillgrund Windfarm. (September):28, 2009.

[24] J-A. Dahlberg and B. Montgomerie. Research program of the utgrunden demonstration offshore wind farm, final report part 2, wake effects and other loads. *Swedish Defense Research Agency*, 2005.

[25] R. Damiani, S. Dana, J. Annoni, P. A. Fleming, J. Roadman, J. van Dam, and K. Dykes. Assessment of Wind Turbine Component Loads Under Yaw-Offset Conditions. *Wind Energy Science Discussions*, 3: 173–189, 2018. doi: 10.5194/wes-2017-38.

[26] B. M. Doekemeijer and R. Storm. FLORISSE_M, 2020. URL `https://github.com/TUDelft-DataDrivenControl/FLORISSE_M`.

[27] B. M. Doekemeijer, J-W. van Wingerden, and P. A. Fleming. A tutorial on the synthesis and validation of a closed-loop wind farm controller using a steady-state surrogate model. *Proceedings of the American Control Conference*, pages 2825–2836, 2019.

[28] B. M. Doekemeijer, S. Kern, S. Maturu, S. Kanev, B. Salbert, J. Schreiber, F. Campagnolo, C. L. Bottasso, S. Schuler, F. Wilts, T. Neumann, G. Potenza, F. Calabretta, and F. Fioretti. Field experiment for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy. *Wind Energy Science Discussions*, (June), 2020. doi: 10.5194/wes-2020-80.

[29] B. M. Doekemeijer, D. van der Hoek, and J-W. van Wingerden. Closed-loop model-based wind farm control using FLORIS under time-varying inflow conditions. *Renewable Energy*, 156:719–730, 2020. doi: 10.1016/j.renene.2020.04.007.

[30] N. R. Draper and H. Smith. *Applied Regression Analysis.* John Wiley & Sons, Inc, third edition, 1998.

[31] B. Efron. Bootstrap Methods : Another Look at the Jackknife. *The Annals of Statistics*, 7(1):1–26, 1979.

[32] B. Efron and R. J. Tibshirani. *An introduction to the Bootstrap.* Chapman and Hall, New York, 1993.

[33] P. A. Fleming, P. M. O. Gebraad, J-W. van Wingerden, S. Lee, M. J. Churchfield, A. K. Scholbrock, J. Michalakes, K. E. Johnson, and P. J. Moriarty. The SOWFA super-controller: A high-fidelity tool for evaluating wind plant control approaches. *European Wind Energy Conference and Exhibition, EWEC 2013*, 3:1561–1570, 2013.

[34] P. A. Fleming, P. M. O. Gebraad, S. Lee, J-W. van Wingerden, K. E. Johnson, M. J. Churchfield, J. Michalakes, P. Spalart, and P. J. Moriarty. Simulation comparison of wake mitigation control strategies for a two-turbine case. *Wind Energy*, 18:2135–2143, 2015. doi: 10.1002/we.1810.

[35] P. A. Fleming, A. Ning, P. M. O. Gebraad, and K. Dykes. Wind plant system engineering through optimization of layout and yaw control. *Wind Energy*, 19(March 2015):329–344, 2 2016. doi: 10.1002/we.

[36] P. A. Fleming, J. Annoni, J. J. Shah, L. Wang, S. Ananthan, Z. Zhang, K. Hutchings, P. Wang, W. Chen, and L. Chen. Field test of wake steering at an offshore wind farm. *Wind Energy Science*, 2(1):229–239, 2017. doi: 10.5194/wes-2-229-2017.

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

[38] P. A. Fleming, J. Annoni, L. A. Martínez-Tossas, S. Raach, K. Gruchalla, A. K. Scholbrock, M. J. Churchfield, and J. Roadman. Investigation into the shape of a wake of a yawed full-scale turbine. *Journal of Physics: Conference Series*, 1037(3), 2018. doi: 10.1088/1742-6596/1037/3/032010.

[39] A. R. M. Forsting and N. Troldborg. The effect of blockage on power production for laterally aligned wind turbines. *Journal of Physics: Conference Series*, 625(1), 2015. doi: 10.1088/1742-6596/625/1/012029.

[40] S. T. Frandsen, R. J. Barthelmie, S. Pryor, O. Rathmann, S. Larsen, J. Højstrup, and M. Thøgersen. Analytical modelling of wind speed deficit in large offshore wind farms. *Wind Energy*, 9(1-2):39–53, 2006. doi: 10.1002/we.189.

[41] J. A. Frederik, R. Weber, S. Cacciola, F. Campagnolo, A. Croce, C. L. Bottasso, and J-W. van Wingerden. Periodic dynamic induction control of wind farms: proving the potential in simulations and wind tunnel experiments. *Wind Energy Science Discussions*, (August), 2019. doi: 10.5194/wes-2019-50.

[42] J. A. Frederik, B. M. Doekemeijer, S. P. Mulders, and J-W. van Wingerden. The helix approach: Using dynamic individual pitch control to enhance wake mixing in wind farms. *Wind Energy*, (March):1–13, 2020. doi: 10.1002/we.2513.

[43] M. Gaumond, P-E. Réthoré, A. Bechmann, S. Ott, G. C. Larsen, Al. Pena Diaz, and K. S. Hansen. Benchmarking of Wind Turbine Wake Models in Large Offshore Windfarms. *Proceedings of Torque 2012, The science of making torque from wind*, 2012.

[44] M. Gaumond, P-E. Réthoré, S. Ott, A. Peña, A. Bechmann, and K. S. Hansen. Evaluation of the wind direction uncertainty and its impact on wake modeling at the Horns Rev offshore wind farm. *Wind Energy*, 17(8):1169–1178, 2014. doi: 10.1002/we.1625.

[45] P. M. O. Gebraad. *Data-Driven Wind Plant Control.* PhD thesis, Delft University of Technology, 2014.

[46] P. M. O. Gebraad and J-W. van Wingerden. A control-oriented dynamic model for wakes in wind plants. *Journal of Physics: Conference Series*, 524(1), 2014. doi: 10.1088/1742-6596/524/1/012186.

[47] P. M. O. Gebraad, F. W. Teeuwisse, J-W. van Wingerden, P. A. Fleming, S. D. Ruben, J. R. Marden, and L. Y. Pao. A data-driven model for wind plant power optimization by yaw control. *Proceedings of the American Control Conference*, pages 3128–3134, 2014. doi: 10.1109/ACC.2014.6859118.

[48] P. M. O. Gebraad, M. J. Churchfield, and P. A. Fleming. Incorporating Atmospheric Stability Effects into the FLORIS Engineering Model of Wakes in Wind Farms. *Journal of Physics: Conference Series*, 753(5), 2016. doi: 10.1088/1742-6596/753/5/052004.

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

[50] P. M. O. Gebraad, J. J. Thomas, A. Ning, P. A. Fleming, and K. Dykes. Maximization of the annual energy production of wind power plants by optimization of layout and yaw-based wake control. *Wind Energy*, 20(May 2016):97–107, 2016. doi: 10.1002/we.1993.

[51] T. Göçmen and G. Giebel. Estimation of turbulence intensity using rotor effective wind speed in Lillgrund and Horns Rev-I offshore wind farms. *Renewable Energy*, 99:524–532, 2016. doi: 10.1016/j.renene.2016.07.038.

[52] T. Göçmen, G. Giebel, N. K. Poulsen, and M. Mirzaei. Wind speed estimation and parametrization of wake models for downregulated offshore wind farms within the scope of PossPOW project. *Journal of Physics: Conference Series*, 524(1), 2014. doi: 10.1088/1742-6596/524/1/012156.

[53] T. Göçmen, G. Giebel, N. K. Poulsen, and P. E. Sørensen. Possible power of down-regulated offshore wind power plants: The PossPOW algorithm. *Wind Energy*, 22(2):205–218, 2019. doi: 10.1002/we.2279.

[54] J. P. Goit and J. Meyers. Optimal control of energy extraction in wind-farm boundary layers. *Journal of Fluid Mechanics*, 768:5–50, 2015. doi: 10.1017/jfm.2015.70.

[55] J. S. González, M. B. Payán, and J. M. Riquelme-Santos. Optimization of wind farm turbine layout including decision making under risk. *IEEE Systems Journal*, 6(1):94–102, 2012. doi: 10.1109/JSYST.2011.2163007.

[56] J. D. Grunnet, M. Soltani, T. Knudsen, M. Kragelund, and T. Bak. Aeolus toolbox for dynamics wind farm model, simulation and control. *European Wind Energy Conference and Exhibition 2010, EWEC 2010*, 2010.

[57] K. S. Hansen. Guideline to wind farm wake analysis. *Risø National Laboratory, Technical University of Denmark*, (Project UpWind, Appendix 7):1–20, 2008.

[58] K. S. Hansen, R. J. Barthelmie, L. E. Jensen, and A. Sommer. The impact of turbulence intensity and atmospheric stability on power deficits due to wind turbine wakes at Horns Rev wind farm. *Wind Energy*, 15:183–196, 2012. doi: 10.1002/we.512.

[59] S. Heier. *Grid integration of wind energy conversion systems.* Wiley, Chichester, UK, 1998. ISBN 047197143x.

[60] M. F. Howland, S. K. Lele, and J. O. Dabiri. Wind farm power optimization through wake steering. *Proceedings of the National Academy of Sciences of the United States of America*, 116(29):14495–14500, 2019. doi: 10.1073/pnas.1903680116.

[61] P. Hulsman, S. J. Andersen, and T. Göçmen. Optimizing Wind Farm Control through Wake Steering using Surrogate Models based on High Fidelity Simulations. *Wind Energy Science Discussions*, (August), 2019. doi: 10.5194/wes-2019-46.

[62] IRENA. Renewable Energy Statistics 2019, 2019. URL `https://irena.org/Statistics/View-Data-by-Topic/Capacity-and-Generation/Statistics-Time-Series`.

[63] N. O. Jensen. A note on wind generator interaction. *Risø-M-2411 Risø National Laboratory Roskilde*, pages 1–16, 1983. doi: Riso-M-2411.

[64] J. Jeppsson, P. E. Larsen, and A. Larsson. Technical Description Lillgrund Wind Power Plant. *Vattenfall Vindkraft AB*, (September):79, 2008.

[65] Á. Jiménez, A. Crespo, and E. Migoya. Application of a LES technique to characterize the wake deflection of a wind turbine in yaw. *Wind Energy*, 13(December 2009):559–572, 2010. doi: 10.1002/we.380.

[66] K. E. Johnson and N. Thomas. Wind farm control: addressing the aerodynamic interaction among wind turbines. *Proceedings of the American Control Conference*, pages 2104–2109, 2009. doi: 10.1109/ACC.2009.5160152.

[67] J. Jonkman and B. Jonkman. NWTC Information Portal (FAST V8), 2018. URL `https://nwtc.nrel.gov/FAST8`.

[68] I. Katić, J. Højstrup, and N. O. Jensen. A Simple Model for Cluster Efficiency. *Proceedings of the European Wind Energy Association Conference and Exhibition*, 1:407–410, 1987.

[69] R-E. Keck, D. Veldkamp, J. J. Wedel-Heinen, and J. Forsberg. *A consistent turbulence formulation for the dynamic wake meandering model in the atmospheric boundary layer.* PhD thesis, DTU, 2013.

[70] R-E. Keck, M. de Maré, M. J. Churchfield, S. Lee, G. C. Larsen, and H. Aa. Madsen. On atmospheric stability in the dynamic wake meandering model. *Wind Energy*, (17):1689–1710, 2014. doi: 10.1002/we1662.

[71] A. C. Kheirabadi and R. Nagamune. A quantitative review of wind farm control with the objective of wind farm power maximization. *Journal of Wind Engineering and Industrial Aerodynamics*, 192(June): 45–73, 2019. doi: 10.1016/j.jweia.2019.06.015.

[72] J. King, P. A. Fleming, R. N. King, L. A. Martínez-Tossas, C. Bay, R. Mudafort, and E. Simley. Controls-Oriented Model for Secondary Effects of Wake Steering. *Wind Energy Science Discussions*, (February): 1–22, 2020. doi: 10.5194/wes-2020-3.

[73] D. Kraft. A Software Package for Sequential Quadratic Programming. Technical Report 28, DLR German Aerospace Center - Institute for Flight Mechanics, Koln, Germany, 1988.

[74] S. D. Kwon. Uncertainty analysis of wind energy potential assessment. *Applied Energy*, 87(3):856–865, 2010. doi: 10.1016/j.apenergy.2009.08.038.

[75] M. A. Lackner, A. L. Rogers, and J. F. Manwell. Uncertainty Analysis in Wind Resource Assessment and Wind Energy Production Estimation Matthew. *American Institute of Aeronautics and Astronautics*, 2014. doi: 10.23919/ecc.2007.7068771.

[76] G. C. Larsen. A Simple Wake Calculation Procedure. *Risø-M*, No. 2760, 1988.

[77] G. C. Larsen, H. Aa. Madsen, F. Bingöl, J. Mann, S. Ott, J. N. Sørensen, V. Okulov, N. Troldborg, N. M. Nielsen, K. Thomsen, T. J. Larsen, and R. F. Mikkelsen. *Dynamic wake meandering modeling*. Roskilde, Denmark: Risø National Laboratory, 2007. ISBN 978-87-550-3602-4.

[78] G. C. Larsen, T. J. Larsen, and A. Chougule. Medium fidelity modelling of loads in wind farms under non-neutral ABL stability conditions - A full-scale validation study. *Journal of Physics: Conference Series*, 854(1), 2017. doi: 10.1088/1742-6596/854/1/012026.

[79] G. C. Larsen, S. Ott, T. J. Larsen, K. S. Hansen, and A. Chougule. Improved modelling of fatigue loads in wind farms under non-neutral ABL stability conditions. *Journal of Physics: Conference Series*, 1037(7), 2018. doi: 10.1088/1742-6596/1037/7/072013.

[80] T. J. Larsen, G. C. Larsen, M. M. Pedersen, K. Enevoldsen, and H. Aa. Madsen. Validation of the Dynamic Wake Meander model with focus on tower loads. *Journal of Physics: Conference Series*, 854(1), 2017. doi: 10.1088/1742-6596/854/1/012027.

[81] A. Lavely, G. Vijayakumar, M. Kinzel, J. Brasseur, and E. Paterson. Space-Time Loadings on Wind Turbine Blades Driven by Atmospheric Boundary Layer Turbulence. *AIAA*, (January), 2011. doi: 10.2514/6.2011-635.

[82] J. C. Y. Lee, P. Stuart, A. Clifton, M. J. Fields, J. Perr-Sauer, L. Williams, L. Cameron, T. Geer, and P. Housley. The Power Curve Working Group's Assessment of Wind Turbine Power Performance Prediction Methods. *Wind Energy Science Discussions*, pages 199–223, 2019. doi: 10.5194/wes-5-199-2020.

[83] S. Lee, M. J. Churchfield, P. J. Moriarty, J. Jonkman, and J. Michalakes. Atmospheric and wake turbulence impacts on wind turbine fatigue loadings. *50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, (December 2011), 2012. doi: 10.2514/6.2012-540.

[84] J. Liew, A. M. Urbán, and S. J. Andersen. Analytical model for the power-yaw sensitivity of wind turbines operating in full wake. *Wind Energy Science Discussions*, (October), 2019. doi: 10.5194/wes-2019-65.

[85] P. B. S. Lissaman. Energy Effectiveness of Arbitrary Arrays of Wind Turbines. *Journal of Energy*, 3(6): 323–328, 1979.

[86] H. Aa. Madsen, N. N. Sørensen, and S. Schreck. Yaw aerodynamics analyzed with three codes in comparison with experiment. *ASME 2003 Wind Energy Symposium, WIND2003*, (January), 2003. doi: 10.2514/6.2003-519.

[87] H. Aa. Madsen, G. C. Larsen, T. J. Larsen, N. Troldborg, and R. F. Mikkelsen. Calibration and validation of the dynamic wake meandering model for implementation in an aeroelastic code. *Journal of Solar Energy Engineering, Transactions of the ASME*, 132(4):1–14, 2010. doi: 10.1115/1.4002555.

[88] B. Maronga, M. Gryschka, R. Heinze, F. Hoffmann, F. Kanani-Sühring, M. Keck, K. Ketelsen, M. O. Letzel, M. Sühring, and S. Raasch. The Parallelized Large-Eddy Simulation Model (PALM) version 4.0 for atmospheric and oceanic flows: Model formulation, recent developments, and future perspectives. *Geoscientific Model Development*, 8(8):2515–2551, 2015. doi: 10.5194/gmd-8-2515-2015.

[89] L. A. Martínez-Tossas, M. J. Churchfield, and S. Leonardi. Large eddy simulations of the flow past wind turbines: actuator line and disk modeling. *Wind Energy*, 18:1047–1060, 2015. doi: 10.1002/we.1747.

[90] L. A. Martínez-Tossas, J. Annoni, P. A. Fleming, and M. J. Churchfield. The aerodynamics of the curled wake: a simplified model in view of flow control. *Wind Energy Science*, 4(1):127–138, 2019. doi: 10.5194/wes-4-127-2019.

[91] D. Medici. *Experimental studies of wind turbine wakes: power optimisation and meandering*. PhD thesis, Royal Institute of Technology, 2005.

[92] A. S. Monin and A. M. Obukhov. Basic laws of turbulent mixing in the surface layer of the atmosphere. *Tr. Akad. Nauk SSSR Geophiz. Inst.*, 24(151):163–187, 1954.

[93] W. Munters and J. Meyers. An optimal control framework for dynamic induction control of wind farms and their interaction with the atmospheric boundary layer. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2091), 2017. doi: 10.1098/rsta.2016.0100.

[94] W. Munters and J. Meyers. 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, 2018. doi: 10.5194/wes-3-409-2018.

[95] J. P. Murcia Leon. *Uncertainty quantification in wind farm flow models*. PhD thesis, DTU Wind Energy, 2017.

[96] A. Niayifar and F. Porté-Agel. Analytical modeling of wind farms: A new approach for power prediction. *Energies*, 9(9):1–13, 2016. doi: 10.3390/en9090741.

[97] J. Nossent, P. Elsen, and W. Bauwens. Sobol' sensitivity analysis of a complex environmental model. *Environmental Modelling and Software*, 26(12):1515–1525, 2011. doi: 10.1016/j.envsoft.2011.08.010.

[98] NREL. FLORIS. Version 2.0.1, 2020. URL `https://github.com/NREL/floris`.

[99] NREL. SOWFA, 2020. URL `https://github.com/NREL/SOWFA`.

[100] N. G. Nygaard. Systematic quantification of wake model uncertainty. *EWEA Offshore*, 2015.

[101] A. M. Obukhov. Turbulence in an atmosphere with a non-uniform temperature. *Boundary-Layer Meteorology*, 2(1):7–29, 1971. doi: 10.1007/BF00718085.

[102] A. Patton, D. N. Politis, and H. White. Correction to automatic block-length selection for the dependent bootstrap by D. Politis and H. White. *Econometric Reviews*, 28(4):372–375, 2009. doi: 10.1080/07474930802459016.

[103] D. N. Politis and J. P. Romano. A circular block-resampling procedure for stationary data. *Technical Report No. 370*, 1991.

[104] D. N. Politis and H. White. Automatic Block-Length Selection for the Dependent Bootstrap. *Econometric Reviews*, 23(1):53–70, 2004. doi: 10.1081/ETC-120028836.

[105] J. Quick, J. Annoni, R. N. King, K. Dykes, P. A. Fleming, and A. Ning. Optimization under Uncertainty for Wake Steering Strategies. *Journal of Physics: Conference Series*, 854(1), 2017. doi: 10.1088/1742-6596/854/1/012036.

[106] J. Quick, J. King, R. N. King, P. E. Hamlington, and K. Dykes. Wake steering optimization under uncertainty. *Wind Energy Science Discussions*, (October), 2019. doi: 10.5194/wes-2019-72.

[107] I. Reinwardt, N. Gerke, P. Dalhoff, D. Steudel, and W. Moser. Validation of wind turbine wake models with focus on the dynamic wake meandering model. *Journal of Physics: Conference Series*, 1037(7), 2018. doi: 10.1088/1742-6596/1037/7/072028.

[108] S. F. Rodrigues, R. Teixeira Pinto, M. Soleimanzadeh, P. A. N. Bosman, and P. Bauer. Wake losses optimization of offshore wind farms with moveable floating wind turbines. *Energy Conversion and Management*, 89:933–941, 2015. doi: 10.1016/j.enconman.2014.11.005.

[109] A. Rott, B. M. Doekemeijer, J. K. Seifert, J-W. van Wingerden, and M. Kühn. Robust active wake control in consideration of wind direction variability and uncertainty. *Wind Energy Science*, 3(2):869–882, 11 2018. doi: 10.5194/wes-3-869-2018.

[110] A Saltelli and T. Homma. Importance measures in global sensitivity analysis of model output. *Reliab. Eng. Sys. Safety*, 52:1–17, 1996.

[111] B. Sanderse. Aerodynamics of wind turbine wakes - Literature Review. *Energy research Centre of the Netherlands*, 2009.

[112] J. Sanz-Rodrigo, P. J. Moriarty, E. Cantero, R. Chávez-Arroyo, M. J. Churchfield, and K. S. Hansen. Guideline for qualification of SCADA data for wake efficiency analysis. In *WAKEBENCH Best Practice Guidelines for Wind Farm Flow Models. 1st Edition*, number April. IEA Task 31 Report to the IEA-Wind Executive Committee, 2015.

[113] A. Sathe, J. Mann, T. Barlas, W. A. A. M. Bierbooms, and G. J. W. van Bussel. Influence of atmospheric stability on wind turbine loads. *Wind Energy*, (16):1013–1032, 2012. doi: 10.1002/we1528.

[114] J. G. Schepers. EU project in German Dutch wind tunnel, Technical Report ECN-RX-01-006. *Energy Research Center of the Netherlands, ECN*, 2001.

[115] J. G. Schepers, T. S. Obdam, and J. Prospathopoulos. Analysis of wake measurements from the ECN Wind Turbine Test Site Wieringermeer, EWTW. *Wind Energy*, (15):575–591, 2012. doi: 10.1002/we.488.

[116] E. Simley, P. A. Fleming, and J. King. Design and Analysis of a Wake Steering Controller with Wind Direction Variability. *Wind Energy Science Discussions*, (July), 2019. doi: 10.5194/wes-2019-35.

[117] I. M. Sobol. On sensitivity estimation for nonlinear mathematical models. *Matematicheskoe Modelirovanie*, 8(1):112–118, 1990. doi: 10.18287/0134-2452-2015-39-4-459-461.

[118] N. N. Sørensen. *General purpose flow solver applied to flow over hills*. PhD thesis, Risø National Laboratory, 1995.

[119] V. Spudic. Hierarchical wind farm control for power/load optimization. *Journal of Physics: Conference Series*, 2010.

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

[121] R. Storn and K. Price. Differential Evolution - A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *Journal of Global Optimization*, 11:341–359, 1997.

[122] R. B. Stull. *An Introduction to Boundary Layer Meteorology*, volume 13. Kluwer Academic Publishers Group, 1988.

[123] H. J. Sutherland. On the fatigue analysis of wind turbines. Technical report, Albuquerque, 1999.

[124] Uknown. Lillgrund wind farm. URL `https://www.poyry.com/sites/default/files/media/related_material/lillgrund_a4_eng_new.pdf`.

[125] M. T. van Beek. Gitlab Wind Energy - MaartenBeek 2020, 2020. URL `https://gitlab.windenergy.dtu.dk/MasterProjects/maartenbeek2020`.

[126] M. T. van Dijk, J-W. van Wingerden, T. Ashuri, and Y. Li. Wind farm multi-objective wake redirection for optimizing power production and loads. *Energy*, 121:561–569, 2017. doi: 10.1016/j.energy.2017.01.051.

[127] J. W. Wagenaar, L. A. H. Machielse, and J. G. Schepers. Controlling wind in ECN's scaled wind farm. *European Wind Energy Conference and Exhibition 2012, EWEC 2012*, 1(April):161–168, 2012.

[128] J. C. Wyngaard. The atmospheric boundary layer. In *Turbulence in Atmosphere*, pages 193–214. Cambridge University Press, 2010. doi: 10.1007/3-211-38078-7{_}4.

[129] D. S. Zalkind and L. Y. Pao. The fatigue loading effects of yaw control for wind plants. *Proceedings of the American Control Conference*, 2016-July:537–542, 2016. doi: 10.1109/ACC.2016.7524969.

[130] X. Y. Zhang, M. N. Trame, L. J. Lesko, and S. Schmidt. Sobol sensitivity analysis: A tool to guide the development and evaluation of systems pharmacology models. *CPT: Pharmacometrics and Systems Pharmacology*, 4(2):69–79, 2015. doi: 10.1002/psp4.6.