

Minute-scale forecasting of wind power using long-range lidar data

A thesis accepted by the Faculty of Aerospace Engineering and Geodesy of the
University of Stuttgart in partial fulfillment of the requirements for the degree of
Doctor of Engineering Sciences (Dr.-Ing.)

by

Ines Würth

born in Stuttgart, Germany

Main referee: Prof. Dr. Po Wen Cheng
Co-referee: Prof. Dr. Charlotte Bay Hasager

Date of defense: 06.12.2021

Institute of Aircraft Design
University of Stuttgart
2022

Contents

Abbreviations	xiii
List of Symbols	xv
Abstract	xvii
Kurzfassung	xix
1 Introduction	1
1.1 Motivation	1
1.2 Research areas in this thesis	2
1.3 Research objectives, methodology and organization	4
2 Background	7
2.1 Intra-hourly variability of wind power generation	7
2.2 Application areas for minute-scale forecasting	9
2.3 Lidar in wind energy applications	11
2.4 How is wind power forecast today?	12
2.4.1 A note on terminology	12
2.4.2 Numerical weather prediction	13
2.4.3 Statistical time series models	15
2.4.4 Where is the gap that needs to be closed?	15
2.5 Summary	17
3 The forecasting chain	19
4 First step in the forecasting chain: from radial velocity to wind field	23
4.1 Measurement setup at the onshore and offshore site	24
4.1.1 Definition of scientific objectives	24
4.1.2 Site selection	25
4.1.3 Site characterization	26
4.1.4 Experiment layout design	29
4.1.5 Infrastructure planning	31
4.1.6 Deployment and calibration procedures	32
4.1.7 Scanning modes design	32

4.1.8	Execution and data collection	33
4.1.9	Decommissioning and post-calibration procedures	35
4.1.10	Data availability	35
4.2	Getting a useful wind speed out of a lidar	37
4.2.1	Data filtering	37
4.2.2	Wind field reconstruction	43
4.3	Lessons learned	57
5	Along the forecasting chain: wind field evolution	59
5.1	Using Taylor's hypothesis in wind energy applications	60
5.2	Using Taylor's hypothesis for minute-scale forecasting	61
5.3	Lessons learned	63
6	Reaching the end of the chain: minute-scale forecasts of wind speed and power	65
6.1	Methodology of probabilistic minute-scale forecasts	66
6.1.1	Converting wind speed to power	69
6.1.2	Choosing a forecast horizon	70
6.1.3	Quantifying the uncertainty with probability density function and cumulative distribution function	70
6.1.4	Evaluating probabilistic forecasts	75
6.1.5	Comparing to benchmark forecasting method persistence	79
6.2	Results of minute-scale forecasts from the onshore campaign	82
6.2.1	Availability of forecasts	83
6.2.2	Forecasts in the time domain	85
6.2.3	Calibration	85
6.2.4	Sharpness and skill	90
6.3	Results of minute-scale forecasts from the offshore campaign	93
6.3.1	Availability of forecasts	93
6.3.2	Forecasts in the time domain	95
6.3.3	Calibration	97
6.3.4	Sharpness and skill	100
6.4	Impact of wind ramps on the forecast accuracy	113
6.4.1	Detecting wind ramps	114
6.4.2	Assessing the ramp impact on the forecast accuracy	118
6.4.3	Case study of a failed ramp forecast in the onshore campaign	118
6.5	Lessons learned	121
7	Challenges for the implementation of lidar-based minute-scale forecasting	127
7.1	How far do we see?	127
7.1.1	Correlating the measurement range and the forecasting horizon	128
7.1.2	Analysing the measurement range	128
7.2	Overcoming barriers to adoption with the help of the community	140
7.2.1	Availability of measurements	141
7.2.2	Reliability and pricing	141
7.2.3	Need for standards and common tools	142

7.3	Lessons learned	143
8	Summary and conclusions	145
8.1	Summary	145
8.1.1	How should the lidar data information be processed to gain a power forecast for a wind turbine?	145
8.1.2	How does the measurement setup and measurement site influence the forecast?	146
8.1.3	What is the forecast horizon of lidar-based forecasts and what influences the forecast horizon? How does the lidar-based forecast perform in comparison to state-of-the-art statistical methods and what are its benefits? .	147
8.2	Discussion and future work	148
8.3	Conclusions	150
A	Appendix	151
A.1	Lidar Data Sheet	151
A.2	Met mast sensor equipment	152
	Bibliography	155

List of Figures

2.1	Example of wind ramps in wind speed and generated power.	9
2.2	Forecasting horizons of diferent wind energy applications.	10
2.3	State-of-the-art forecasting approaches.	13
2.4	Different classes of weather models.	14
2.5	Forecast error for different forecasting techniques.	16
3.1	The lidar forecasting chain.	21
4.1	Measurement sites marked on map of Germany.	25
4.2	Wind characteristic in Stötten measured at the met mast.	26
4.3	Map of the measurement site in Stötten.	27
4.4	Wind characteristic in alpha ventus.	28
4.5	StreamLine XR lidar on the top level platform of the radio tower.	29
4.6	StreamLine XR lidar on the nacelle of the AV04 turbine in alpha ventus.	30
4.7	Schematic drawing of the lidar measurement setup.	33
4.8	Typical time series of LOS lidar data unfiltered and filtered.	39
4.9	Radial wind speed over CNR sorted by the measurement range.	40
4.10	Image of the Stuttgart TV tower with edge detection filter.	41
4.11	Demonstration of edge filter for a generic wind speed time series.	42
4.12	Percentage of good points after filtering with edge filter.	42
4.13	Radial velocity over CNR filtered with the edge filter.	43
4.14	Overview of the coordinate systems used for the wind field reconstruction.	46
4.15	Wind direction and wind speed using global-local reconstruction method.	47
4.16	Regression of wind data using the global-local reconstruction method.	48
4.17	Sketch of the local moving window reconstruction method.	49
4.18	Regression using the moving window reconstruction method.	50
4.19	Regression parameters using the moving window reconstruction method.	51
4.20	Scan plots of wind direction and wind speed from different reconstruction methods.	53
4.21	Overview of the coordinate systems used in alpha ventus.	54
4.22	Reconstructed horizontal wind speed in alpha ventus.	56
4.23	Correlation of the yaw angle and the wind direction.	56
5.1	Sketch of a turbulent eddy passing a wind turbine.	60
5.2	Scan plots of wind speed propagation.	64
6.1	Timeline of measured and forecasted wind speed.	66

6.2	Probabilistic forecast process.	68
6.3	Power curves of the onshore and offshore site.	71
6.4	Example of uncertainty information for forecasted power data.	73
6.5	Examples of PIT histograms of different forecasts.	76
6.6	Predictive PDF and CDF with low sharpness and high sharpness.	78
6.7	Comparison of the wind speed measurements at the met mast and turbine.	79
6.8	Correlation of the wind speed measured at turbine and met mast.	80
6.9	Example of a probabilistic persistence power forecast.	82
6.10	Availability of power data and forecasts onshore.	84
6.11	Example of the probabilistic wind speed and power forecast.	86
6.12	PIT histograms for onshore wind speed forecasts.	87
6.13	PIT histograms for onshore power forecasts.	88
6.14	CRPS of onshore forecast.	91
6.15	Histograms of wind speed CRPS onshore.	92
6.16	Histograms of power CRPS onshore.	92
6.17	Availability of power data and forecasts offshore.	94
6.18	Example of the probabilistic forecast offshore.	96
6.19	PIT diagrams wind speed forecasts offshore.	98
6.20	PIT diagrams power forecasts offshore.	99
6.21	CRPS of offshore forecasts.	101
6.22	Histograms of CRPS of wind speed forecast offshore.	102
6.23	Histograms of CRPS of power forecast offshore.	102
6.24	Example CDFs of wind speed and power forecast.	103
6.25	Histograms of error of deterministic forecast.	104
6.26	Polar plots of CRPS depending on wind direction.	105
6.27	CRPS depending on the turbine wind speed.	106
6.28	Increase of measurement height against the measurement range.	107
6.29	Deviation of wind speed between lidar and hub height.	108
6.30	CRPS for different measurement ranges.	110
6.31	CRPS depending on the number of measured wind speeds vectors.	112
6.32	CRPS depending on stability classes.	113
6.33	Schematic diagram of a ramp matrix.	115
6.34	Number of upward and downward ramps.	116
6.35	Timeline of power generation with ramps marked.	117
6.36	CRPS for periods with and without ramp event.	119
6.37	Timeline of a ramp event at the onshore site.	120
6.38	Wind speed propagation during a ramp event.	122
7.1	Forecast horizon calculated based on Taylor.	129
7.2	Schematic visualization of two methods to determine the maximum range.	130
7.3	Comparison of the Sum Range and Weighted Range method.	131
7.4	Maximum measurement range over time.	131
7.5	Comparison of of the valid measurement points over the measurement range.	133
7.6	Boxplot of maximum measurement range for different number of pulses.	135
7.7	Histograms of maximum range for different number of pulses.	135
7.8	Maximum measurement range versus different environmental conditions.	137
7.9	Correlation of humidity and rain sensor data.	137

7.10	MIC between maximum range and environmental variables.	138
7.11	Maximum measurement range on June, 8th 2016.	138
7.12	Webcam pictures of the measurement site on June 8th, 2016.	139
A.1	Sensor equipment of the met mast in Stötten, effective 13.10.2015	152

List of Tables

4.1	Overview of lidar scans at the onshore campaign in Stötten.	34
4.2	Overview of lidar scans at the offshore campaign in alpha ventus.	34
4.3	Overview of available data for the onshore campaign.	36
4.4	Overview of available data for offshore measurement campaign.	36
6.1	Filter criteria for onshore power curve.	69
6.2	Example timeline of turbine power data.	81
6.3	Example of probabilistic persistence power data sets.	81
6.4	Total number of available forecasts onshore.	83
6.5	Number of available forecasts offshore.	95
6.6	Median of wind speed CRPS offshore.	100
6.7	Stability classes using Ri_S	111
7.1	Overview of measurement periods with different number of pulses.	134
A.1	Data sheet for Stream Line XR lidar.	151
A.2	Sensors on the met mast of the offshore research platform FINO1.	153

Abbreviations

BSH	Bundesamt für Seeschifffahrt und Hydrographie (Federal Maritime and Hydrographic Agency of Germany)
CDF	Cumulative Distribution Function
CFD	Computational Fluid Dynamics
CNR	Carrier-to-Noise Ratio
CRPS	Continuous Ranked Probability Score
DBS	Doppler Beam Swinging
DOF	Degree Of Freedom
FFT	Fast Fourier Transform
FWHM	Full Width at Half Maximum
IFB	Institut für Flugzeugbau (Institute of Aircraft Design)
LES	Large Eddy Simulation
lidar	Light detection and ranging
LOS	Line-Of-Sight
MAE	Mean Average Error
MIC	Maximum Information Coefficient
NWP	Numerical Weather Prediction
PDF	Probability Density Function
PIT	Probability Integral Transform
PO	Project Organizer
PPI	Plan Position Indicator
PV	Photovoltaics
RAVE	Researach At Alpha Ventus
RHI	Range Height Indicator
RMSE	Root Mean Square Error
SCADA	Supervisory Control and Data Acquisition
SWE	Stuttgart Wind Energy
UTC	Universal Time Coordinated
VAD	Velocity Azimuth Display

List of Symbols

Greek letters

α	wind shear coefficient
Δ	range [m]
$\bar{\gamma}$	turbine tilt angle [°]
λ	diameter of eddy [m]
$\bar{\Theta}_v$	average potential temperature [K]

Roman letters

A	position matrix of measurement points [-]
b	offset of regression
d	distance to measurement point [m]
f	probability density function
F	cumulative density function
g	acceleration due to gravity [m s^{-2}]
h	smoothing parameter
K	Kernel function
m	slope of regression
n	counter
O	observation
p	pressure
P	power [% of rated power]
q	percentile
Q	quantile
r	vector of measured radial wind speeds [m s^{-1}]

R^2	coefficient of determination
Ri_s	Richardson number speed
s	vector of unknown wind components [m s^{-1}]
t	time
u	longitudinal wind component [m s^{-1}]
v	lateral wind component [m s^{-1}]
VEL	radial velocity [m s^{-1}]
w	vertical wind component [m s^{-1}]
x	longitudinal coordinate [m]
X	random variable
y	lateral coordinate [m]
z	vertical coordinate [m]

Subscripts

$(\cdot)_i$	referring to a specific measurement point
$(\cdot)_I$	referring to the inertial coordinate system
$(\cdot)_L$	referring to the lidar coordinate system
$(\cdot)_r$	referring to a wind ramp definition
$(\cdot)_W$	referring to the wind coordinate system

Operators

$(\cdot)^{-1}$	inverted
\sim	proportional

Abstract

With the introduction of renewable energies, the power grid has transformed from a centralised to a decentralised system. To balance the supply and demand of power in the energy grid at all times in spite of the volatile nature of wind and solar power, grid operators have to rely on accurate forecasts. However, state of the art wind power forecasting methods are not able to forecast changes of power in the minute-scale accurately. Therefore new methods are needed. This thesis investigates the use of a long-range lidar to forecast wind power on the minute-scale.

To that aim, two measurement campaigns were carried out. One was an onshore campaign, where the lidar was installed fixed on a radio tower next to a turbine that a forecast was made for. The second was an offshore campaign where the lidar was installed on top of the nacelle of a wind turbine. Both campaigns lasted over several months and the wind speed was measured in several kilometers in front of the turbine. During this time the turbine's own data system also recorded the 10-minute average power from the turbine.

In this thesis, a wind power forecast process is established. Lidar data is transformed from radial velocity to filtered horizontal wind speed and wind direction. The wind field information is then propagated to the wind turbine with an advection model based on Taylor's hypothesis. The forecasted wind speed at the turbine is then transformed into a forecasted power with the help of the power curve of the turbine.

To account for the uncertainty in the wind speed and power forecast, probabilistic forecast methods are applied. The results show that lidar-based forecasts at the offshore site are accurate in a forecast horizon up to ten minutes and outperform the benchmark forecast method persistence. Longer forecast horizons are biased because only small wind speeds measured further away from the wind turbine arrive with a delay of more than ten minutes. At the onshore site, persistence outperforms the lidar-based method in all forecast horizons, including the forecast horizon up to 10 minutes. The reason is that the Taylor based advection model does not model the actual propagation at the complex onshore site well enough.

During ramp events, the lidar-based forecast demonstrates its strength: information from the

wind speed measured a few kilometers in front of the turbine allows us to forecast changes of power. In comparison, persistence only uses old information and therefore cannot forecast any future changes. It is concluded that the added value of using a lidar for minute-scale forecasts lies in forecasting changes of power. As wind ramps are potentially critical to the grid stability, or can affect the cost of balancing the power system if they are not forecast well, using lidars at wind farms to improve the power forecast is advised.

However, challenges to the implementation of lidar-based forecasts remain. Lidar measurements depend on the aerosol content in the air and therefore the availability of the measurements for a forecast is not guaranteed. A fallback solution is needed such as statistical models or numerical weather prediction. To achieve forecast horizons of more than 10 minutes, the lidar measurement range needs to be extended beyond 10 kilometers. And to establish lidars as a state-of-the-art forecasting tool, standards are needed, which could be enabled by groups such as the IEA Wind community.

Wind lidar data coupled with propagation models and power curves has fundamental advantages for minute-scale wind power forecasting. Although this thesis has shown that current approaches may not be perfect, the rapid pace of wind lidar technology development, the increasing number of users, and the growing network of third party service providers, suggests that wind lidar is the future of minute-scale wind power forecasting.

Kurzfassung

Mit der Einführung der erneuerbaren Energien hat sich das Stromnetz von einem zentralen zu einem dezentralen System gewandelt. Um trotz der volatilen Natur von Wind- und Solarenergie das Angebot und die Nachfrage von Strom im Energienetz jederzeit auszugleichen, sind Netzbetreiber auf genaue Vorhersagen angewiesen. Die aktuellen Methoden zur Vorhersage der Windkraft sind jedoch nicht in der Lage, Leistungsänderungen im Minutenbereich genau vorherzusagen. Daher werden neue Methoden benötigt. In dieser Arbeit wird der Einsatz eines long-range Lidars zur Vorhersage der Windleistung im Minutenbereich untersucht.

Zu diesem Zweck wurden zwei Messkampagnen durchgeführt. Die erste war eine onshore Kampagne, bei der das Lidar fest auf einem Funkturm neben einer Windenergieanlage installiert wurde, für die eine Vorhersage gemacht werden sollte. Die zweite war eine offshore Kampagne, bei der das Lidar oben auf der Gondel einer Windkraftanlage installiert wurde. Beide Kampagnen dauerten mehrere Monate und die Windgeschwindigkeit wurde in mehreren Kilometern vor der Anlage gemessen. Während dieser Zeit zeichnete das anlageneigene Datenerfassungssystem auch die 10-minütige Durchschnittsleistung der Anlage auf.

In dieser Arbeit wird ein Verfahren zur Windleistungsvorhersage entwickelt. Die Lidardaten werden von der Radialgeschwindigkeit in die gefilterte horizontale Windgeschwindigkeit und Windrichtung transformiert. Die Windfeldinformationen werden dann mit einem Advektionsmodell, das auf der Taylor-Hypothese basiert, auf die Windenergieanlage übertragen. Die prognostizierte Windgeschwindigkeit an der Anlage wird dann mit Hilfe der Leistungskurve der Anlage in eine prognostizierte Leistung umgewandelt.

Um die Unsicherheit in der Windgeschwindigkeits- und Leistungsvorhersage zu berücksichtigen, werden probabilistische Vorhersagemethoden angewendet. Die Ergebnisse zeigen, dass lidarbasierte Vorhersagen am offshore Standort in einem Vorhersagehorizont von bis zu zehn Minuten genau sind und Persistenz als Benchmark-Vorhersagemethode übertreffen. Längere Vorhersagehorizonte sind fehlerbehaftet, da nur kleine Windgeschwindigkeiten, die weiter entfernt von der Windkraftanlage gemessen werden, mit einer Verzögerung von mehr als zehn

Minuten eintreffen. Am onshore Standort übertrifft die Persistenz die lidarbasierte Methode in allen Vorhersagehorizonten, einschließlich des Vorhersagehorizonts bis zu 10 Minuten. Der Grund dafür ist, dass das Taylor-basierte Advektionsmodell die tatsächliche Ausbreitung an dem komplexen onshore Standort nicht gut genug abbildet.

Bei Rampenereignissen spielt die lidarbasierte Vorhersage ihre Stärke aus: Die Information aus der wenige Kilometer vor der Anlage gemessenen Windgeschwindigkeit erlaubt es, Leistungsänderungen zu prognostizieren. Im Vergleich dazu nutzt die Persistenz nur alte Informationen und kann daher keine zukünftigen Änderungen vorhersagen. Es wird gefolgert, dass der Mehrwert der Verwendung eines Lidars für Prognosen im Minutenbereich in der Vorhersage von Leistungsänderungen liegt. Da Windrampen potenziell kritisch für die Netzstabilität sind oder die Kosten für den Ausgleich des Stromsystems beeinflussen können, wenn sie nicht gut vorhergesagt werden, ist der Einsatz von Lidaren in Windparks zur Verbesserung der Leistungsvorhersage ratsam.

Allerdings bleiben Herausforderungen bei der Implementierung von lidarbasierten Vorhersagen bestehen. Lidarmessungen sind abhängig vom Aerosolgehalt in der Luft und daher ist die Verfügbarkeit der Messungen für eine Vorhersage nicht garantiert. Es wird eine Ausweichlösung benötigt, wie statistische Modelle oder numerische Wettervorhersagen. Um Vorhersagehorizonte von mehr als 10 Minuten zu erreichen, muss der Lidarmessbereich außerdem auf mehr als 10 Kilometer erweitert werden. Und um Lidare als modernes Vorhersageinstrument zu etablieren, werden Standards benötigt, die durch Gruppen wie die IEA Wind Community ermöglicht werden könnten.

Windlidardaten, die mit Ausbreitungsmodellen und Leistungskurven gekoppelt sind, haben fundamentale Vorteile für die Vorhersage von Windenergie im Minutenbereich. Obwohl diese Arbeit gezeigt hat, dass die derzeitigen Ansätze nicht perfekt sind, legen die rasante Entwicklung der Wind-Lidar-Technologie, die steigende Anzahl von Nutzern und das wachsende Netzwerk von Drittanbietern nahe, dass Windlidar die Zukunft der Windleistungsvorhersage im Minutenbereich ist.

Introduction

1.1 Motivation

The use of renewable energies is an essential part of reducing climate-changing emissions. During the last years, the energy grid transitioned from centralized and manageable energy production with conventional sources such as coal, nuclear and gas, to decentralised energy production from solar and wind power.

Both technologies work with energy sources that nature provides for free, but they come with a catch. Solar and wind power are fluctuating energy sources and are governed by local changes. Without intermediate storage, they do not provide a steady energy output. However, in order to keep the energy grid stable, the supply from power plants and the demand from consumers need to be balanced at all times. As a result of the energy transition, weather forecasting became a crucial tool to tackle the challenge of grid balancing, because it helps to manage the variable energy supply.

If one looks up weather forecasting in Wikipedia, it says that “Weather forecasting is the application of science and technology to predict the conditions of the atmosphere for a given location and time.” In fact, people have tried to forecast the weather for millennia and since the 19th century national weather services successfully forecast meteorological conditions [1]. With the introduction of renewable energies into the energy grid, weather forecasts became important to forecast wind and solar power.

Wind power generation forecasts hours or days ahead of real time are based on weather forecasts which are calculated with Numerical Weather Prediction (NWP) models. NWP models are accepted as baseline forecasts and use input data such as temperature, humidity and pressure from weather stations and simulate future weather conditions by solving physical and mathematical equations [2]. NWP models are computationally expensive and so tend to be run by large governmental organisations, such as the Deutscher Wetterdienst (German Meteorological Service), or commercial providers.

However, for wind power forecasts up to 60 minutes, different forecast methods are needed because due to their spatial resolution of several kilometers, NWP models are not accurate for short time scales. The state-of-the-art is to use statistical methods that are based on historic measurement data of wind and power at the target wind turbine or wind farm. These forecasts are reliable, cheap and do not need much computing power [3], but have a disadvantage. Because they use old measurement information to predict the future performance, they are not able to predict changes in power generation.

The successful integration of wind energy into the grid therefore requires a different approach to wind and power forecasting. Ideally this would be an accurate forecast that provides temporally-resolved data at the point of interest and can be updated within a few seconds or minutes.

Wind lidar may be an ideal tool for this application. Long-range lidars can measure the wind speed several kilometers upstream of a wind turbine or wind farm. This preview information of the wind speed that will affect the turbine's power generation, can be used to generate power forecasts for the next minutes. The lidar is able to measure the variation in wind speed and should therefore be able to predict changes in generated power more accurately than conventional time series based methods.

This thesis will therefore investigate the use of a long-range lidar to predict the power output of a wind turbine in the minute-scale.

1.2 Research areas in this thesis

Lidars are used for many different applications in wind energy such as site assessment, turbine or wind farm control and power curve measurements [4]. All applications have one thing in common: the lidar data needs to be processed to be usable. This is the case for all data from measurements, and not specific to lidar data. The goal of lidar data processing for wind power forecasting is to make sure that only accurate wind measurements are used to generate the forecast.

The first step is to check for outliers in the data and filter the data accordingly. Lidars measure radial wind speed, so the second step is to apply wind field reconstruction methods to retrieve wind parameters such as horizontal wind speed and wind direction from the measurements [5]. When the work for this thesis was started, long-range lidars were still new on the market and there was not much experience in terms of data processing. Therefore, the processing chain first had to be developed in this thesis based on existing methods for ground-based and short-range nacelle-based lidar measurements.

Having chosen wind lidar as the tool to deliver the wind data, it is necessary to consider how to convert this into a forecast.

The forecasting method that is applied in this thesis is probabilistic forecasting. Probabilistic forecasting is also called uncertainty forecasting and aims to provide a forecast of a value and quantifies the uncertainty of the forecast [6] at the same time. The goal is to help and facilitate decision making processes. For wind power forecasting, this means that not only the power output of a wind turbine is forecast for a specific time in the future, but also a probability that this power will be generated is given. Probabilistic power forecasting is able to quantify the uncertainty resulting from the volatile nature of the wind, which is responsible for the power generation. This thesis builds on existing methods for probabilistic forecasting and adapts them for the application of lidar-measurements.

In parallel to the work carried out for this thesis, similar research was carried out by colleagues at other institutes. In 2019, Elliot Simon finished his PhD thesis on "Minute-Scale Wind Forecasting Using Lidar Inflow Measurements" [7]. He focused on using scanning-lidar to examine space-time correlations of wind patterns measured onshore, upstream of a lidar. In the same year Laura Valdecabres defended her PhD thesis. She showed in her work that lidars can forecast wind speeds at near-coastal conditions [8] and then focused on using radars to carry out minute-scale probabilistic power forecast [9], [10] and forecast ramp events [11]. The author of this thesis worked together with Elliot Simon and Laura Valdecabres and together they organised a collaborative workshop between the IEA Wind Tasks 32 and 36 on "Very short-term forecasting of wind power", which took place in 2018. The outcome of this workshop was an overview paper of methods for minute-scale forecasting, led by the author of this thesis [12] and reported elsewhere in this thesis.

1.3 Research objectives, methodology and organization

This thesis aims to improve minute-scale forecasts of wind power using wind speed data from a long-range lidar and power data from wind turbines. The following research questions arise:

- How should the lidar data information be processed to gain a power forecast for a wind turbine?
- How does the measurement setup and measurement site influence the forecast?
- What is the forecast horizon of lidar-based forecasts and what influences the forecast horizon?
- How does the lidar-based forecast perform in comparison to state-of-the-art statistical methods and what are its benefits?

The data needed to answer these questions were collected in two measurement campaigns. In one campaign a wind lidar was installed next to an onshore wind turbine. In another, a wind lidar was installed on top of an offshore wind turbine. Both campaigns were carried out over a period of several months. Simultaneous to the lidar data, turbine data was made available by the operator and used in the forecast process to generate the power curve of the turbines and validate the forecasts. At both sites, data from a meteorological mast was available which recorded meteorological conditions.

The methods used in this thesis to process the data and to generate the forecasts were previously established for other applications. This thesis introduces each method, discusses its context and previous application, and explains how it was adapted to be used for lidar-based minute-scale forecasting in this work. Assumptions that were made for simplification are stated clearly.

The methodology applied in this thesis is valid for the power forecast of a single wind turbine. To forecast the power output of a wind farm, the methods applied here need to be extended and wind farm effects need to be considered.

The thesis starts by explaining background information in Chapter 2. An overview of the use of lidar in wind energy applications is given and the need for minute-scale forecasts is explained by discussing the variability of intra-hourly wind power generation. By describing the state of the art forecasting methods, the gap that lidar-based forecasts can close is described.

The forecast chain that is established in this thesis is explained in Chapter 3. The chain is an overview of which steps are taken in order to process lidar data and to gain a wind power forecast of a wind turbine. The thesis structure then follows the chain links step by step and explains the processing steps.

- In the first link of the forecasting chain, wind speed measurements from the lidar are converted into horizontal wind speeds, which are the basis for the forecast. The methods for filtering lidar data, the wind field reconstruction, as well as the measurement setup are explained in Chapter 4.
- In order to forecast the power of a wind turbine, these distant measurements need to be propagated through time and space to the wind turbine's location. This is the second link in the forecasting chain and the use and implementation of a propagation model are explained in Chapter 5.
- The final link in the forecast chain is to transform the predicted wind speeds into predicted power, and to calculate the forecasts. The methods applied are explained in 6 and the results of the forecasts from an onshore and offshore campaign are compared and discussed.

The challenges for the implementation of lidar-based minute-scale forecasting are discussed in Chapter 7. To that end first the correlation of the measurement range and the forecasting horizon is analysed, and then solutions to overcome barriers are suggested and discussed.

The thesis concludes with Chapter 8, where conclusions and recommendations for further research are given.

Background

This chapter explains the background that is needed to understand the need for minute-scale forecasting and the reason why a lidar was chosen as a tool. First the reason for the intra-hourly variability of wind power is explained in Section 2.1 and the need for minute-scale forecasting in different application areas that arises from the power variability is discussed in Section 2.2. The use of lidar in different wind energy applications is explained in Section 2.3. In Section 2.4 the state-of-the-art forecasting methods are described and the need for lidar-based forecasting is derived. Section 2.5 summarises the key information needed to understand the background of this thesis.

2.1 Intra-hourly variability of wind power generation

The need to forecast wind power on timescales in the order of minutes arises from the variability of the wind speed in these time scales. The cause for the variability can be found in atmospheric phenomena with spatial scales from one up to tens of kilometers that introduce the fluctuations in the wind. There are four dominant phenomena [13]:

- Open cellular convection is an offshore phenomenon that occurs when cold air is convected over warm sea water. Clouds are formed that have a honeycomb-like structure and rising air in the cloudy cell walls leads to a horizontal wind speed variability. When such cells move over an offshore wind farm, the farm experiences the spatial wind speed variability as a temporal variability.

- For coastal wind farms the land and sea breeze can lead to power fluctuations in the minute scale. The breeze is caused by the surface temperature difference of mainland and water causing sea breezes during the day and land breezes at night. Although the sea breeze circulation is predominantly a diurnal effect, it can also influence smaller scales when, for example, convergence of the sea breeze front with the background flow results in squall lines or showers near the coast.
- Gravity waves in the atmosphere are a phenomenon that occurs over land and over sea. Gravity waves are waves generated in a fluid medium or at the interface between two media when the force of gravity or buoyancy tries to restore equilibrium. They cause wave-like cloud structures in the atmosphere and similar wave structures in the wind speed pattern with amplitudes of several meters per second.
- Low level jets are regions of increased wind speed in the lower hundred meters of the atmospheric boundary layer. They occur onshore and offshore and are of special interest to wind energy as the extreme wind shear over the rotor causes extreme loading on the turbine and unexpectedly high power generation. As the mechanism for low-level jet formation is not fully understood, they are difficult to forecast.

All these phenomena cause fluctuations in the wind speed in the minute scale that are augmented to the third power when transferred to the power generated by a wind turbine or wind farm. In Figure 2.1 an example time series of one day of wind speed and generated power of a single wind turbine is given. Small fluctuations in wind speed result in large fluctuations in generated power. The figure also shows phenomena referred to as ramp events. These rapid and strong changes in generated power are caused by extreme changes in wind speed or direction and are associated with the passage of weather fronts. Ramp events can also have other causes than extreme changes in wind speed. Ramp events are especially critical around cut-out wind speed. If the cut-out wind speed of a wind farm is reached the whole wind farm shuts down. In this case a small fluctuation in wind speed can lead to a strong decrease in the power production.

Ramp events are critical for the grid balancing as they are difficult to forecast and thus often occur unexpectedly. In [14] an overview of recent ramp forecasting techniques is presented. The article also points out that despite being critical, there is no common definition of ramp events. Instead, the definition usually depends on location as well as the size of the considered wind turbine or wind farm. However, the basic parameters that need to be forecast are clear: Δt is the duration of the ramp and ΔP is the minimum change of power within Δt . Positive ΔP implies an upward ramp, negative ΔP a downward ramp. Downward ramps tend to be more critical for grid stability than upward ramps. A downward ramp causes an energy shortage that needs to be compensated using available power generations, while an upward ramp can




Figure 2.1: Example time series of wind speed and generated power of a single wind turbine with wind ramps marked for a time window of 60 min and a change of power of 40%. Each data point in the time series corresponds to a 10-minute average. Reproduced without modifications from Würth et al. [16] with permission.

be managed by curtailing the wind farm [15]. Wind ramp forecasting errors are specified as level errors and phase errors. Level error describe the magnitude of the power change that was forecast inaccurately, while phase errors describe the deviation in time when the ramp occurred.

2.2 Application areas for minute-scale forecasting

Minute-scale forecasts of wind speed and power are needed in three application areas in the wind energy sector [12]:

- Wind farm control: the controller of the wind farm uses preview information of the wind speed and direction to optimize the power output and reduce loads.
- Power grid balancing: the Transmission System Operators (TSO) need information about changes of the produced power in their grid to balance power production and load at all times and to manage power reserves.
- Energy and ancillary services markets: wind power energy trade takes place in intra-day markets and forecasts are needed to reduce imbalance costs and increase revenue.

Each application area is associated with a different minute-scale forecast horizon in the minute scale (Figure 2.2).

Wind turbine and wind farm control uses forecast information of the wind parameters of




Figure 2.2: Overview of forecast horizons of different wind energy applications in the second and minute scale. Reproduced with modifications from Würth et al. [12] with permission.

up to 10 minutes before they arrive at the turbine or wind farm. These forecasts can be used to control wind turbines by modifying turbine alignment or blade pitch. For instance for yaw control, an estimation of the wind direction is used to align the rotor with the wind direction. The preview wind information can also be used for wake steering which uses the misalignment effect to steer the wake of a wind turbine away from a downstream turbine.

TSOs have the task to balance the power supply and power demand at all times to keep the frequency of the grid stable at 50 Hz. If the power production differs from the demand, balancing actions are needed. A shortage of power supply (or a rise in consumption) would otherwise lead to a frequency drop, a rise in power supply (or a drop in consumption) would lead to a frequency increase. To balance the supply and demand of power all times TSOs therefore need a control reserve, which is also called balancing power. There are three different control mechanisms established to activate the balancing power and they are categorized according to the activation time: primary control (within 30 sec), secondary control (within 5 min) and tertiary control (directly activated or supplied in schedules of up to 4 times 15 min). Forecasts are therefore useful in time ranges of up to 60 minutes to reduce the need for balancing actions and to lower the amount of balancing power that a TSO needs to keep on hold.

Short-term wind energy trading takes place in intra-day markets. The regulations for the market depend on the country and also the lead times are country specific. In Germany and Australia for example the lead time for the trade is 5 min and the trading block of power is 15 min. This means that if a TSO or wind farm operator wants to sell their produced power on such a market, they need to know the amount of power they can sell as accurately as possible in advance. Forecasts are therefore necessary in the range of 5 min to 60 min. Minute-scale forecasts help to reduce the risk for penalties, that arise when the actual produced power differs from the power that was traded. This can happen if the produced power differs significantly from the forecast.

Apart from the wind energy sector, wind speed forecasts are also relevant for other sectors, e.g.

the construction sector. Wind speed forecasts could help increase the safety during commission and operations on construction sites. If it is known in advance that a gust is coming, accidents can be avoided if heavy objects are lifted by a crane.

2.3 Lidar in wind energy applications

Doppler wind lidars use the Doppler frequency shift of light to measure the speed of airborne particles. The measuring principle is based on the backscattering of the emitted laser light to a receiver [17]. The backscattering takes place at particles in the air, so-called aerosols, which scatter back the light frequency shifted by their own speed. This frequency shift is based on the Doppler effect. Due to the fact that the Doppler shift only affects particles moving in the direction of the laser beam, only this directional component of the wind speed can be measured. The radial wind speed is known as the line-of-sight wind speed and is denoted v_{los} .

The aerosols that the lidar measuring principle is based on are [18]:

- sea salt aerosols, especially near the sea.
- dust aerosols, mineral origin
- secondary aerosols, nitrates and sulfates
- biological aerosols, e.g. fungal spores and pollen
- smoke aerosols, from forest fires or anthropogenic caused
- volcanic aerosols, especially in higher air layers

These aerosols with a size in the order of the laser wavelength of 1.5 μm are particles that have very low settling velocities and thus are suspended by the wind. The quantity of the individual aerosols in the air however is strongly dependent on the location. The optical properties of the particles are subject to strong fluctuations. By accumulation of moisture from the air, for example, the diameter of the particles can fluctuate strongly and thus change the backscattering properties [18]. As a result, the lidar measurement properties are subjected to fluctuations and depend on the existence and the properties of the aerosols.

Lidars are used for different applications in wind energy such as site assessment, turbine or wind farm control and power curve measurements [4]. For each application, a different lidar type is used, because the requirements for the wind measurements vary. For site assessment, ground-based vertical measuring lidars are ideal, which measure the wind speed and wind direction over several heights [19]. For turbine control, nacelle-based lidar systems are used, which use the preview information from the incoming wind field several hundred meters in front of the rotor, to determine control actions at the wind turbine to optimize the power performance and

reduce loads [20]. For wind farm control, forecasting of the wind is necessary, but not on a minute-scale, but on a second-scale [12]. Ground-based, and nacelle-based lidars are also used for power curve assessment to correlate the inflow wind speed over the whole rotor area to the generated power of the wind turbine [21]. For this application long-range lidars are used and mounted e.g. on the transition piece of offshore wind turbines, to measure the inflow wind speed outside of the induction zone of the turbine over several heights [22]. Minute-scale forecasting is a new application for wind lidars, and so far it is not established as common practice. All investigations conducted so far used long-range scanning lidars to measure the upstream wind velocity and forecast future power generation [12].

2.4 How is wind power forecast today?

The classic method to generate a wind power forecast comprises one or several of the following ingredients (Figure 2.3):

- A numerical weather prediction (NWP) model which computes the state and evolution of the atmosphere e.g. weather forecast,
- Observational data, e.g. Supervisory Control and Data Acquisition (SCADA) data from the wind turbine or wind farm that is either used to compute a forecast using statistical time series methods or serves as input for the forecast model,
- Terrain information and information about the wind farm layout,

These ingredients are combined using a forecast model that uses the information from NWP, observational data, and terrain information to generate a power forecast for a wind turbine or wind farm [23, 2].

In the following the ingredients are explained in more detail and their limitations for minute-scale forecasts are discussed. The goal is to derive the need for the lidar-based forecasting method. But first, the forecast terminology is clarified in the next section.

2.4.1 A note on terminology

When talking about forecasting of wind power, the concept of the forecast horizon is often mentioned. The forecast horizon describes the time period or point in time in the future, for which the forecast is generated. When it comes to forecast horizons from seconds to a few hours, the terminology used for the description varies between nowcasting, very short-term, and short-term forecasting with no clear definition of the length of the time period. Therefore, for the sake of clarity, it was decided in the collaborative IEA Wind Task 32 and 36 Workshop on




Figure 2.3: Elements in the forecasting approaches. Reproduced without modifications from Giebel et al. [2] with permission.

“Very short-term forecasting of wind power”, to use the exact forecast horizon as a description. Therefore, this thesis only uses the term minute-scale forecast which describes a forecast horizon from one minute up to one hour. Longer forecast horizons would then be called hour-ahead or day-ahead forecasts.

2.4.2 Numerical weather prediction

Numerical weather prediction models divide the atmosphere up into cells. Physical models are used to describe the state of the atmosphere in each of these cells. The parameters describing the atmosphere can then change over space and time. Depending on the size of the cells and the domain that is covered, the models are roughly divided into different classes (Figure 2.4). More accurate physics or bigger domains require more computational effort. As a result, forecasting the weather over large areas and longer periods of time is usually only possible for large organisations such as national weather services [23].

It is possible to provide high resolution time series forecasts by decreasing the size of the domain, or reduce the details of physics (but this might lead to increased uncertainty). Minute-scale forecasting requires a combination of high spatial resolution and domain sizes in the order

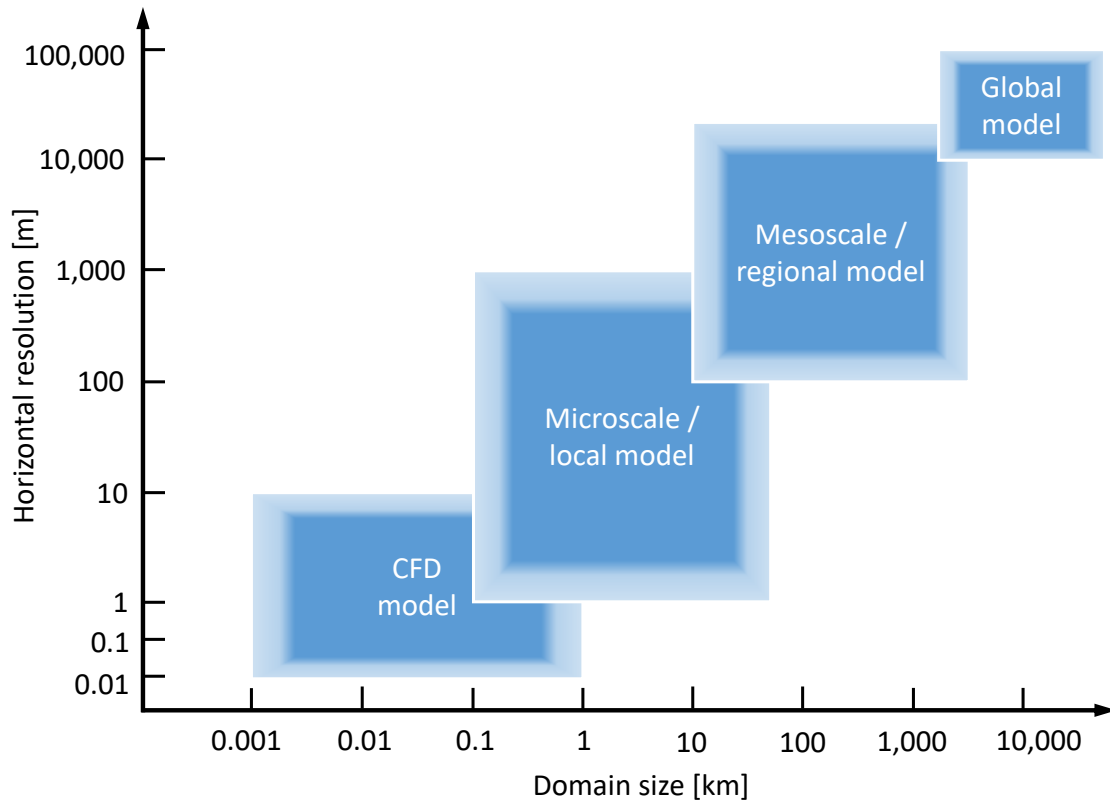


Figure 2.4: Different classes of weather models. The borders of the boxes are fuzzy, because the size specifications are to be understood as approximate values.

of 10 km to 50 km. This makes numerical weather forecasting challenging, because it requires a combination of high resolution and comprehensive physics that is computationally expensive.

NWP models also require boundary and initial condition data to deliver accurate minute-scale forecasts. Unfortunately these data are often not available for wind farms which makes it difficult to provide accurate forecasts. NWP forecasts are usually time consuming, meaning that it may be unable to deliver a minute-scale forecast in time for the operators to take actions. For these reasons, minute-scale forecasting is more typically based on simple algorithms based on the available on-site data.

2.4.3 Statistical time series models

Statistical approaches to forecasting mainly rely on deducing patterns from past observational data and extrapolating these relationships to predict future values over a desired time step [12]. Forecasts in wind energy are carried out for one dimensional time series signals such as a wind speed measurement, or SCADA data such as wind turbine or wind farm active power signal. The chosen forecast horizon should relate to the time resolution of available input data, and at a minimum be one sample (time step) ahead to avoid errors introduced by interpolation.

Minute-scale statistical forecasting methods are largely identical to techniques employed for longer horizons. The main differences are the temporal resolution of the data and the variability of the physical process being predicted [12].

Benchmark statistical time series models are persistence and climatology. Persistence is a very simple forecast method and assumes that the forecasted conditions are the same as the present conditions. This means the most recent power measurement of a wind turbine is used for the power forecast. Climatology uses statistics from historic measurements, e.g. an average of the last n hours of measured power generation, to create a forecast.

2.4.4 Where is the gap that needs to be closed?

NWP models are optimized to produce forecasts in the hour- and day-scale. They can produce weather forecasts for up to 15 days ahead. However, they are not very accurate in the minute-scale, where statistical time series models perform better (Figure (2.5)). In fact, persistence frequently outperforms hour-ahead or day-ahead forecasts in the time range up to 60 min.

This means new methods for minute-ahead forecasts need to be more accurate than persistence and persistence is therefore the benchmark.

Persistence however has one disadvantage: it uses historic measurements to forecast future events. This approach produces large errors if the future event deviates significantly, such as a wind ramp. Considering that these changes in power are crucial information for TSOs who need to balance the grid, or wind farm operators selling their produced power, a better forecast method than persistence is needed.

It is therefore of great interest to investigate if lidar-based forecasting can close the gap of forecasting power changes in the minute-scale and outperform the persistence model.

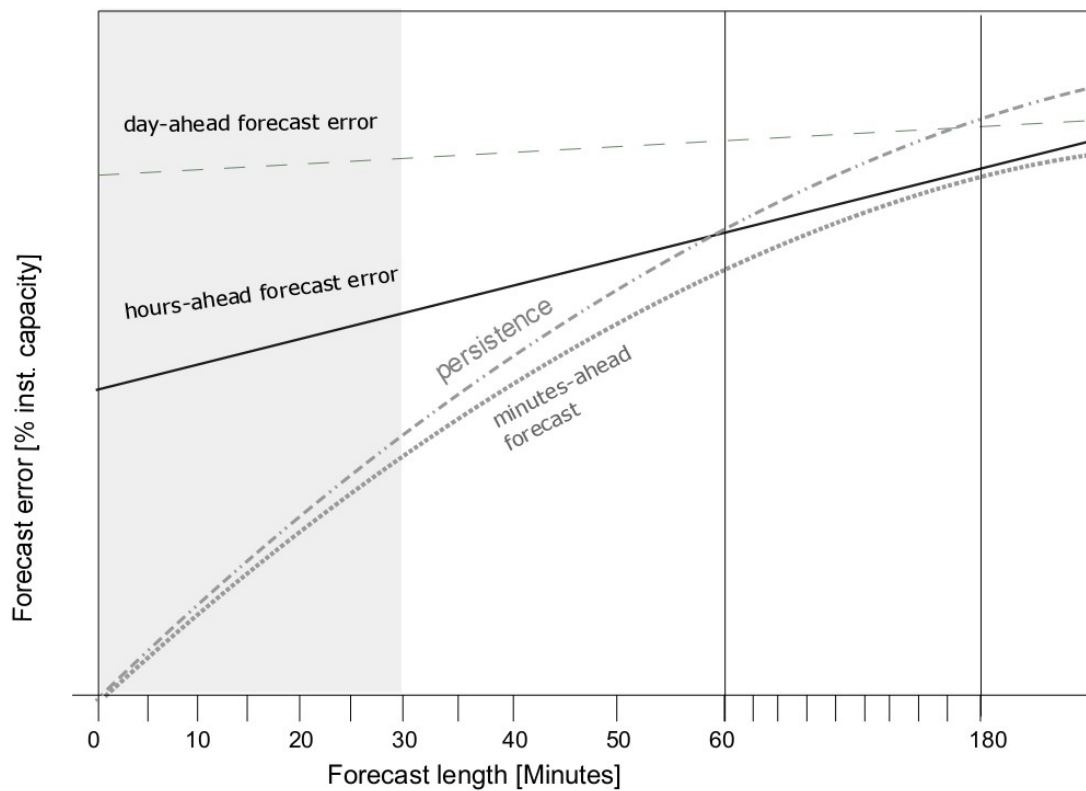


Figure 2.5: Qualitative visualization of the forecast error development over the first hours of a forecast for different temporal forecast techniques. Reproduced without modifications from Würth et al. [12] with permission.

2.5 Summary

Minute-scale forecasts of wind speed or power are important for TSOs to keep the grid stable and to reduce balancing power, or for wind farm operators to optimize wind farm control or energy trading if they sell power on rolling markets.

State-of-the-art NWP models are optimized to forecast wind conditions in the hour and day-scale. Time-series based forecasts produce more accurate minute-scale forecasts than NWP models, but rely on historical measurement data from the wind farms to forecast the power. Therefore they cannot forecast large changes in future power output, e.g. from wind ramps.

Long-range lidars measure the wind speed remotely and can be used to measure the upstream wind speed of a wind turbine or wind farm. Lidar-based measurements therefore contain pre-view information of future wind speed changes. Therefore it is of great interest to investigate if lidar-based minute-scale forecasts are able to outperform conventional methods such as persistence.

The forecasting chain: from radial velocity to minute-scale wind power forecasts

In order to gain minute-scale forecasts of wind power of a wind turbine from long-range lidar data, a forecasting chain had to be established in this thesis. The thesis will follow the chain link by link in the next chapters and explain the steps in detail.

This chapter presents an overview of the steps and serves as orientation. In principle, the chain shows that data exists in one form and needs to be processed in order to reach a new form of data (Figure 3.1). An overview of the data forms and the steps of processing is given in the following.

- **LOS.** Lidars measure the wind speed in Line-Of-Sight (LOS) direction along the laser beam and therefore measure the radial component of wind field in this direction.
- **Wind field reconstruction.** As the lidar measures the radial component of the wind field at each measurement point, the three wind vector components u, v, w need to be deduced from this measurement at each point. This process is called wind field reconstruction. In order to reconstruct the wind field components from the LOS measurements, assumptions have to be made and algorithms have to be applied to the data, see Chapter 4.

- **Wind field.** Long-range lidars are able to measure the wind field simultaneously in several measurement points up to a range of several kilometers. After wind field reconstruction, the wind field components u, v, w and the wind direction in those points are known. This data of the wind field is the basis for the minute-scale forecasting, because the measured wind field contains the preview information of the wind conditions the turbine will experience in the following minutes, if the lidar measures the inflow of the wind turbine.
- **Propagation model.** The wind vectors measured in the distance are transported through space and time. This is known as wind field propagation. A model needs to be applied to the measured wind field data that determines how the propagation is happening. With the help of the propagation model, the goal for minute-scale forecasting is to determine which wind vector measured in the inflow of the turbine will reach the turbine and at which future point in time this will happen, see Chapter 5.
- **Predicted wind speed at turbine.** With the help of the propagation model, the measured wind field is transported through time and space and the predicted wind speed at turbine level is determined. This means, the wind conditions at the turbine location for the minutes after the forecast is issued are known. This wind speed forecast contains information from the lidar measurements several minutes prior to when the forecast is issued.
- **Power curve.** In order to gain the power forecast of the turbine, the predicted wind speed at the turbine needs to be converted to power. This is achieved through means of the turbine's power curve. The power curve is a unique property of each turbine type and sets the inflow wind speed in relation to the power the turbine produces for this wind speed. The power curve of a turbine is determined in a measurement campaign using free stream wind speed measurements and measured power values.
- **Predicted power.** The predicted wind speed of the turbine is converted into the predicted power using the turbine's power curve. Similar to the predicted wind speed, this means the power output of the turbine for the minutes after the forecast is issued is known. The forecast can be generated for different forecast horizons, depending on the desired horizons and depending on for how many minutes in the future the forecasted power is actually available, see Chapter 6.

This chain represents a very high level overview. To get from link to link, many more steps need to be considered, which is the task of the next chapters. It should also be noted that in this thesis the forecast of only one wind turbine is investigated. However, if the forecast is extended to wind farm level in the future, the chain can be extended with considerations of taking into account wake effects to aggregate the predicted power of several wind turbines.





Figure 3.1: The lidar forecasting chain. Data exists in one form (circle) and needs to be processed (arrow) in order to each a new form of data. Reproduced without modifications from Würth et al. [16].

First step in the forecasting chain: from radial velocity to wind field



In the first step in the lidar forecasting chain, the radial wind speed measurements from the lidar are converted in several processing steps into horizontal wind speeds, which are the basis for the forecast. This chapter explains the processing steps but starts with an overview of the measurement setup in Section 4.1. Section 4.2 then explains the methods applied for filtering and wind field reconstruction in order to obtain a useful wind speed and wind direction signal from the lidar.

4.1 Measurement setup at the onshore and offshore site

The data used in this thesis were generated in the projects VORKAST and ParkCast. VORKAST was a German national funded research project to optimise the design and operational management of hybrid power plants and energy storage technologies by means of wind and Photovoltaics (PV) power minute-scale forecasting [24]. The focus of the project was on onshore sites for wind and PV power forecasting. It was running from September 1st, 2014 to October 31st, 2017. ParkCast is the follow-up project to VORKAST and started in November 2018 with the goal to develop, optimize and evaluate new methods for minute-scale forecasts of offshore wind farms. ParkCast will end in October 2021. The author of this thesis was Project Organizer (PO) for the University of Stuttgart's contributions for both projects.

To describe the measurement setup that was used in the projects and that resulted in data that are used in this thesis, a 10-step methodology has been applied that was introduced by Vasiljevic et al. in 2017 [25]. This methodology provides guidance on how to carry out a lidar measurement campaign in order to ensure its success. The next sections will follow these steps and explain how they apply to the present campaigns, in order to give the reader a comprehensive understanding of the measurement setup.

4.1.1 Definition of scientific objectives

Three scientific goals drove the planning of the VORKAST measurement campaign. First, the idea for the project was born in a time where commercial lidar measurement systems were first brought on the market that had an extended range and were able to measure the wind speed in distances of several kilometers. Therefore the first goal of the project (and thus the measurement campaign) was to test one of the new systems for functionality, and its applicability for minute-scale forecasting.

Second, during the campaign the lidar should be able to measure as far as possible. The goal was to find out how to set up the lidar and carry out the measurements in order to obtain the best wind speed measurement for forecasting. This included the development of laser scan strategies to extract the wind field components from the radial wind speed, testing the data analysis process, and the maximum measurement range that could be reached.

Third, the goal was to carry out minute-scale forecasts of wind speed and also wind power of an onshore wind turbine using the data from the long-range lidar. To that end, new forecasting methods using the lidar data should be developed, which had to take into account the dynamic variability of the wind. Thus the forecasts should be able to capture the variability of the power fluctuation.



Figure 4.1: Measurement sites marked on map of Germany. Reproduced with modifications from [27] with permission.

In ParkCast, the results from VORKAST should be used and transferred to an offshore wind farm. The goal was to investigate how the lidar-based minute-scale forecast methods developed for one turbine onshore could be transferred to a wind farm with several turbines offshore and how the forecasts perform under offshore conditions.

It should be pointed out, that this thesis compares the methods developed for minutes-scale forecasts for one turbine onshore and offshore and does not take into account the forecasts of the whole offshore wind farm.

4.1.2 Site selection

The site that was selected for the onshore measurement campaign was near Stötten in the south of Germany (Figure 4.1 bottom marker). The site was chosen due to its proximity to Stuttgart and prior use in related studies [26]. The proximity to the institute was important, because for this campaign a new lidar system had to be tested. When testing a new measurement system, it has proven beneficial to have easy access to be able to adjust its settings or repair it in case of failure. Due to related studies, the site also offered easy access to local wind turbine data and access to meteorological data from the institute's meteorological (met) mast.

As offshore site the alpha ventus wind farm in the Northern Bight 45 km north of the German island of Borkum was chosen (Figure 4.1 top marker). For ParkCast it was important to find a site that offered the possibility to install the lidar and get access to turbine data and also meteorological data. In 2003 the research platform FINO 1 (Forschungsplattformen in Nord- und Ostsee) was erected at the site, measuring the meteorological conditions with a met mast. In 2007 the German Federal Ministry for the Environment, Nature Conservation and




Figure 4.2: Wind characteristic in Stötten measured at the met mast. (a) Reproduced with modifications from Hofsäb et al. [26] with permission.

Nuclear Safety (BMU) launched the RAVE (Research At Alpha Ventus) research initiative. The initiative’s goal was to support and facilitate research projects at the offshore wind farm. In 2009 alpha ventus started operating with 12 turbines and became Germany’s first offshore wind farm [28]. University of Stuttgart is member of the RAVE initiative and has conducted several research projects at the site. Access to the wind farm data and the FINO 1 data is obtained through data portals hosted by the Federal Maritime and Hydrographic Agency of Germany (BSH). Alpha ventus was chosen as a site for the ParkCast project, because as a RAVE member, easy access to the turbines to install and maintain a lidar and access to turbine and meteorological data was guaranteed.

4.1.3 Site characterization

Stötten is in the Swabian Alps; the location is a very hilly area consisting of high plateaus surrounded by a pronounced 100 – 150 m tall wooded escarpment known in the region as the Albtrauf. A detailed study of the local meteorology [26] shows that the main wind direction is west to north-west and the most frequent wind speed is around 5 m s^{-1} (Figure 4.2). The measurements described in [26] are centred on a 100 m high met mast in relatively flat land less than 1 km easterly from a section of Albtrauf (Figure 4.3). On the plateau several wind turbines and a radio tower are located.

alpha ventus is an offshore wind farm consisting of 12 turbines with two different turbine types set out in a 4-by-3 grid (Figure 4.4 (c)). The two northern rows are 5 MW turbines of type REpower 5M, with a rated power of 5 MW, a hub height of 92 m, and a rotor diameter of 126 m. The two southern rows are 5 MW turbines of the type Adwen AD 5-116, with a rated power of 5 MW, a hub height of 90 m, and a rotor diameter of 116 m. The turbines are enumerated row






Figure 4.3: Map of the measurement site in Stötten. Map data: ©OpenStreetMap-Mitwirkende, SRTM | map display ©OpenTopoMap (CC-BY-SA)

wise to identify them, starting with the AV01 which is the turbine at the top left corner, and AV12 which is the turbine at the bottom right corner. FINO 1 is located directly west of the AV04 turbine in a distance of 405 m.


alpha ventus was the first wind farm installed in the area, but not the last. Around the site, several other wind farms were commissioned after alpha ventus (Figure 4.4 right) and they influence the inflow conditions of alpha ventus. With 60 MW total capacity alpha ventus is small in comparison to the wind farms around it. Direct neighbors of alpha ventus are Borkum Riffgrund I with 312 MW capacity and Borkum Riffgrund II with 448 MW capacity in the south-west, and Merkur with 396 MW capacity in the north-west. The closest distance between alpha ventus and the surrounding wind farms is around 2 km. The main wind direction at alpha ventus is south-west (Figure 4.4a), the mean wind speed at 91 m is 8.52 m s^{-1} (Figure 4.4b). The wakes of the surrounding wind farms in the main wind direction affect the power production of alpha ventus and lead to an increased turbulence intensity at the site [29].



(a) Wind rose



(b) Weibull distribution



(c) Layout of alpha ventus. Reproduced with modifications from [30].



(d) alpha ventus surrounded by other wind farms. Reproduced with modifications from [31].

Figure 4.4: Wind characteristic in alpha ventus measured at FINO1 at 91 m (top) and layout of alpha ventus situated in the North Sea (bottom).



Figure 4.5: StreamLine XR lidar on the top level platform of the radio tower.

4.1.4 Experiment layout design

The lidar system that was used for the measurement campaign in VORKAST is a StreamLine XR pulsed doppler scanning-wind-lidar from the company Halo Photonics. The lidar was chosen because of its measurement range of 10 km and because of its light weight and compact form. The lidar was mounted on the top level platform of the radio tower (Figure 4.5) at a height of 736 m above sea level. The unobstructed view towards the main wind direction west/north-west was the reason for installing the lidar on the 91 m tower (Figure 4.3).

Together with the lidar, a webcam was installed on the platform in order to take pictures of the view in westerly direction. The 102 m met mast was located in 1350 m distance westerly from the tower at 652 m above sea level and is fully equipped with meteorological sensors. Adjacent to the tower a reference turbine is located for which the minute-scale power forecasts are carried out. The hub height of the turbine is at the same level with the lidar mounting. Due to a confidentiality agreement with the turbine owner, data of this turbine will be displayed normalized by relevant parameters such as rated power.

The met mast was operated by University of Stuttgart as part of the project LidarComplex to record the environmental conditions at the site. It is equipped with sensors at 5 m, 50 m, 75 m, and 98 m that record high resolution meteorological data. Wind speed and wind direction data from the mast are used in this thesis to verify the reconstructed wind speed from the lidar data (cf. Section 4.2.2). Other meteorological data from the mast such as temperature, relative humidity, and precipitation are used to assess environmental conditions at the site. This will



Figure 4.6: StreamLine XR lidar on the nacelle of the AV04 turbine in alpha ventus.

be relevant for the assessment of the analysis of the lidar measurement range and the possible forecast horizon (cf. Section 7.1). Details of the met mast's sensor equipment are given in the appendix. The met mast was dismantled in August 2016, which means only a few months of concurrent lidar and met mast data were available.

In alpha ventus, the same lidar system (StreamLine XR) was installed on the nacelle of the AV04 turbine (marked in Figure 4.4c). It was installed behind the rotor in a corner of the service platform of the turbine (Figure 4.6). The lidar was raised up on a 2.5 m high frame, to measure above the railing of the service platform and to have more clearance before the lidar beam hits the nacelle when measuring in a vertical pattern, e.g. with an Range Height Indicator (RHI) scan. It was decided to mount the lidar on top of a turbine - and not for example on the transition piece - as the lidar then rotates with the yaw angle of the turbine and faces the inflow direction. Another benefit of installing the lidar on the nacelle is that a horizontal measurement is automatically at the hub height and therefore can be directly used to assess the power production of the turbine. The turbine AV04 was chosen, as it is a turbine at the outside of the wind farm layout facing the main wind direction, and located directly opposite of the met tower FINO1. This configuration ensured that the lidar was able to measure directly the inflow of the wind farm and at the same time this inflow could be further characterised using the meteorological data from the met mast.

FINO1 is a platform equipped with a met mast that measures wind speed, wind direction, air temperature, air pressure, precipitation and relative humidity at several heights up to a

100 m [32]. Wind speed and wind direction are measured with cup anemometers and sonic anemometers. The wind speed data from the cup measurements are available with a mast correction, which account for lateral speed-up effects, upwind flow retardation and downwind wake effects due to the mast construction itself [33]. A table with all sensors and their respective measurement heights can be found in Table A.2 in the Appendix.

The minute-scale forecasts at alpha ventus in this thesis are performed for the turbine AV04. Sensitive data of the turbine, such as the power curve are displayed normalized for confidentiality reasons.

4.1.5 Infrastructure planning

Two factors are important to consider when measuring at remote sites: the power supply of the measurement device and the remote access to the device to change the settings and transfer data.

During the campaign at Stötten the lidar was plugged into the power supply of the radio tower and remote data access was ensured via a modem. Measurement data from the lidar and met mast was automatically downloaded every night. New trajectories to test different configurations and methods for the forecasting were set via the remote access. The met mast was plugged into a nearby wind turbine and had access to the turbine's internet connection. The data from the reference turbine was supplied by the owner.

At alpha ventus, the lidar was connected to the turbine's power supply and the remote access to the device was established by connecting the lidar directly to the network of the wind farm. Data was downloaded directly to the institute server every night. Access to the turbine data was established via the data portal hosted by BSH. A data user agreement regulated the conditions between BSH and University of Stuttgart. It is important to note that the turbine data of the wind farms is not automatically uploaded to the data portal, but has to be handed over by the wind farm operator. This causes delays in access to the data of several months. The FINO1 data is available through another data portal hosted by the BSH, where just a registration but no user agreement is necessary.

4.1.6 Deployment and calibration procedures

After initial tests, the lidar was mounted for the first campaign on the top level platform of the radio tower in Stötten in October 2015 and measured there until August 2017. To find out the exact device alignment, a step-by-step procedure was established in the beginning of the campaign to ensure that the exact position of the laser beam is well known. The procedure is described in [16] in detail and was not part of this thesis. The wind speed calibration of this device was part of this thesis in so far, as they are needed to reconstruct the horizontal wind speed using measurements from the radio tower. To test these methods, the results were compared to the horizontal wind speeds measured at the met mast. The details are explained in Section 4.2.2.

For the second offshore campaign in alpha ventus, the lidar used in the VORKAST project was installed on the nacelle of the AV4 turbine in March 2019. Unfortunately, the device had a malfunction and did not measure more than around 1 km. Therefore it had to be decommissioned and sent to repair to the manufacturer. In the meantime an identical unit of a StreamLine XR lidar could be installed in the same spot and started measuring in October 2019. At the time of writing of this thesis, the measurement is still ongoing and planned until summer 2021. These data are available for the ParkCast project and this thesis.

4.1.7 Scanning modes design

The StreamLine XR lidar is able to measure radial wind speeds along the line-of-sight direction of the emitted laser beam. The beam can be steered to any direction in the hemisphere above the device and in an angle of 15° below. The steering angles are defined as elevation for vertical movements and azimuth for horizontal movements (cf. Figure 4.7). The measurement distance of the lidar is divided into range gates and as it is a pulsed device, the measurements are carried out simultaneously. The length of the range gates can vary from 18 to 60 m. A radial velocity is measured for each of the range gates. A sequence of emitted beams along a predefined azimuth and elevation angle is called a scan. For more technical device parameter, see Table A.1 in the appendix.

The scanning modes carried out during the first onshore campaign had the goal to measure the horizontal wind speed component which defines a wind turbine's power conversion and not the vertical component. Therefore only horizontal scans, with varying azimuth and zero degree elevation angle were carried out. The so-called Plan Position Indicator (PPI) scans were directed westerly into the main wind direction. An overview of the different scans is given in Table 4.1. Each scan in the table is tagged with a scan ID and the time period, the number




Figure 4.7: Schematic drawing of the lidar measurement setup.

of rays per horizontal scan, the number of range gates for each ray and its number of pulses, the azimuth angle of the scan and the time it takes to perform the scan is given. In the first few months of the campaign, only the number of pulses was varied, to test the influence on the measurement range (Section 7.1). It should be noted that along with the number of pulses, also the scan time varied, as these two parameters are directly linked. Later, also the range of azimuth angle was broadened. All scans were carried out in a step-stare mode, which means that the scan motors stop, and only move on after the measurement is carried out.

For the offshore campaign a similar approach for the scan modes was chosen. The goal was to measure the horizontal wind component, therefore horizontal PPI scans were carried out. The range of the azimuth angle was set broad, in order to capture the inflow of the whole wind farm. The number of pulses was only changed once. An overview of the different scans is given in Table 4.2.

4.1.8 Execution and data collection

During the onshore campaign, the lidar measured reliably on the radio tower platform from October 2015 to August 2017. Data from lidar and webcam was downloaded every few days via the modem connection or on regular maintenance checks on site. New scan modes were tested every few weeks. Data from the met mast was collected automatically every night.

During the offshore campaign, the lidar measured from March 2019 to August 2019 with a malfunction and was decommissioned. The replacement lidar was installed in October 2019 and measured reliably from then on. The data was collected automatically every night. Only in June 2020, the lidar was shut down due to a wind farm shut down and could only be started

ID	Period	No. rays	No. range gates	No. pulses	Azimuth angle	Scan time
1	28.04.2016 00:00 26.05.2016 23:40	41	167	10,000	250°– 290°	1'12"
2	07.06.2016 00:00 02.09.2016 12:20	41	167	30,000	250° – 290°	2'30"
3	02.09.2016 14:00 26.09.2016 08:20	41	167	40,000	250° – 290°	3'10"
4	26.09.2016 08:30 05.10.2016 08:20	41	167	60,000	250° – 290°	4'27"
5	05.10.2016 08:10 03.11.2016 10:00	41	111	60,000	250° – 290°	4'27"
6	19.01.2017 10:27 26.04.2017 14:10	16	111	60,000	252° – 282°	1'42"
7	26.04.2017 14:20 01.08.2017 00:00	11	111	45,000	282°– 302°	0'52"

Table 4.1: Overview of lidar scans at the onshore campaign in Stötten. Azimuth angle given in geographical coordinate system. Maximum measurement range always set to 10 km.

ID	Period	No. rays	No. range gates	No. pulses	Azimuth angle	Scan time
1	14.01.2020 00:00 27.02.2020 09:30	91	7980	20,000	240° – 60°	4'21"
2	11.01.2021 15:42 11.01.2021 15:30	71	400	40,000	260° – 40°	2'47"

Table 4.2: Overview of lidar scans at the offshore campaign in alpha ventus. Azimuth angle given in geographical coordinate system. In the first scan, gate overlapping was turned on. Maximum measurement range always set to 12 km.

again a month later. This was the only gap in the measurements.

4.1.9 Decommissioning and post-calibration procedures

For the onshore campaign, the lidar was removed from the radio tower in August 2017 as the VORKAST project is finalized. Post-calibration procedures were not performed, as no met mast was available at that time. The met mast was decommissioned a year before in August 2016, as the project Lidar Complex finished.

The offshore campaign is still ongoing at the time of this writing. It is planned to send the lidar for maintenance to the manufacturer, after decommissioning at alpha ventus.

4.1.10 Data availability

An overview of available data used from the onshore and offshore campaigns is given in Table 4.3 and Table 4.4 respectively. Data is available from different sources: lidar, webcam, met mast, and turbine.

The lidar data is available for both campaigns with a time resolution that depends on the scan configuration. For the onshore campaign, a webcam was available which recorded a picture every minute.

For the onshore campaign, met mast data is available only until August 2016 but with a high time resolution which depends on the sensor. Offshore, FINO1 met mast data is only available as 10-minute mean values.

The turbine data at Stötten are available as 10-minute averages for a period from July 2016 to December 2017. Available data are the turbine power, nacelle anemometer wind speed and the corresponding time stamp. For the AV04 offshore turbine, besides the above mentioned data the yaw angle of the nacelle is also available. For the alpha ventus wind turbine, also only 10-minute averaged data is available.

Device	Recorded signal	Time resolution	Available period
Lidar	Radial wind speed	Scan dependent	Oct. 2015 – Aug. 2017
Webcam	Pictures in westerly direction	1-min	Apr. 2016 – Aug. 2017
Met mast	Wind speed	50 Hz	Oct. 2015 – Aug. 2016
	Wind direction	50 Hz	
	Temperature	1 Hz	
	Relative humidity	1 Hz	
Turbine	Precipitation	1 Hz	Oct. 2015 – Aug. 2016 /
	Power	10-min mean	
	Nacelle wind speed	10-min mean	

Table 4.3: Overview of available data for the onshore campaign. More details on the met mast instrumentation is given in Appendix A.2.

Device	Recorded signal	Time resolution	Available period
Lidar	Radial wind speed	Scan dependent	Jan. 2020 – Jan. 2021
Met mast	Wind speed	10-min mean	Jan. 2020 – Sep. 2020
	Wind direction	10-min mean	
	Temperature	10-min mean	
	Relative humidity	10-min mean	
Turbine	Precipitation	10-min mean	Jan. 2020 – Apr. 2021
	Power	10-min mean	
	Nacelle wind speed	10-min mean	
	Nacelle azimuth angle	10-min mean	

Table 4.4: Overview of available data for offshore measurement campaign. More details on the met mast instrumentation is given in Appendix A.2.

4.2 Getting a useful wind speed out of a lidar

Conventional ground-based, profiling lidar devices are prepared by the manufacturer and are ready for costumers to use. For applications such as wind resource assessment the relevant data are 10 minute averages of horizontal wind speed wind direction and turbulence intensity. In these cases, the lidar is treated as a black box and algorithms that are applied for the data processing are not modified.

However, the StreamLine XR lidar that is used in this thesis requires customized data processing before the data can be used for power forecasting.

4.2.1 Data filtering

The lidar measurement principle is based on the reflection of laser pulses on particles in the air which backscatter the light with a frequency shift due to the airspeed of the particles (cf. Section 2.3). The measurement depends on the existence of these aerosols. If the concentration in the air is too high or too low, or the laser energy of the scan is set too low, the device may measure an incorrect signal. Especially in far range gates, the backscattered signal intensity is often low. As a consequence lidar data needs to be filtered and the measurement range may deviate from the maximum range given in the product data sheet.

When designing a filter for lidar data, filter requirements should be defined first, as the requirements vary from application to application. In this thesis the following requirements apply:

1. Conservative filtering with least possible data loss
2. Adaptability to varying environmental conditions
3. Highly efficient processing for real time capability.

The requirements are specified for the application of minute-scale forecasting. Therefore as many corrupted data as possible are to be filtered out (conservative filtering), but at the same time the least possible amount of data should be lost in order to reach the maximum measurement distance. The filter should also be robust and work in varying environmental conditions, and in order to be able to use it for real time application of the forecast, the processing speed needs to be accordingly high.

Carrier-to-Noise Ratio (CNR) filter

The standard approach for detecting outliers and noise in the lidar raw data is to use the CNR, which is an indicator for the signal quality. The CNR is an output signal of the lidar device. It will be used in this thesis with its normalized unit in decibel [dB]. The CNR is saved for to every radial velocity measurement for every range gate. Very high CNR is the result of the beam hitting a hard target. A low CNR value can be the result of either low aerosol concentration at a certain gate and therefore not enough backscattered signal or the result of too low signal intensity reaching the gate.

A typical timeline of unfiltered measured radial velocity VEL over the range of a typical measurement period is plotted in Figure 4.8a . The measured wind speed data in the first kilometers is coherent. After a transition range where outliers enter the homogeneous data, it becomes very noisy in the far measurement ranges. In the 30-hour time series that is shown, the sign of the wind speed changes from negative (wind speed towards LOS) to positive (wind speed away from LOS). From this, a change in wind direction from westerly to easterly can be deduced.

The radial wind speed is plotted over the corresponding CNR sorted by the measurement range in Figure 4.9. The data in the first 1000 m start off as a cloud in a CNR range from -12 to 0 dB. This cloud then moves gradually towards lower CNR levels with increasing measurement distance. Between 3000 to 4000 m range the first wind speed outliers occur. These outliers increase with increasing distance and spread from -20 to 20 m s^{-1} until the former data cloud vanishes in the outlier spread. There is a CNR threshold at around -22 dB below which wind speed outliers occur.

In CNR filter algorithms this threshold is used as a filter parameter. One can be sure that radial wind speed measurements that are tagged with a CNR above this threshold, are valid measurements. The only exception are data with CNR values above the threshold and a wind speed around 0 m s^{-1} . These data can result from hard targets, i.e. solid objects. A standard approach to filter lidar data using the CNR values is to set the device specific CNR threshold and mark all wind speed data below the threshold as invalid. This approach is used for example for applications such as lidar-assisted control for wind turbines where short-range wind lidars are used, and the use of any incorrect lidar data must be avoided [20]. However, the disadvantage is that potentially valid wind speed data with CNR below the threshold (and therefore marked as invalid) are lost and the measurement range therefore is cut short unnecessarily (Figure 4.8b). The challenge for the application of minute-scale forecasting is therefore to implement a filter algorithm that removes only the incorrect wind speed data, and includes valid wind speed measurements although they might have low CNR.




Figure 4.8: Typical time series of LOS lidar data unfiltered and filtered.




Figure 4.9: Radial wind speed over CNR sorted by the measurement range.

Edge detection filter

The solution to filter lidar wind speed data comes from the area of image processing.¹ When processing any image, chart or photo, often there is an edge detection algorithm involved. An edge within an image is classified as a significant local change in the image intensity, which typically occurs right on the boundary of two adjacent areas within the image [34]. For this reason “edge detection is frequently the first step in recovering information from images” [34].

Figure 4.10 shows a photograph of the Stuttgart TV tower on the left and the result of the edge detection filter on the right. Local changes in the color data are detected with the filter. Transferred to the application of wind speed filtering, an edge detection algorithm detects local changes in the wind speed by calculating the difference of maximum and minimum of radial wind speed ΔVEL within a predefined window. A window size of [1 3] for the edge filter means that three velocity values are included in the window from the same range gate of three neighbouring beams. The window then moves over all range gates and beams. Thus, a matrix with differences for each measurement point is stored. Afterwards it is checked whether the differences ΔVEL exceed a predefined ΔVEL threshold. If this is the case for both adjacent values of a measurement point, the corresponding measurement is marked as invalid. This results in a logical matrix which contains an entry for valid or invalid measurement points for all range gates and beams. It is important to note that by checking ΔVEL of adjacent measurement points, only the outliers are marked as invalid, and not necessarily the neighbours.

To demonstrate the functionality of the filter, an example of a generic wind speed time series

¹This section is based on the Bachelor thesis "Adaptive filtering of long range lidar data" carried out by Malte Justus Niemeier, which was handed in at SWE in 2016 and supervised by the author of this PhD thesis and her colleague Maayen Wigger.



Figure 4.10: Image of the Stuttgart TV tower (left) and edge detection filtered (right). Left image reproduced with modifications from [35] with permission.

is given in Figure 4.11. The Figure also gives the wind speed difference ΔVEL for a window size of [1 3]. For the first and last beam in the scan, the left and right values next to the center point of the moving window are ignored. For a ΔVEL threshold of $>2\text{ m s}^{-1}$, the measurement points from beam 1, 5 and 6 would be marked as invalid in this example. For a wind threshold of $>3\text{ m s}^{-1}$ only beam 1 would be excluded.

The filter parameters that are relevant for the performance of the edge filter are the window size and the wind speed limit ΔVEL . To test filter requirement 1, a parameter study is carried out, to find out how many data remain after filtering with the edge filter for different window sizes and wind speed limits (Figure 4.12). This study was carried out for the time series in Figure 4.8a. The results show that for a window size of 3 adjacent data points (smallest possible window), the least amount of data are filtered. By increasing the window size to 5, around 10% more data are filtered. Increasing the window more, leads to more data loss. The wind speed limit ΔVEL also has a significant influence on the data availability after filtering. Setting the limit to 0.5 m s^{-1} , thus allowing only very small wind speed fluctuations, leads to a rigorous filtering of the data. When increasing the limit gradually, the amount of available data increases rapidly at first, and then evens out.

Figure 4.8c shows the result of the edge filter for a window size [1 3] and a wind speed limit ΔVEL of $>3\text{ m s}^{-1}$. Compared to the CNR filter (26% data availability), much less data is discarded with the edge filter (65% data availability). Hence, the measurement range is increased significantly. Wind speed values with a CNR below the threshold of -22 dB are then

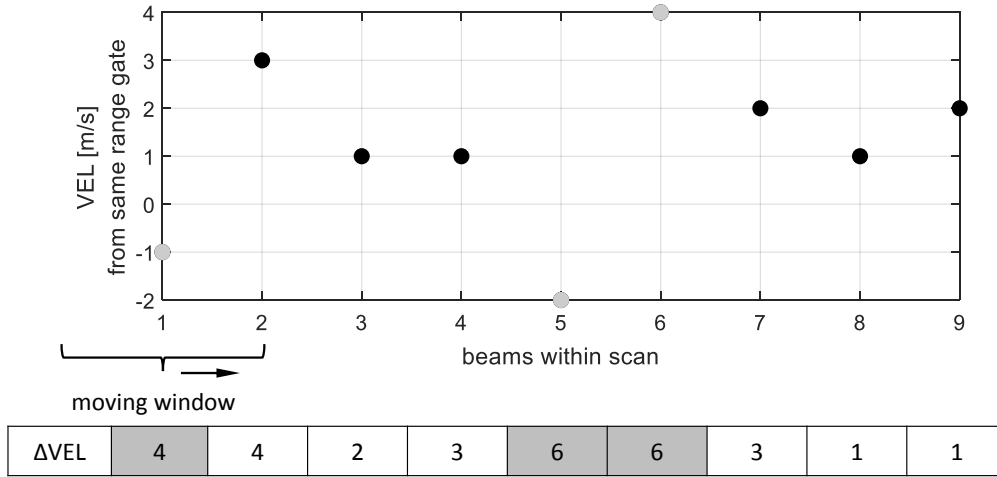


Figure 4.11: Demonstration of edge filter for a generic wind speed time series for window size [1 3] and Δ VEL threshold of $> 2 \text{ ms}^{-1}$. Filtered values marked grey.

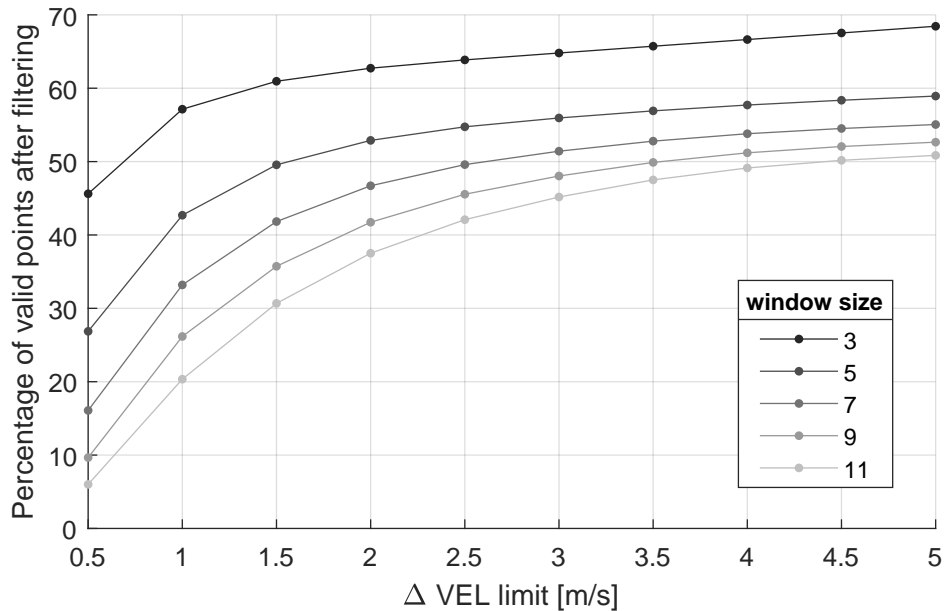


Figure 4.12: Percentage of good points after filtering with edge filter with different Δ VEL limits (thresholds) and window sizes. Study carried out for time series in Figure 4.8




Figure 4.13: Radial velocity over CNR filtered with the edge filter.

“rescued” when applying the edge detection filter with those filter parameters (Figure 4.13). Although some outliers remain in the data, it is decided to apply the edge detection filter with a window size $[1 \ 3]$ and a wind speed limit ΔVEL of $>3 \text{ m s}^{-1}$, for the sake of the increased measurement range for both the onshore and offshore data.

The edge filter can adapt to varying environmental conditions such as the wind direction change, and is very simple so that it is computationally highly efficient. Thus all the filter criteria are fulfilled.

4.2.2 Wind field reconstruction

A lidar measures the radial component VEL of the wind vector in the laser beam direction LOS, in contrast to classical wind measuring systems such as a cup anemometer that measure horizontal wind speed. The challenge is to derive the horizontal wind velocity from the measured radial wind velocities of the discrete measuring points of the lidar scan; this is known as wind field reconstruction.

The basic equation to reconstruct the three wind components u, v, w from the radial component is

$$VEL = \frac{x_i}{d_i}u_i + \frac{y_i}{d_i}v_i + \frac{z_i}{d_i}w_i. \quad (4.1)$$

where the lidar measures at the coordinate point $[x_i \ y_i \ z_i]$ in a distance d_i the unknown wind vector $[u_i \ v_i \ w_i]$ [5]. At least three measurement points are necessary to solve this equation.

Therefore assumptions must be made. For example assuming a constant and homogeneous wind field during the measurement, the previous equation can be described in matrix form for all measured radial wind velocities as

$$\underbrace{\begin{bmatrix} VEL_1 \\ \vdots \\ VEL_n \end{bmatrix}}_r = \underbrace{\begin{bmatrix} \frac{x_1}{d_1} & \frac{y_1}{d_1} & \frac{z_1}{d_1} \\ \vdots & \vdots & \vdots \\ \frac{x_n}{d_n} & \frac{y_n}{d_n} & \frac{z_n}{d_n} \end{bmatrix}}_A \cdot \underbrace{\begin{bmatrix} u \\ v \\ w \end{bmatrix}}_s. \quad (4.2)$$

The wind field vectors u , v and w are assumed to be constant for the time it takes to measure points $1 \dots n$. The matrix equation can then be simplified in abbreviated form to

$$r = A s. \quad (4.3)$$

If sufficient linearly independent measuring points are available, the wind speed components u , v and w can be estimated using the inverse A^{-1} if A is square, or the least-square Penrose pseudo inverse A^+ when there are more than three measurements available

$$\begin{bmatrix} u \\ v \\ w \end{bmatrix} = A^{-1} \begin{bmatrix} VEL_1 \\ \vdots \\ VEL_n \end{bmatrix}. \quad (4.4)$$

Onshore campaign

In this thesis two different approaches for wind field reconstruction had to be applied for the onshore and offshore campaign. For the onshore campaign, the lidar was installed fixed on the radio tower and the scan direction was therefore fixed. There was no information about the wind field available from an external source. Consequently wind direction and horizontal wind speeds used for the minute-scale forecasting had to be derived solely from the radial wind speed itself using Equation 4.4.

The question that Equation 4.4 presents is how many data points are necessary, in order to estimate reliably a wind vector at a specific location.

To answer the question, two methods are developed and investigated in this thesis for the onshore campaign: i. the global-local reconstruction and ii. the moving window reconstruction. Both approaches are applied scanwise, thus only using data from one horizontal sweep at a time. One scan consists of 41 beams with 167 measurement distances (Table 4.1). Thus a total number 6,847 (41x167) measurement points are available to solve the linear equation system

and it needs to be investigated for which points the solution of the equation provides the best estimates of the wind speed vector. Data from the local met mast serves as reference for results of the wind field reconstruction methods. The mast was located westerly of the lidar and records both wind speed and wind direction (cf. Section 4.1). The mast data is the basis for the evaluation of the wind field reconstruction methods. The idea is to compare data from the mast and the reconstructed lidar data measured closest to the mast location. It needs to be noted that the highest sonic measurements on the met mast are still 73 m below the lidar scan.

To ensure a correct reconstruction, it is important to know in which coordinate system the lidar data is analyzed. Four coordinate systems are defined, which are shown in Figure 4.14.

1. The lidar coordinate system (L) in which the raw data are recorded. This depends on the alignment of the lidar device and the system specifications. The StreamLine XR measures in a left hand coordinate system.
2. The geographic coordinate system in which the wind direction is specified. The wind direction is defined by the direction the wind is coming from.
3. The inertial coordinate system (I) is a fixed right-hand system with arbitrary orientation. The origin is set at the location of the lidar measurement and the x -axis points west in the main wind direction.
4. The wind coordinate system (W), which rotates with the reconstructed global wind direction and in which the lidar raw data are projected onto this global wind direction vector. Thus, the reconstructed local wind velocity components u and v per measuring point are obtained.

Global-local reconstruction method

The first wind field reconstruction approach developed and tested is the *global-local method*. In a first step a global wind field reconstruction per lidar scan is performed using Equation 4.4, to determine the horizontal wind vector components u , v over the entire measuring range. It is assumed that the vertical component w of the wind field can be neglected. From the reconstructed velocity components u , v a global wind direction is derived. In a second step, the local horizontal wind speeds are determined by projecting the radial wind speeds VEL_i on the global wind direction, assuming that the wind direction does not change for the duration of a scan.

Before discussing the results, first a list of lessons learned from the implementation of the global-local reconstruction method is given.

- The precise wind direction reconstruction is critical for a correct wind speed reconstruc-

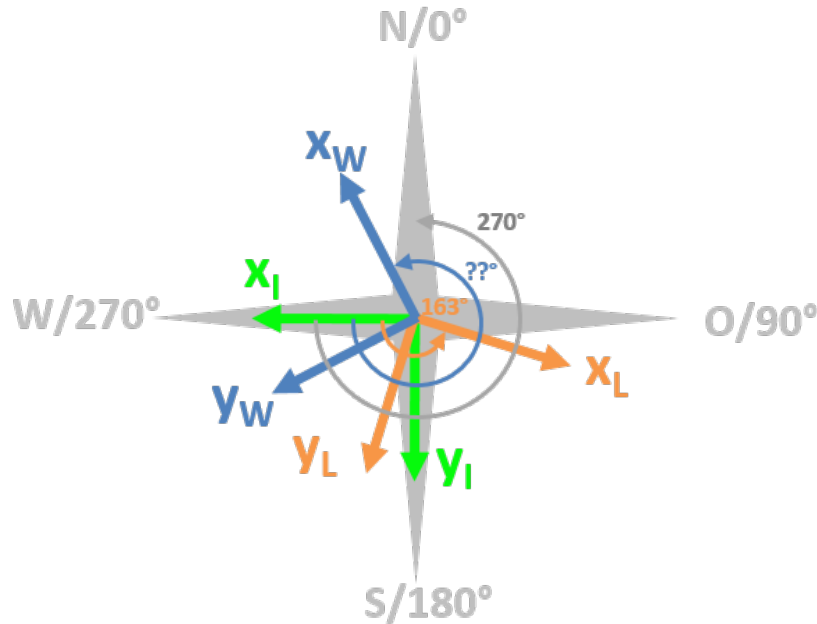


Figure 4.14: Overview of the coordinate systems that are necessary for the wind field reconstruction. The inertial coordinate system x_I, y_I (green) is fixed; the lidar coordinate system x_L, y_L (orange) is fixed and depends on the alignment of the lidar; the wind coordinate system rotates x_W, y_W (blue) along with the wind direction.

tion.

- In order to reconstruct the wind direction correctly, the exact position of the lidar coordinate system needs to be known.
- Scans with poor data availability must be neglected, as it can lead to wrong reconstruction of the wind direction if too little data are available to solve the equation system. Therefore a filter was installed to check the data availability per scan and to reject scans with an availability less than 15%.
- The reconstruction of the local wind speed fails when the wind direction is 90° offset to the lidar scanning direction. In this case, the lidar device can only measure very small values of the radial wind speed, and large errors occur during reconstruction, since the sensitivity for reconstruction errors is high. Therefore a filter was implemented which detects areas in which the scanning direction is in a range of $90^\circ \pm 15^\circ$ of the wind direction and filters out the erroneous data of the reconstructed wind speed.

After the implementation, the reconstruction algorithm is applied to a test data set from July 2016, where data from lidar and met mast are available for validation. The 10-minute averaged reconstructed global wind direction is compared with the wind direction measurement at the mast in the time domain (Figure 4.15 top) and as regression (Figure 4.16a). The




Figure 4.15: Time series of wind direction (top) and wind speed (bottom) of sonic (black) and lidar measurements (grey) calculated using the global-local reconstruction method. Data averaged over 10 minutes.

local reconstructed horizontal wind speed closest to the mast is retrieved, averaged over 10 minutes and compared to the horizontal wind speed of the mast in the time domain (Figure 4.15 bottom) and as regression plot (Figure 4.16b). In order to avoid the corruption of the calculated regression parameter due to the north jump of the wind direction around 360° and 0° and to give credit to its circular nature, the sonic wind directions are shifted 360° if lidar and sonic data are on the opposite sides of the north jump respectively.

The comparison in the time domain shows that lidar and sonic wind directions are well aligned. Changes in wind direction are captured when reconstructing the global wind direction using the data from one scan. The regression plot shows however that outliers reduce the regression parameter R^2 significantly. Overall, the assumption of a global wind direction which does not change during the time it takes to carry out one scan and which is valid over the whole measurement range is reasonable. The comparison of lidar and sonic wind speed shows that the lidar data follows the trend of the sonic measurements. However, the regression plot reveals a high scatter in the data. The positive offset and Root Mean Square Error (RMSE) can be explained with the distance in measurement height between sonic and lidar. Overall, the projection of VEL on the global wind direction to obtain a local wind speed is a reasonable reconstruction method but leaves room for improvement.




Figure 4.16: Regression of wind data between sonic and lidar measurements calculated using the global-local reconstruction method. A linear fit is also shown (grey). Data averaged over 10 minutes.

Moving window method

The second wind field reconstruction approach developed and tested is the *moving window method*. Assuming that the vertical component w of the wind field can be neglected, the local horizontal wind vector components u , v are estimated using Equation 4.4 with *VEL* data from a moving window. The method uses a window which is centered on each measurement point of a scan and moves through the data (Figure 4.17). Thus for each center point of the window the local horizontal wind vector and wind direction is estimated. The question that arises is which data points should be used for the reconstruction, i.e. the size of the window. Therefore different window sizes are tested. To evaluate the results, the same test data set as for the global-local reconstruction is used and the reconstructed horizontal wind speed at the measurement location closest to the mast is compared to the sonic data.

The linear regression of the 10-minute averaged data of wind direction and wind speed (Figure 4.18 shows that the window size has a significant influence on the quality of the reconstruction. A reconstruction in a 3×3 window using 9 radial velocity measurements, results in a big scatter in wind direction ($R^2 = 0.71$), and reconstructed wind speeds that are on average twice as high as the sonic wind speeds ($m = 2.03$). Outliers with wind speed over 60 m s^{-1} can occur as well. The more data is used for the reconstruction, the better the regression becomes. A saturation seems to be reached however with a window size of 19×19 . Increasing the window size further

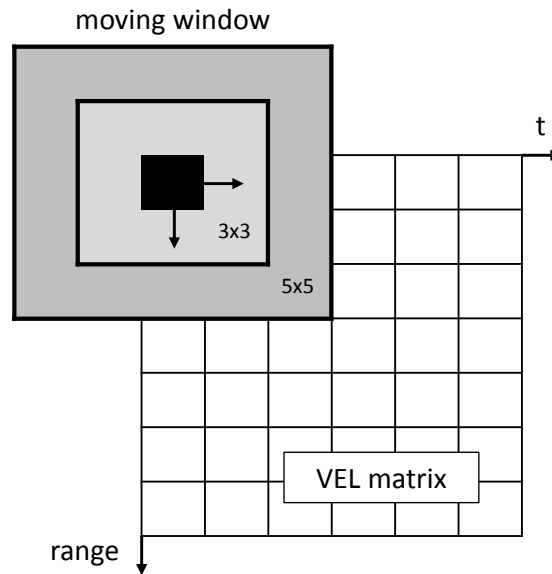


Figure 4.17: Sketch of the local moving window reconstruction method with two exemplary window sizes.

to the maximum of 41x41 measurement points, does not change the regression significantly.

Figure 4.19 gives an overview of the regression parameter for all tested window sizes ranging from 3x3 to 41x41. Only odd numbers and square windows are tested because of the need to center the moving window. The maximum window size tested corresponds with the maximum number of beam per scan. Each point in the figure is derived from a linear regression between 10 minute averaged lidar and sonic data of wind direction and horizontal wind speed respectively. Generally the wind direction reconstruction works better than the wind speed reconstruction (R^2 is higher and slope m is closer to 1). For small window sizes, the reconstruction is inaccurate (low R^2 , slope well over 1 and a significant offset). When increasing the window size, that is taking more data into account for the reconstruction, the reconstruction improves considerably. For window sizes > 25 the reconstruction improves only slightly.

Comparison of global-local and moving window reconstruction method

The two different reconstruction methods have been introduced and tested. The question that arises, is which method is more appropriate for minute-scale forecasting. To compare the global-local and moving window reconstruction method, Figure 4.19 also shows the regression results of the global-local reconstruction. The comparison shows an improvement of the reconstruction using the moving window method for window sizes > 19 (R^2 is higher, offset b and RMSE lower).




Figure 4.18: Regression of wind direction (left) and wind speed (right) between sonic and lidar measurements calculated using the moving window reconstruction method with different window sizes. A linear fit is also shown (grey). Data averaged over 10 minutes.




Figure 4.19: Regression parameters of 10 minute averaged wind direction (black) and wind speed data (grey) between lidar and sonic measurements using the moving window reconstruction method with different window sizes. Results of the global-local reconstruction are given for comparison.

The analysis so far was carried out with 10-minute averaged reconstructed wind speeds from the location closest to the met mast. However, for the application of minute-scale forecasting high resolution measurements of the horizontal wind speed are necessary to forecast the power output of wind turbines in a minute-scale time resolution. Therefore, as a last step to evaluate the reconstruction methods, the reconstructed wind speed is analysed scanwise without averaging. Figure 4.20 gives an overview of wind direction (left) and wind speed (right) for the global-local method and for three different window sizes of the moving window method. For comparison the measured radial wind speed is given in the top row. The plotted area is limited to a measurement distance of 3 km as there are data gaps in further distances due to filtered values.

The particular scan is chosen as the radial wind speed shows an inhomogeneous wind field and is a typical example for the complex measurement site. The wind direction for this scan is 270° , thus the scan direction and wind direction are aligned. Therefore the magnitude of the measured radial wind speed is directly comparable to the magnitude of the reconstructed horizontal wind speed. The result of the global-local reconstruction method reflects the inhomogeneity of the wind field accurately in the reconstructed horizontal wind speed. As expected, the reconstructed wind direction of the global-local method is constant over the scan as this assumption is the basis of the method.

The results for the moving window reconstruction reflect the results of the previous statistical analysis with 10-minute averaged data. For small window sizes, the reconstruction fails at many measurement locations and leads to high horizontal wind speed outliers. The reconstructed wind direction in this case is inhomogeneous. With growing window sizes, the wind speed outliers become less, and wind speed and wind direction of the scan become more homogeneous. As the analysis of the 10-minute averaged data has suggested, the reconstruction with window sizes bigger than 19×19 does not change the result. The disadvantage of the moving window method becomes obvious: the window that runs through the data and takes all data within the window into account to solve Equation 4.4 acts as a filter. It smoothes the wind speed fluctuation within the window as it forces one valid result of the equation for all data within the window. The inhomogeneity of the wind speed is lost.

At this point, the expectations have not been fulfilled. The idea to localize the wind field reconstruction in order to account for to the inhomogeneity of the wind field did not work in practice. Smoothing the wind speed fluctuations by forcing one result within a certain window is counterproductive. For the most accurate minute-scale forecast, the fluctuations of the wind speed are assumed to be crucial, in order to forecast correctly. It is therefore decided, to apply the global-local reconstruction method for the following analysis in this thesis.




Figure 4.20: Scan plots of wind direction (left) and wind speed (right) from different reconstruction methods. 1st row: radial wind speed for comparison, 2nd row: global-local method, 3rd-5th row: wind method with different window sizes.

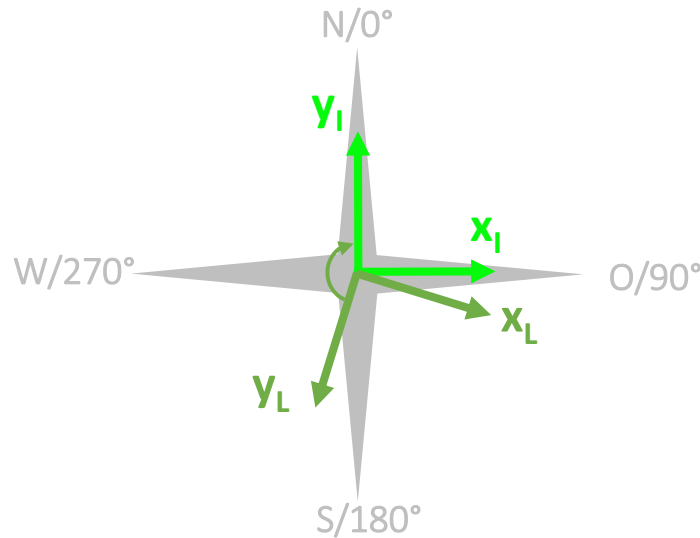


Figure 4.21: Overview of the coordinate systems used in alpha ventus. The inertial coordinate system x_I, y_I (light green) is fixed, the lidar coordinate system x_L, y_L (dark green) rotates along with the nacelle of the turbine.

Offshore campaign

For the offshore campaign, a different, simplified wind field reconstruction approach is applied in order to obtain the horizontal wind speed and wind direction. This approach has been introduced in previous studies for nacelle-based lidar systems [36]. The lidar on alpha ventus is installed on top of the nacelle of the wind turbine AV04. The difference compared to the onshore campaign is that the lidar rotates along with the wind direction, because the turbine's yaw control ensures, that the turbine is always facing into the wind. This means first that information about the turbine's yaw position is necessary, in order to know the lidar orientation. Secondly, the information of the yaw or wind direction can be used to reconstruct the horizontal wind speed. Before the reconstruction method can be applied, however, it is again important to define the coordinate systems (Figure 4.21). The lidar measures in the lidar coordinate system which rotates with the turbine yaw angle. In order to transfer the lidar measurement coordinates into a fixed inertial coordinate system, the yaw angle has to be known. With this knowledge, the reconstruction task can start.

However, in order to reconstruct the wind field and to solve the basic Equation 4.4, two assumptions are required:

- The vertical wind speed w is zero.
- The wind turbine is always perfectly aligned with the wind direction, therefore the lateral wind speed component v in the lidar coordinate system is zero.

With these assumptions, Equation 4.4 is transformed from an over-determined set of linear

equation to the following equation:

$$u_i = \frac{f_i}{x_i} V E L_i \quad (4.5)$$

That means for every lidar measurement point i , the radial velocity is projected on the perpendicular vector to the rotor, in order to obtain the longitudinal wind speed component u in the lidar coordinate system. If the wind direction is assumed to be the yaw angle of the turbine, then the u component in the lidar coordinate system can be transferred to the fixed inertial coordinate system in order to obtain the horizontal wind speed components u, v in the inertial frame of reference.

The assumption that there is no yaw misalignment and that only the longitudinal wind speed component u is reconstructed makes sense for the minute-scale forecasting. For the conversion to power, only this longitudinal component is relevant. Yaw misalignment leads to a decrease in power output, but since only the relevant wind speed component u is considered with this approach, the decrease in wind speed and therefore power is considered automatically.

Figure 4.22 shows the reconstructed horizontal wind speed in alpha ventus from two simultaneous scans that measured the inflow and wake of the wind farm at the same time. The scan was chosen because it shows the typical measurement conditions at the site. The wind direction is south-west and the measurement range is around 5 km.

The reconstructed wind field shows the typical inhomogeneity in the inflow of the alpha ventus which result from the wakes of surrounding wind farms. The wind field also shows that the reconstructed wind speed has very large fluctuations, when the measurement is orthogonal to the wind direction. This phenomenon was already observed in the global-local reconstruction method, and results from the fact that the measured radial wind speed is very small in these conditions and therefore the sensitivity to reconstruction errors is very high.

As with the Stötten data, a filter was implemented which detects areas in which the scanning direction is in a range of $90^\circ \pm 15^\circ$ of the wind direction and filters out the erroneous data of the reconstructed wind speed.

Another filter had to be implemented for the turbine's yaw angle data. A comparison of the yaw angle and the wind direction measured at FINO1 revealed deviations for wind directions around 0 deg and 360 deg (Figure 4.23). It was concluded that the averaging algorithm that calculated 10-minute averages from high resolution azimuth data did not take into account the correct calculation when the wind direction crosses the north during the averaging period. These data were therefore filtered, as the correct yaw angle is crucial when using it to transfer the lidar measurement coordinates into a fixed inertial coordinate system.




Figure 4.22: Reconstructed horizontal wind speed in alpha ventus from two scans from 15.01.2020 20:23. Turbines are marked with black dots.

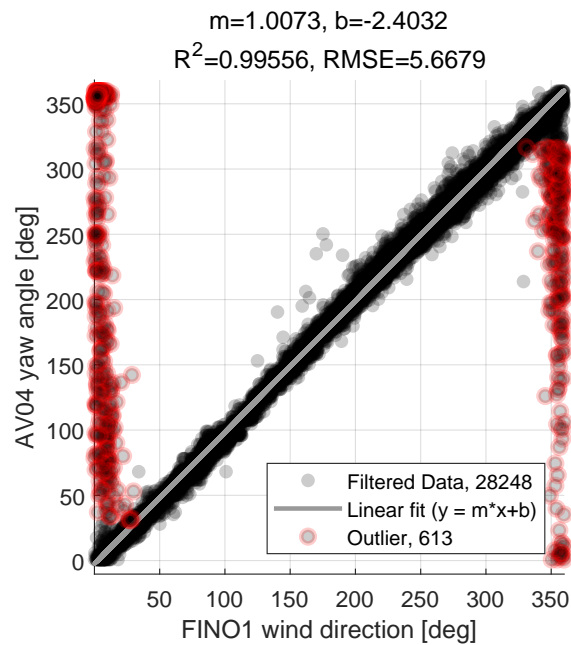


Figure 4.23: Correlation of the azimuth angle of the AV04 turbine and the wind direction measured at FINO1, including outliers (red). The correlation coefficients are calculated for the filtered data set.

4.3 Lessons learned


Lidar data needs to be filtered before they can be used for forecasting. Traditional filter techniques using CNR are too conservative for long-range measurements and remove valid data from farther measurement ranges. Using only radial wind speed measurements as filter input and searching for outliers within a size optimized data window leads to increased valid measurement ranges.

Radial wind speed measurements need to be reconstructed to obtain horizontal wind speed and wind direction information. The optimized reconstruction strategy for measurements that are fixed and do not rotate along with the turbine azimuth is to solve the reconstruction equation per scan. It is then assumed that the wind direction is constant per scan and the radial wind speed is project to the wind direction to gain the horizontal wind speed. Localized reconstruction within a size optimized data window leads to a smoothing of the natural wind speed fluctuations, which is not desired.

The optimized reconstruction strategy for measurements from the nacelle where the lidar rotates along with the wind direction, is to assume that the nacelle is aligned with the wind direction and then reconstruct only the longitudinal wind component in the lidar coordinate system. For minute scale forecasting of power this assumption is reasonable, as the longitudinal wind component is the most dominant component for the power conversion and therefore needs to be forecast.

Scan angles of 90° to the wind direction should be avoided because the measured radial velocity is very small and therefore the reconstruction leads to large fluctuations of the reconstructed values. Measuring in those areas costs unnecessary scan time and the data needs to be removed.

Along the forecasting chain: wind field evolution



In the previous chapter, it was shown how horizontal wind speed data can be obtained from lidar measurements. In order to forecast the power of a wind turbine, these distant measurements need to be propagated through time and space to the wind turbine's location. The goal of this chapter is to find a method for wind field evolution and to apply it to the data set, taking into account the local flow conditions of the boundary layer.

The simplest approach to propagating the wind field is to assume that the wind field measured at a distance simply propagates with its existing characteristics. This is known as Taylor's frozen turbulence hypothesis and is introduced in Section 5.1. Its implementation for minute-scale forecasting is discussed in Section 5.2.




Figure 5.1: Sketch of a turbulent eddy passing a wind turbine according to Taylor’s frozen turbulence hypothesis. If the length of the eddy is 100 m and the horizontal wind speed is 10 m s^{-1} , it takes 10 s to pass.

5.1 Using Taylor’s hypothesis in wind energy applications

In 1938 the British physicist and mathematician Geoffrey Ingram Taylor published the article “The spectrum of turbulence” in the Proceedings of the Royal Society [37]. In the article he postulated that for special conditions turbulent eddies in the atmosphere are transported with the mean flow and do not evolve with time but remain unchanged. This hypothesis became known as “Taylor’s frozen turbulence hypothesis”.

Eddies are turbulent structures in the atmosphere that are created by forces acting on the flow. In the lower boundary layer turbulence is mostly generated by forces from the ground [38]. For instance, buoyancy forces act during sunny days when air is heated by the ground. Surface roughness causes friction drag which leads to wind shear, which often generates turbulence. Obstacles such as trees, forests, or in the case of the measurement site in Stötten an escarpment, cause turbulent wakes downwind of the obstacles. The largest eddies in the boundary layer are roughly the same size as the depth of the boundary layer, i.e. 100 m to 3000 m.

Following Taylor, one can calculate the time it takes an eddy with the diameter λ to pass a sensor, or in our case a wind turbine, if the mean wind speed v_{hor} of the eddy is known (see Figure 5.1). The time period Δt it takes the eddy to pass is then given by

$$\Delta t = \lambda / v_{hor} \quad (5.1)$$

with v_{hor} as the horizontal wind speed.

In [38] Stull reminds the reader to remember that turbulence is not really frozen and that Taylor’s simplified hypothesis is only valid for eddies that evolve with a time scale larger than

it takes the eddy to pass the sensor. A recent study on the testing of Taylor's hypothesis for wind energy applications conducted by Schlipf et al. has shown that the hypothesis is valid for eddy sizes relevant for multi megawatt wind turbines [5]. In the study, a nacelle-based lidar was used to track the inflowing wind speed on the turbine and to calculate the coherence and phase between measurement points from different measurement distances. The result showed that for eddy length scales of around 250 m, Taylor's hypothesis was found valid with a 90 % accuracy.

In this thesis Taylor's hypothesis is applied to propagate the reconstructed wind field measured by the lidar through time and space to obtain the wind speed at the turbine. It is recognised that the conditions to apply Taylor's hypothesis are not ideal: the long-range lidar measures wind speed in several kilometers distance, over a large area. Moreover the onshore measurement site in Stötten is situated in complex terrain and it is expected that inhomogeneous wind conditions may not meet the assumption of Taylor's hypothesis of frozen turbulence. Nevertheless, the approach is attractive because of the low computational effort, which is crucial for minute-scale forecasting.

There are alternative approaches possible for modeling wind evolution in complex terrain. For example, the measured wind fields could be used as the initial conditions for a numerical weather prediction-based simulation [39] using local terrain and land cover data. However, this approach would require much increased computational time compared to Taylor's hypothesis. The extra steps required for processing would add time and increase the risk of failure of the forecast. For these reasons Taylor's frozen turbulence hypothesis was used as a pragmatic first engineering solution to investigate potential for lidar-based minute-scale forecasting.

5.2 Using Taylor's hypothesis for minute-scale forecasting

The goal of using a wind field evolution model in minute-scale forecasting is to find out which wind conditions reach the turbine at which point in time. Applying Taylor's hypothesis, this means that for each horizontally wind speed vector v_{hor} measured by the lidar, the distance to the turbine Δx is used to calculate the time period δt it takes the wind speed vector to reach the turbine (Equation 5.1).

What needs to be determined is which of the measured wind speed vectors actually reach the turbine.

Two factors are important to find the answer to this question:

1. The measured wind speed vector needs to be moving towards the turbine. Only if the actual inflow of the turbine is measured, the measurement is relevant for the forecasting.

This means the measurement direction and the wind direction need to be aligned.

2. The measured wind speed vector needs to go through the wind turbine's rotor diameter. Only wind speed vectors that are relevant for the turbine's power production are taken into account.

In the case of the offshore campaign, where the lidar was measuring on top of the nacelle of the turbine and therefore rotating along with the wind direction, the first condition is always met. Only for periods where the wind turbine is stopped, e.g. for maintenance, the nacelle and the lidar is not facing the incoming wind direction.

In case of the onshore campaign, where the lidar was fixed and did not rotate along with the wind direction, the measured flow field was not necessarily blowing towards the turbine. Only in the case of westerly winds, when the measurement trajectory was aligned into the wind direction, the measured wind field would actually arrive at the turbine.

To ensure that the wind vectors go through the turbine's rotor diameter, the turbine's rotor is implemented as a target: only if the measured wind speed vectors pass through the target, they are taken into account for the forecasting of the power production. The target is set as a horizontal line, with a length of two times the rotor diameter of the respective turbine, for both measurement campaigns. Two times the rotor diameter was found to be the ideal area of influence to optimize dual Doppler radar forecasts [10] and is therefore used here as well. The results can be applied to doppler lidar measurements, because the propagation principle is the same: a measured wind vector is advected to the turbine.

Taking the above criteria into account, the wind field evolution algorithm is implemented as follows:

- At the time a forecast is issued, lidar data from 30 minutes before the forecast is issued are taken into account to generate the forecast.
- Every measured wind speed vector is checked whether it actually reaches the target.
- For each vector that reaches the target, the time is calculated when the target is reached.
- If this time is before the forecast is issued, the vector is not used for the forecast.

Figure 5.2 visualizes the wind field evolution for an example period from the onshore campaign. The target wind turbine is shown in red, and the visualization starts when the lidar scan starts sweeping clockwise. For each following minute of the measurement a snapshot of the wind field is shown, where the measured vectors would move through time and space according to the Taylor hypothesis.

In this example, the vectors are moving towards the turbine, as the wind direction was aligned in the scan direction. On the second sweep (starting at 02:31), the wind direction

changed slightly and the blue vectors now come from a different direction. At around 2:32 the first wind speed vectors reach the turbine, and more vectors are passing through the target in the following minutes.

This is exactly the information that is needed for the forecasting: with the help of Taylor's frozen turbulence model, it is determined at which point in time the measured wind speed vectors reach the turbine after the forecast is issued. Only wind speed vectors that actually pass through the turbine are relevant for the turbine's power output. These wind speed vectors are taken into account for the forecast.

5.3 Lessons learned

To predict the wind speed at the turbine, the wind speed measured by the lidar needs to be propagated through space and time. Therefore a wind field evolution model is necessary.

Taylor's hypothesis can be used to calculate the time it takes the eddy (and its representative wind speed vector) to move towards a target, under the assumption that a turbulent eddy does not change but remains frozen.

Only wind speed vectors which actually reach the turbine should be taken into account for the wind speed and power prediction. Using Taylor, it can be checked whether the wind speed vectors pass the turbine's rotor.

Nacelle-mounted wind lidars have the advantage that the measurement is always aligned to the wind direction and therefore the measured wind speeds are always propagated towards the wind turbine. Hence, there are always wind speed vectors available for the prediction (if the lidar is measuring).

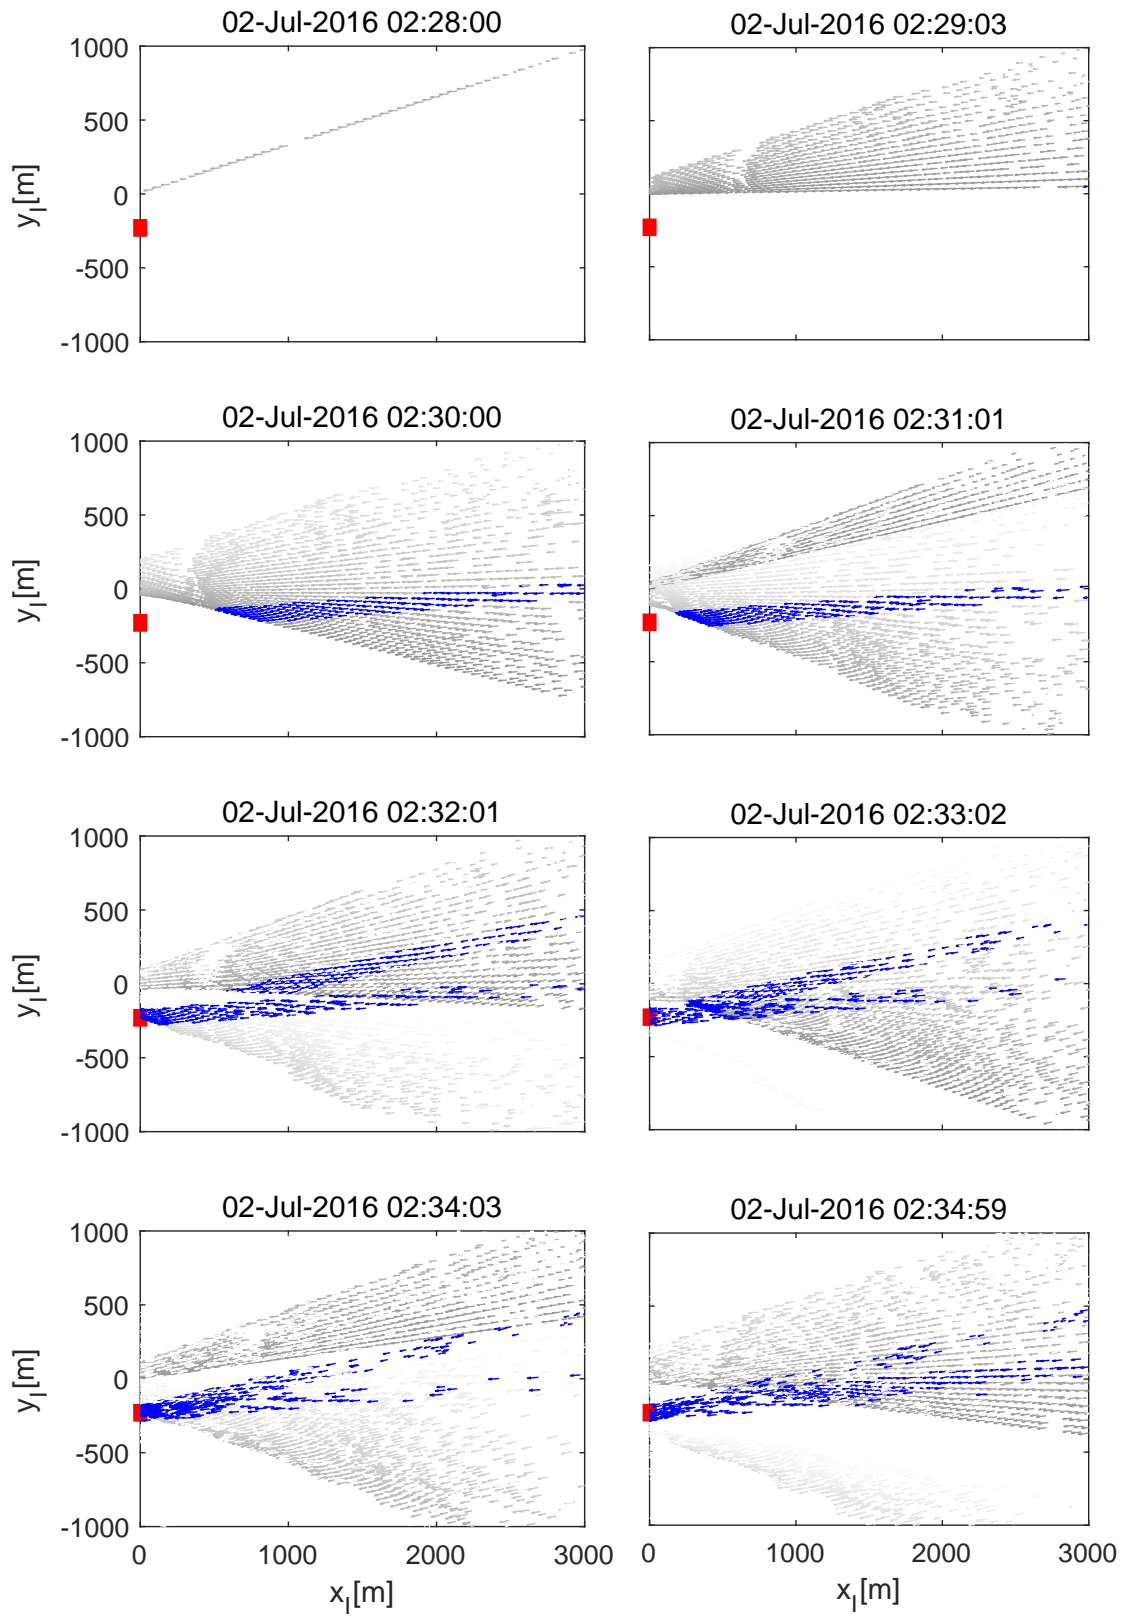



Figure 5.2: Scan plots of wind speed propagation from seven consecutive minutes of measurements. Wind direction is toward wind turbine (red). Wind speed vectors in grey miss the turbine (faded out with time), wind speed vectors in blue pass through turbine (do not fade out).

Reaching the end of the chain: minute-scale forecasts of wind speed and power



In the previous chapter it was discussed how useful wind field information can be gained from raw lidar measurements and how this wind field can be propagated through time and space in order to predict wind speed at the turbine.

Now the end of the chain is reached. In this chapter the last link of the forecasting chain is explained: how the predicted wind speeds are transformed into predicted power, and how the forecasts are then calculated. To this end, first the methodology of probabilistic minute-scale forecasting is introduced and explained in Section 6.1. Then the results from the onshore and offshore campaign are discussed in Section 6.2 and Section 6.3, respectively. Finally the impact of wind ramps on the minute-scale forecasts is assessed in Section 6.4.

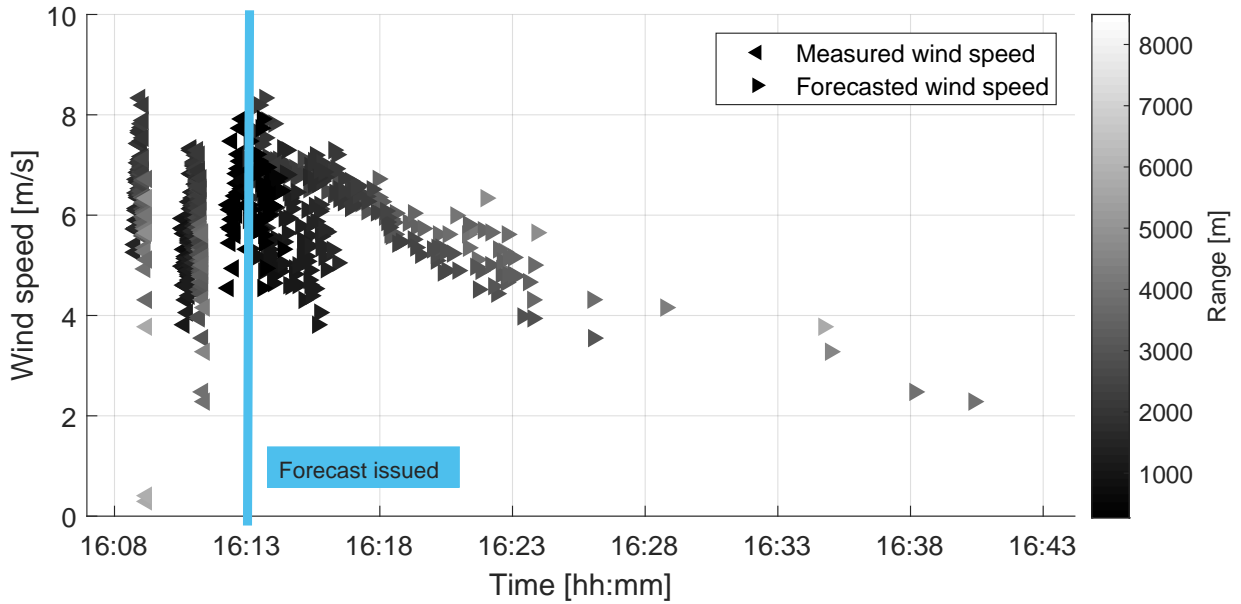


Figure 6.1: Timeline of measured wind speed that reaches the turbine and forecasted wind speed using Taylor’s frozen turbulence hypothesis. Measurement range that data originates from is given in shades of grey.

6.1 Methodology of probabilistic minute-scale forecasts

In Chapter 5.1 the method for the propagation of the measured wind speeds of the lidar is explained. This propagation is the basis for the forecasts which are the goal of this work. Each wind speed vector is propagated through time and space, and so it can be calculated at which time the wind speed vector reaches a certain target. This target is the wind turbine.

When the forecast is issued at a certain time, the last 30 minutes of measured lidar data are taken into account for the forecast. For each measured wind speed vector, it is checked whether it reaches the turbine and at which time after the forecast was issued. In the example data set in Figure 6.1 the wind speed measurements that reached the turbine is depicted. The corresponding forecasted wind speeds propagate along the timeline according to Taylor. Wind speed vectors that were measured close to the turbine with a high wind speed value are forecasted to arrive only a few minutes after the forecast was issued. Wind speed vectors measured in a farther distance from the turbine with a low wind speed value are forecasted to arrive later in time.

The spread of forecasted wind speeds is typical for this forecasting method. What it actually represents is an uncertainty in the wind speed forecast due to the turbulent nature of the wind, which also transfers into an uncertainty in the power forecast. For a specific time in the future, there is not only one forecasted wind speed, but a range of wind speeds. To deal with the


uncertainty in the forecast and to quantify it, probabilistic forecasting methods are applied.

Probabilistic forecasts are used in several different areas. The most common one is weather forecasts [40] that predict for instance the probability of rainfall and therefore help to decide whether to bring an umbrella or not. Probabilistic forecasts are also important for risk management in finances [41] or epidemiological studies [42]. In each area the goal is to assess the uncertainty of the forecast and use this information for decision making.


Probabilistic wind power forecasting was introduced in the early 21st century and theories and methods have been established since [6]. The method that was developed and applied in this thesis for probabilistic forecasting using lidar data is based on those findings. It follows the procedure depicted in Figure 6.2:

1. Transformation of wind speed to power: Using the power curve of the turbine, the forecasted wind speed is transformed into a forecasted power (Figure 6.2a). Details are in Section 6.1.1.
2. Binning of data: The forecasted wind speed and power are binned according to the forecast horizon (Figure 6.2b). For each bin a probabilistic forecast is calculated, which takes into account the spread of the wind speed or power data in the bin. Details are in Section 6.1.2.
3. Calculation of uncertainty of forecast as probability intervals: In each bin, the probability of the wind speed or power forecast is calculated using a Probability Density Function (PDF) and Cumulative Distribution Function (CDF). Using these tools, confidence intervals can be calculated that indicate the likelihood of a wind speed or power forecast value (Figure 6.2c). Details are in Section 6.1.3.
4. Evaluation of forecast: Using the actual measured wind speed of power observation of the respective forecast horizon, the probabilistic forecast is evaluated (Figure 6.2d). Details are in Section 6.1.4
5. Comparison with benchmark method: The probabilistic lidar forecast is compared to persistence, a state-of-the-art minute-scale forecasting method. Persistence is a deterministic model and is converted into a probabilistic forecast with historic wind speed and power measurements, to be able to compare it to the lidar forecast. Details are in Section 6.1.5.


Results of the forecasts for both sites are shown and discussed in Section 6.2 and Section 6.3.




(a) Transformation of wind speed into power forecast.



(b) Binning of data according to forecast horizon.



(c) Calculation of confidence intervals.



(d) Evaluation of forecast with observation

Figure 6.2: Probabilistic lidar-based forecast process.

6.1.1 Converting wind speed to power

To transform the forecasted wind speed into power, the power curve of the wind turbine is used. The power curve is the relationship between the produced power of the turbine and wind speed. The power curve is defined according to IEC standard 61400-12 by the wind turbine manufacturer. According to this standard, the wind speed needs to be measured in the free stream, ideally 2.5 rotor diameters in front of the turbine. The power data is given as 10-minute averages and binned in wind speed bins of 0.5 m s^{-1} . The power curve is calculated from the wind speed and power mean of all data in each bin and the uncertainty is given as the standard deviation of the power in each bin.

The power curve of the turbine in Stötten and the power curve of the AV04 were not provided by the turbine operator. Therefore they had to be generated from the available measurement data.

The power curve in Stötten is not IEC conform as wind speed data was only available from the nacelle anemometer and not from a free stream measurement. Figure 6.3 top, shows the scatter plot of normalised wind speed and power data. This data needed to be filtered in order to obtain a clean curve. Three filter criteria were applied and data was filtered out accordingly (Table 6.1).

Filter	Power	Wind speed
1	$P \leq 0$	$WS > 0$
2	normal production	too low
3	normal production	too high

Table 6.1: Filter criteria that were applied to clean the power curve in Stötten.

In the first category the turbine power is zero or below zero but the wind speed is above zero. Data in this category account for over a third of the total amount of available data and occurs throughout the whole year. It is assumed that during these periods the turbine was shut down due to malfunctioning or for maintenance. Data of the second category, where the wind speed is too low for the generated power, only occur in winter times. It is a typical sign that the anemometer was slowed down due to icing and therefore recorded lower wind speeds. Data of the third category, where the wind speed is too high for the generated power, are very few and occur throughout the year. They are discarded as temporary speed up of the anemometer which the turbine did not follow due to its inertia.

The power curve of the AV04 was calculated in the same way as the power curve for the Stötten turbine, but instead of the nacelle wind speed, the free stream wind speed from FINO1 was taken, and only undisturbed wind direction sectors were taken into account (Figure 6.3,

bottom).

The filtered data are then used to generate the binned power curve which is also shown in Figure 6.3. This curve is used as a look up table to transfer the forecasted wind speed into a forecasted power. The uncertainties of the power curve are neglected.

6.1.2 Choosing a forecast horizon

The second step in the probabilistic forecast process is to choose a forecast horizon and bin the forecasted wind speed accordingly (Figure 6.2b). The forecast horizon quantifies the period of time for which the forecast is valid. It is important to note that the forecast horizon as per definition is not a single time point in the future that relates to the forecasted power output in that instant, but a period of time during which the forecasted power will occur.

In this thesis a practical approach to choosing the forecast horizon is taken: as the forecast needs to be validated with actual power output measurements, the horizon is chosen according to the sampling rate of the turbine measurements. In case of the onshore site 10 minute averages of power output measurements are available. The same is the case for the offshore site alpha ventus. Hence 10 minute forecast periods are chosen. This means the forecast is calculated for a 0-10 minute period directly after the forecast is issued, and a 10-20 minute and 20-30 minute period respectively. Forecasts for more than 30 minutes ahead are not issued as no data is available for this horizon (cf. Section 7.1 on the influence of the measurement range on the forecast horizon).

When the forecast horizon is chosen, the forecasted wind speed and power data spread is binned accordingly. The data spread in each bin is analysed using probabilistic methods explained in the following.

6.1.3 Quantifying the uncertainty with probability density function and cumulative distribution function

The third step in the probabilistic forecast process is to quantify the uncertainty of the forecast which is represented in the data spread in each forecast horizon bin. The method to quantify the uncertainty is to estimate a probability density function (PDF) $f_X(x)$ in each bin. The values of the function f can be interpreted as a relative probability that a random variable X in the sample space (the set of possible values x taken by the random variable X) would equal that sample [43].

It is differentiated between parametric, and non-parametric approaches to define a PDF.




Figure 6.3: Binned normalised power curve (grey line) from the turbine in Stötten (top) and the AV04 in alpha ventus (bottom) with standard deviation and filtered data. Data are 10 minute averages. Numbers in brackets give the number of data points in each category.

Parametric approaches to PDFs for forecasting purposes are based on the assumed shape of probability density, e.g. Gaussian distribution [6] and can be described by an analytical expression. They are low in computation time but the assumed shape of the distribution may be incorrect and should change with time, if e.g. the wind speed conditions change. A well known PDF in wind energy is the Weibull distribution used for site assessment [44].

Non-parametric approaches to estimate the density distribution do not assume a shape of the density curve but are data driven. They estimate the density at a finite number of points and then a PDF is obtained through interpolation between these points [6].

A very simple non-parametric approach is to calculate the histogram of a data set. In a histogram the height of each bar indicates the number of data point in each bin. To obtain a PDF the histogram is normalised so that the area of each bar is the relative number of data points. The sum of the bar areas then is less than or equal to 1 (Figure 6.4, top).

A very popular non-parametric approach to PDF estimation often used in forecasting is the Kernel Density Estimation (KDE) method. Similar to a histogram, it is a data driven approach but smooths the curve [6], cf. Figure 6.4, top. The KDE, also known as the Parzen-Rosenblatt window method is defined as [45]

$$f_h(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - x_i}{h}\right)$$

where x_1, x_2, \dots, x_n are random samples from an unknown distribution, n is the sample size, K is the kernel function, and h is the bandwidth or smoothing parameter. Similar to a histogram, the kernel distribution builds a function to represent the probability distribution using the sample data. But unlike a histogram, which places the values into discrete bins, a kernel distribution builds the PDF by creating an individual probability density curve for each data value, then summing the smooth curves. This approach creates one smooth, continuous probability density function for the data set [46].

Both approaches, the histogram and the Kernel approach, are used for wind power forecasting to estimate a PDF of a forecasted wind speed. In Figure 6.4, top, the histogram and Kernel approach to estimate a PDF are compared. In this example, a data set of 0-10 minute forecasted power data from the onshore site is used, with a number of 85 forecasted power values. The histogram and KDE have a similar shape, but the smoothing effect of the KDE is clearly visible.

Testing the forecast with both approaches for a longer period, it was decided that the histogram approach will be used for this thesis. For forecast horizons with low number of forecasted values, the smoothing effect of the KDE led to unrealistic power forecasts. For example when only a few power values around rated power were forecast, the KDE PDF estimated values




Figure 6.4: Example of uncertainty information for forecasted power data from Figure 6.2. Top: Comparison of PDFs from histogram and Kernel Density Estimation. Frequency is normalised so that the height of each bar is (number of observations in the bin) / (total number of observations * width of bin). The area of each bar is the relative number of observations. The sum of the bar areas is 1. Bottom: Comparison of CDFs from histogram and Kernel Density Estimation. Exemplary quantiles and intervals and how they can be read from PDF and CDF are marked in blue.

above rated power to smooth the curve and make it continuous. The simple approach of the histogram with no assumptions and no smoothing applied, produced more reliable results in such cases.

After the PDF is calculated, the next step is to calculate the Cumulative Distribution Function (CDF) (Figure 6.4 bottom). Similar to the PDF, the CDF $F_X(x) = P(X \leq x)$ gives the probability P that the random variable X is less or equal to x . The CDF of a continuous random variable X can be expressed as the integral of its probability density function f_X as follows:

$$F_X(x) = \int_{-\infty}^x f_X(t) dt$$

With the help of PDF and CDF, the uncertainty of the forecast can be quantified. The tools for that are quantiles or intervals which are used to describe uncertainty. A quantile $Q_\alpha(x)$ is defined so that the probability of the observation X being smaller than $Q_\alpha(x)$, is equal to the proportion α , with $\alpha \in [0, 1]$. This property can also be expressed with the CDF so that $F_X(Q_\alpha(x)) = \alpha$ but unlike the CDF the unknown value that has to be predicted is x , not P [30].

The most common quantile with $\alpha = 0.5$ is the median, a value separating the higher half from the lower half of the data sample. Intervals on the other hand, intend to provide a range where the random variable (or rather the forecast in this case) lies, defined by a given probability. Therefore they are referred to as confidence intervals. For example the 50% confidence interval lies between the 0.25 and 0.75 quantile (6.4). That means that 50% of the forecasted values lie within this interval. For the forecast it means that the probability is 50% that the actual wind speed or power is in that range.

Intervals are ideal to visualize the probability of a wind power forecast within a certain range of wind power output (cf. Figure 6.2). Even though they have less information content than PDFs and CDFs, they are often applied, as the range of a forecast is crucial information for decision making [6]. Since they do not make any assumption of the distribution shape, approaches using quantiles or intervals are classified as non-parametric as well. Figure 6.4 shows how PDFs, CDFs, quantiles, and intervals can be visualized and how they are related.

6.1.4 Evaluating probabilistic forecasts

Once the forecast is issued and the uncertainty information is obtained, the next step is to evaluate the forecast using appropriate criteria.

The evaluation is carried out using the actual observation, which means the measurement of the forecasted wind speed or power. For a deterministic or spot forecast an evaluation is simple, as only the one forecasted value is compared to the observation, using error measures such as the mean average error (MAE) or the root mean square error (RMSE). For probabilistic forecast, the evaluation is not so straightforward, because there is a spread of forecasted values. This section first explains what the methods for the evaluation of probabilistic forecasts are and then describes the necessity to prepare the data so it can be used as a reliable reference.

Methods for evaluation

For probabilistic forecast evaluation, different methods are used and different properties of the forecast are evaluated. The main required properties of a probabilistic forecast are [47]:

- **Reliability:** is referred to as the statistical consistency of the predictive distribution and the observation. This property is also called probabilistic calibration. It means that the observed and forecasted probabilities should be as close as possible and the probabilistic forecasts is therefore unbiased. It should be noted that this is a joint property of the prediction and the observation measurements. If the measurements are not well calibrated, the forecast will be biased due to the error in the measurements [48].
- **Sharpness:** is the concentration of the predictive distribution. It is a property only of the predictive distribution and not of the observation. The smaller the scatter in the forecasted data, the smaller is the predictive interval and the sharper is the forecast [48].
- **Skill:** is reliability and sharpness in one score and provides the whole information of uncertainty forecasting performance. The definition depends on the scoring rule, but one requirement is that the skill score is proper, which means that the best forecast has the best score [6].

A reliable or well calibrated forecast which is not biased and the forecasted probability is as close as possible to the actual observed probability. A tool to check the calibration is the Probability Integral Transform (PIT) histogram [48]. The PIT histogram is a simple visual tool to evaluate the reliability of a probabilistic forecast. The PIT is defined as

$$PIT = F(O)$$




Figure 6.5: Examples of PIT histograms of different forecasts. The forecasts can be miscalibrated as in (a)-(c) or well calibrated as in (d). The ideal frequency is marked as horizontal line. As the probabilities are binned in 10 bins, the ideal frequency for each bin is 0.1.

where PIT is the value that the predictive CDF F of the forecast takes at the observation O . The value of PIT is therefore in the interval $[0, 1]$. The cumulative distribution function F of a probabilistic forecast is well calibrated, if PIT is uniformly distributed [49]. This means that statistically, all probabilities are forecast an equal amount of time.

The values PIT of a forecast are analyzed in a PIT histogram and for a reliable or well calibrated forecast, the observed frequency of the PIT is even (Figure 6.5 (d)). If the forecast is not well calibrated, the frequencies are not evenly distributed. Examples of miscalibrated forecasts and a calibrated forecast are shown in Figure 6.5 (a) to (c). Forecasts can be biased, under- or over-dispersed. In all cases, the PIT is not evenly distributed, but the observation tends to take a specific value of the CDF. For example in case of an under-dispersed forecast, the forecasted values do not have enough spread (or dispersiveness) and therefore the outer bins of the PIT histogram are overpopulated.

The sharpness and skill of a forecast is checked with the Continuous Ranked Probability

Score (CRPS). The CRPS is one of the most widely used accuracy metrics where probabilistic forecasts are involved [6]. It is defined as the area between the predictive CDF and a step-function at the observation. In mathematical terms it is defined as

$$CRPS(F, O) = \int_{-\infty}^x (F(x) - \mathbb{1}\{x \leq O\})^2 dx$$

where $\mathbb{1}\{x \leq z\}$ denotes the Heaviside step-function, returning one if $x \leq O$ and zero otherwise [50].

The CRPS is expressed in the same unit as the observed variable. The CRPS generalizes the mean absolute error (MAE); in fact, it becomes the MAE if the forecast is deterministic. Just like the MAE, the theoretically ideal CRPS for the power prediction is 0 and the worst possible score is 1.

Figure 6.6 shows an examples of two predictive PDFs and the corresponding CDFs with high and low sharpness. A sharp PDF results from a low dispersion of the forecasted values which is indicated by the narrow width of the 50 % confidence interval, meaning the interval that covers 50 % of the data. In the CDF, the sharper PDF corresponds to a steeper cumulative curve. The steeper CDF results in a lower CRPS as the area between the CDF and the step function at the observation becomes smaller.

Preparing observational data

The wind speed and power data measured at the turbine is the basis for the evaluation of the forecast. Therefore it is important to prepare it with the same care as the lidar data. This means

1. filtering out outliers that would corrupt the forecast evaluation,
2. making sure that lidar data and turbine data have the same time stamp.

In Section 6.1.1 it was already explained how the wind speed and power data of the onshore and offshore site were filtered, in order to calculate the power curve of the respective turbines. For evaluating the forecast, only observational data is used where the same filter criteria were applied: the turbine was actually running and producing power; the power was not curtailed and the turbine was operating at full rated power; and the wind speed was not corrupted by a malfunctioning wind speed sensor due to icing.

For the offshore data, a comparison of the wind speed data of the turbine and the wind speed data of the met mast in the time domain shows a time delay (Figure 6.7). It was noted that a measured wind speed peak occurred in the turbine data with a time delay of 10 min 48 sec. As the met mast and the turbine at the site were not directly next to each other, a delay




Figure 6.6: Predictive PDF (top) and CDF (bottom) of a forecast with low sharpness (dashed line) and high sharpness (solid line). The step function is taking the value 1 at the observation (blue solid line). The CRPS is indicated as the area in blue between the step function and the respective CDF. The 50% confidence interval is indicated in each PDF.




Figure 6.7: Comparison of the wind speed measurements at the met mast (black) and the turbine (grey) at the onshore site in Stötten for a time period with a strong change in wind speed. The original turbine data (dark grey) was time shifted (light grey) to fit the UTC time stamp of the met mast.

was expected (the wind direction in that period was coming from the met mast towards the turbine). However, with a separation of 1350 m and a wind speed peak of 12.68 m s^{-1} , applying Taylor’s hypothesis, this delay should be around 1 min 46 sec.

Both the lidar and the met mast were operated by the University of Stuttgart and it was known that they recorded the time stamp in Universal Time Coordinated (UTC). Therefore the time stamp of the turbine data was shifted to match the the lidar’s time stamp. A correlation of the wind speed between mast and turbine before (Figure 6.8a) and after time correction (Figure 6.8b) proved that the time correction was valid and necessary for the whole turbine data set, as the slope, offset and coefficient of determination as shown in the figures improved.

6.1.5 Comparing to benchmark forecasting method persistence

After describing how to calculate the probabilistic lidar forecast, and how to evaluate it, the question is how to compare this new forecasting method to state-of-the art minute-scale forecasting techniques.

In this thesis, the lidar forecasts are compared to the persistence method. Persistence is the classical benchmark method because it is the simplest and therefore most cost effective method [51]. Persistence assumes that the system is in equilibrium and only changes slowly, so that it is most likely that the forecasted value will be the same as the present value. In other words, persistence assumes that the forecasted wind speed or power in the future will be the same as when the forecast is issued. If $v(t)$ and $P(t)$ are the wind speed and power at the time t the forecast is issued, then the forecasted values can be formulated as the following terms:

$$v(t + \Delta t) = v(t) \quad (6.1)$$




Figure 6.8: Correlation of the wind speed [m/s] measured at turbine and met mast before and after time correction, to prove that the correction of the turbine time stamp is valid.

$$P(t + \Delta t) = P(t) \quad (6.2)$$

where Δt is the forecast horizon.

Since wind speed and power measurements are available for every wind turbine through the SCADA system, persistence is the chosen method for minute-scale forecasts. Especially in the minute-scale forecast time range the forecasts are very accurate, but the accuracy decreases rapidly when the time-scale of the forecast increases [12].

The above described persistence method only works as benchmark for deterministic forecasts as only one value is forecasted for a specific forecast horizon. To obtain a predictive distribution of forecasted values, the point forecast can be dressed with recent observed values of the persistence error, as suggested in [48]:

$$\{\max(O_t - O_{t-h} + O_{t-h-n}, 0) : h = 0, \dots, 17\} \quad (6.3)$$

Here O are the recent values of the observed wind speed or power, t is the time stamp when the forecast is issued, h denotes the time steps of recent data that should be taken into account and n is the number of forecast horizons. As suggested in [48] the last 18 time steps, corresponding to the last 3 hours of available wind and power data, are used to calculate a set of forecast values. As the forecast is calculated for three forecast horizon bins (0-10min,

time	12:00	12:10	12:20	12:30	12:40	12:50	13:00
P [kW]	300	400	600	200	800	1000	1200
h	6	5	4	3	2	1	0

Table 6.2: Example timeline of turbine power data used to calculate a predictive distribution of a persistence forecast. h are the time steps that are taken to calculate the predictive distribution from the time, when the forecast is issued at $h = 0$.

n (forecast horizon)	h	0	1	2	3	4	5	6
		1000	1000	600	1600	1000	1100	-
1 (13:00-13:10)		800	400	1000	1400	900	-	-
2 (13:10-13:20)		200	800	800	1300	-	-	-
3 (13:20-13:30)								

Table 6.3: Example of probabilistic persistence power data sets. The values are calculated from data in Table 6.2 and Equation 6.3 for three forecast horizons n and seven time steps h . The values are given in kW.

10-20min, 20-30min), n is 1,2 or 3. For each bin, the suggested approach results in a set of predictive wind speed and power values that can be used to calculate a PDF and CDF that represent the probabilistic persistence forecast.

To explain in more detail, how the predictive distribution is calculated, Table 6.3 gives the results of the persistence power forecast data sets which are calculated using Equation 6.3 and data from Table 6.2. In this example 7 time steps are available ($h_{max} = 6$), the forecast is issued at the last time stamp of the example data set, and three forecast horizons n are considered ($n_{max} = 3$). The first value in the forecast set for $h = 0$ and $n = 1$ is calculated as $1200\text{kW} - 1200\text{kW} + 1000\text{kW} = 1000\text{kW}$, and the other values are calculated accordingly. The result is a set of forecasted power values for each forecast horizon n . In this example, the sets do not have an equal amount of values. To calculate a full set of power values for each forecast horizon, at least $h_{max} + 1 + n_{max}$ measured power values need to be available at the time the forecast is issued.

Figure 6.9 shows the persistence forecast for the same day and time as used as an example in Figure 6.2 to explain the lidar forecast procedure. The forecast horizons are set equally as to the lidar forecast (0-10min, 10-20min, 20-30min) and for each horizon, a set of persistence power forecasts is calculated according to Equation 6.3. In each forecast horizon, it is made sure that an equal amount of 18 predictive power values are available. In Figure 6.9, the time for each persistence forecast is set to the bin center, but in reality, the values of the persistence forecast in each forecast bin are not allocated to a specific time in the bin.

The procedure to extract probabilistic information from the persistence forecast, is then




Figure 6.9: Example of a probabilistic persistence power forecast.

exactly the same as for the lidar forecast: first a PDF and CDF for the forecasted set of values in each bin is calculated using the histogram approach. With this probabilistic data, the uncertainty can be quantified and visualised, e.g. using intervals. The evaluation of the forecast is then carried out using the observed wind speed or power to calculate the PIT and CRPS (compare Section 6.1.4). Using the same methods for quantifying the uncertainty in the probabilistic lidar and persistence forecast, and using the same tools for the evaluation, allows both methods to be directly compared.

6.2 Results of minute-scale forecasts from the onshore campaign

In the last sections, the methodology of minute-scale forecasting was explained. This section analyses the results of the minute-scale forecasts of the onshore campaign. First the forecast availability is discussed, followed by the results of the forecast in the time domain, and the statistical analysis of the forecast calibration, sharpness and skill.

6.2.1 Availability of forecasts

For the onshore campaign, lidar data from the measurements on top of the radio tower were analysed for the period from July 1st, 2016 to July 31st, 2017. During this time the lidar measured with a horizontal scan towards the main wind direction (cf. Table 4.1 in Section 4.1.7). Between November 3rd, 2016 and January 19th, 2021, no measurements were available. As the scan direction was fixed, but the wind direction and thus the turbine turned, forecasts could only be carried out when the lidar actually measured into the wind direction. Also, forecast were only carried out when there was a wind speed and power observation from the turbine available to validate the forecast.

In Stötten, the power data was available for evaluation of the forecast for 45.3% of the time (Figure 6.10, top). The gaps result from filtering of the data, as explained in Section 6.1.4. The forecasts with the forecast horizon 0-10min were only possible 17.2% of the time (Figure 6.10, bottom). This means in 28.1% of the cases when a forecast could have been carried out, either the wind direction did not match the scan direction, or the lidar did not measure (e.g. due to foggy conditions) hence, a forecast was not possible.

The forecast were carried out for three forecast horizons 0-10min, 10-20min and 20-30min, and for each horizon the total number of available forecasts is analysed (Table 6.4). For the lidar-based forecasts, each of the horizons has a different availability, as the number of forecasts depends on the availability of the forecasted wind speed vectors. The availability of the forecasts for the persistence method also varies, but not as strongly. When a persistence forecast is possible depends only on the number of the historic measurements of turbine data. If the last 18 wind speed or power measurements are available, a 0-10min forecast is possible (19 and 20 measurements for the 10-20min and 20-30min forecast respectively), cf. Section 6.1.5.

	0-10min	10-20min	20-30min
Lidar-based	9877 (71.9 %)	4137 (61.9 %)	1103 (56.6 %)
Persistence	22,892 (94.6 %)	22,354 (93.6 %)	21,510 (92.2 %)

Table 6.4: Total number of available lidar-based and persistence forecasts for different forecast horizons during the onshore campaign. The availability in percent is given in respect of the availability of power data.

To conclude, the availability of lidar forecasts at Stötten is significantly lower compared to the availability of the persistence forecast for all forecast horizons. The reason is the measurement setup, as the lidar measures with a fixed trajectory that does not change the measurement direction with the wind direction. To improve the availability, the lidar should be installed




Figure 6.10: Availability of data during the onshore campaign. Top: power data after filtering. Bottom: 0-10min forecasts for periods where power data was available. Periods where no data is available are marked white.

on top of the nacelle (as in the offshore campaign in this thesis), or the trajectory should be adapted automatically, depending on the wind direction.

6.2.2 Forecasts in the time domain

Figure 6.11 shows an example of a 0-10min forecast horizon from March 1st, 2017. During the period, the actual wind speed was initially around 10 m s^{-1} and then decreased to around 6 m s^{-1} . During the period shown, the lidar-based forecast is able to forecast the wind speed well, in that the observation is lying within the forecasted confidence interval. Changes in the wind speed are anticipated and there are not many 10-minute periods where there is no forecast available.

The power forecast in the period first fluctuates around the maximum power, as the wind speed is close to rated wind speed. Along with the wind speed, the observed power then decreases and the forecasts follow this trend. The confidence intervals for each 10-minute forecast are however much broader than the intervals of the wind speed forecast. This results from the fact that power change is proportional to the third power of the wind speed change in the below rated area of the power curve.

The period shown in Figure 6.11 is actually not a typical example of the forecast results in Stötten. It was chosen because there were only few 10-minute periods where there was no forecast available, and the observation was within the confidence interval of the forecast. The visualization of the forecasts with intervals is however, not a good method to analyse the results statistically, as it only shows a limited forecast period that might not be representative for the whole forecast method. The statistical analysis will be done in the following, using appropriate tools.

6.2.3 Calibration

As explained in Section 6.1.4, to evaluate a forecast first the reliability, or calibration of the forecast should be checked. In order to do that, the value of the predicted wind speed and power cumulative distribution function (CDF) at the point of the actual wind speed and power observation was determined for each forecast horizon. As explained in Section 6.1.4, this is called the Probability Integral Transform (PIT). The PIT values were then plotted in a PIT histogram for the wind speed (Figure 6.12) and power forecast (Figure 6.13). The frequency is normalised, so the sum of the bar heights is less than or equal to 1. This analysis was carried out both for the lidar-based forecast and for the benchmark forecast method (persistence).

For a well calibrated forecast, the bars in the PIT histogram should have an equal height.




Figure 6.11: Example of the probabilistic wind speed (top) and power forecast (bottom) for the 0-10min forecast horizon at the onshore site in Stötten. Data are for March 1st, 2017. Three different confidence intervals are given (shades of blue). The actual wind speed and power observations are given for comparison (black). Periods where no forecast is available are marked (grey).




Figure 6.12: PIT histograms for onshore lidar wind speed forecasts (left) and persistence wind speed forecasts (right) for different forecasting horizons. The ideal frequency is marked as horizontal line. As the probabilities are binned in 10 bins, the ideal frequency for each bin is 0.1.




Figure 6.13: PIT histograms for onshore lidar power forecasts (left) and persistence power forecasts (right) for different forecasting horizons. The ideal frequency is marked as horizontal line. As the probabilities are binned in 10 bins, the ideal frequency for each bin is 0.1.

This height depends on the number of bins that is used for the histogram. The analysis of the PIT histograms is a visual tool to check calibration and does not include a statistical analysis of the PIT values.

The PIT histograms of the lidar-based wind speed and power forecast shown here, are not very well calibrated. The 0-10min and 10-20min forecasts are under-dispersed, which means the forecasted distribution does not have enough spread (or dispersion). Therefore the actual observation occurs often either at the lower end of the CDF, or even outside of it, or at the high end of the CDF (or outside of it). In other words, the lidar-based forecasts are biased and have a tendency to be either too high or too low.

The lidar-based forecasts of the 20-30min forecast horizon are heavily biased. In this forecast-horizon, the actual observation is most of the time at the high end of the CDF or outside of it. This means that the lidar-based forecasts predicts a wind speed and power that is too low. There are much less forecasts available compared to the other forecast horizons, which is a results of the lidar measurement: a forecast is only available if there are measurements from far ranges, or with low wind speeds that will actually arrive at the turbine with a delay of 20-30 minutes after the measurement took place. This also explains the bias: if only low wind speeds are available in this forecast horizons, the forecasted power will be too low and therefore biased. This further emphasizes the need to increase the range of wind lidars used for this application.

To sum up, the lidar-based forecasts of the onshore campaign are reasonably well calibrated in the first two forecasting horizons, but heavily biased for the forecast horizon of 20-30min. The calibration shows the limitations of lidar-based forecasts using Taylor's hypothesis: for bigger forecast horizons, there are less forecasts available, because the availability depends on the measurement range and the wind speed. Only low wind speeds of far ranges need such a long time to arrive at the turbine, and therefore the forecast in the 20-30min horizon is biased and too low. A more complete physical model might reduce this effect.

The calibration of the persistence forecast is better for all forecast horizons. In the persistence PIT diagrams for wind speed and power, the bars are evenly distributed, indicating that there is no bias in the forecast. Since this method just relies on historic turbine measurements and projects the persistence error of available historic data into the future, there are also an equal amount of forecasts available for all forecast horizons.

6.2.4 Sharpness and skill


To evaluate the sharpness and skill of the lidar-based forecasts, the Continuous Ranked Probability Score (CRPS) was calculated for the wind speed and power forecasts of each forecast horizon (as explained in Section 6.1.4). The CRPS is an error metric of the probabilistic forecast, with the same unit as the forecasted values, i.e. wind speed and power. The smaller the CRPS, the smaller the deviation between the actual observation and the prediction.

To statistically analyse the CRPS of the forecasts, a boxplot was created for both the wind speed and power forecasts (Figure 6.14). In addition, histograms are calculated of the wind speed (Figure 6.15) and power forecasts (Figure 6.16) to analyse the distribution of the respective data sets. The forecasts are analysed for different forecast horizons, and compared to the benchmark forecast method.


The lidar-based wind speed and power forecasts of the onshore campaign have a higher CRPS value than the persistence forecast for all forecast horizons. This means that the predictive wind speed and power spread used to calculate the CDF of the lidar-based forecast is higher. Therefore, the forecast is less sharp than the persistence forecast and the actual wind speed or power does not lie within the predicted interval.

It is noted that for both the lidar-based and the persistence forecast, the CRPS increases with increasing forecast horizon. This is expected, as the calibration has shown that (especially for the 20-30min forecast horizon) the lidar-based forecast is significantly biased.

For the onshore campaign, it is concluded that the lidar-based forecast method as it is carried out in this thesis does not offer value over the persistence method. Due to the fact that the campaign was carried out in complex terrain, the forecast error of the lidar-based method is higher in all forecast horizons. The simple Taylor frozen turbulence assumption does not account for the complex flow at the site and the wind speeds that are measured in a distance in front of the lidar are not transported to the turbine using the Taylor hypothesis. Therefore, the preview information provided by the lidar does not add value to the forecast.



(a) Wind speed forecast



(b) Power forecast

Figure 6.14: CRPS of lidar and persistence forecast from the onshore campaign for different forecast horizons. Central mark indicates the median; bottom and top edges of box indicate $q_1 = 25th$ and $q_3 = 75th$ percentiles; whiskers extend to extreme data points not considered outliers; points are considered outliers if they are greater than $q_3 + 1.5 \cdot (q_3 - q_1)$ or less than $q_1 - 1.5 \cdot (q_3 - q_1)$; outliers are plotted individually.




Figure 6.15: Histograms of CRPS from onshore wind speed forecast for lidar (top) and persistence (bottom) forecasts for different forecasting horizons. Median of data is marked as grey line. Note: y-axes limits are not the same for lidar and persistence histograms; x-axes are limited to 5 m/s and do not comprise all outliers.




Figure 6.16: Histograms of CRPS from onshore power forecast for lidar (top) and persistence (bottom) forecasts for different forecasting horizons. CRPS given in percentage of rated power. Median of data is marked as grey line. Note: y-axis limits are not the same for lidar and persistence histograms; x-axes are limited to 60% and do not comprise all outliers.

6.3 Results of minute-scale forecasts from the offshore campaign

This section analyses the results of the minute-scale forecasts of the power generated by the AV04 turbine of the offshore campaign at alpha ventus. First the forecast availability is discussed, followed by the results of the forecast in the time domain, the statistical analysis of the forecast calibration, and assessment of sharpness and skill for different conditions.

6.3.1 Availability of forecasts

For the offshore campaign at alpha ventus, lidar data from the measurements on top of the AV04 turbine were analysed for a period from March 1st to December 31st, 2020. During this time the lidar measured the inflow of the turbine on top of the nacelle with a horizontal scan (cf. Table 4.2 in Section 4.1.7). Between May 27 and June 25, 2020, no measurements were available, because after a shut down of the wind farm the measurement systems of the turbine could not be restarted. Forecast were only carried out when there was a wind speed and power observation from the turbine available to validate the forecast.

In alpha ventus, the power data of the AV04 was available for evaluation of the forecast for 64.0% of the time (Figure 6.17, top). The gaps result from filtering of the data, as explained in Section 6.1.4. The forecasts for the forecast horizon 0-10min were available at 62.2% of the time (Figure 6.17, bottom). This means in only 1.8% of the cases, when a forecast could have been carried out, a forecast was not possible. The availability of the forecasts is much higher compared to the onshore campaign. In contrast to the onshore campaign, the lidar on the AV04 was always looking into the wind direction, because it rotated along with the nacelle. This proved to be a big advantage, because it means the reason that a forecast was not possible could not have been that scan direction and wind direction were not aligned, but that the lidar did not measure due to e.g. foggy conditions. This did not happen very often (as mentioned, only in 1.8% of the time).

Table 6.5 summarised the availability of forecasts at the offshore campaign. As for the onshore campaign, the forecasts of the AV04 were carried out for three forecast horizons 0-10min, 10-20min and 20-30min, and the number of available forecast was analysed. The lidar-based forecast has a different availability for each of the horizons, as the number of forecasts depends on the availability of forecasted wind speed vectors. The number of available persistence forecast is actually lower in the 0-10min horizon compared to the lidar-based forecast. This results from the fact that the probabilistic persistence forecasts needs several measurements from the last minutes to be able to be carried out. If these are not available, e.g. due to a technical reasons, a persistence forecast cannot be issued, whereas a the lidar forecast can be carried out. However,




Figure 6.17: Availability of data during the offshore campaign. Top: power data after filtering. Bottom: 0-10min forecasts for periods where power data was available. Periods where no data is available are marked white.

in contrast to the lidar forecast availability, the availability of the persistence forecast, decreases only slightly for the other forecast horizons.

	0-10min	10-20min	20-30min
Lidar-based	27,776 (98.2%)	7327 (52.4%)	963 (38.2%)
Persistence	26,996 (96.5%)	26,724 (95.8%)	26,132 (95.5%)

Table 6.5: Total number of available lidar-based and persistence forecasts for different forecast horizons during the offshore campaign. The availability in percent is given in respect of the availability of power data.

To conclude, the availability of the lidar-based forecast in the first forecast horizon at the offshore campaign is significantly higher than at the onshore campaign (98.2% compared to 71.9%) and is also slightly higher compared to persistence. The reason for the increased availability is the measurement setup on top of the nacelle. As the lidar automatically measures into the wind direction, the only reason why a forecast is not possible is if no or not enough measurements are available (e.g. due to foggy conditions). In the 10 months measurement campaign this happened only about 1.8% of the time. However, the availability of the lidar based forecast for forecast horizons of more than 10 minutes is significantly lower compared to the persistence method in the offshore case.

6.3.2 Forecasts in the time domain

To analyse the offshore forecasts, the forecasted confidence intervals of wind speed and power are analysed in the time domain. To that end, the 90%, 50% and 10% confidence interval are plotted together with the actual wind speed and power observation (Figure 6.18). In the example, the measured wind speed first fluctuates around rated wind speed, it then decreases within 40 minutes to around 7 m s^{-1} and even further in the next hours.

The lidar forecast is able to anticipate the wind speed and power fluctuations and also the decrease in wind speed (and power) is forecasted. During the whole period, the actual wind speed and power was in the forecasted intervals. It is noticeable, that the forecasted wind speed intervals are broader for higher wind speeds. Thus the corresponding forecasted power intervals are proportionally larger as well, as the fluctuation in wind speed is proportional to the third power of the change in power forecast.

The period shown in Figure 6.18 is representative of the forecasts at alpha ventus for the AV04. Compared to the onshore campaign, it was easy to find time periods, where forecasts were available during several consecutive hours without gaps. A trend was noted that for high




Figure 6.18: Example of the probabilistic wind speed (top) and power forecast (bottom) for the 0-10min forecast horizon of the AV04 at the offshore site at alpha ventus. Data are for March 15, 2020. Three different confidence intervals are given (shades of blue). The actual wind speed and power observations are give for comparison (black). The wind direction was 265° and the AV04 was not affected by the wakes of its neighbouring turbines.

wind speeds, the lidar-based forecast tended to overestimate the wind speed (as also notable in the first hour of the example period). However, to analyse the forecasts statistically for the whole period, the calibration and sharpness are checked in the following step.

6.3.3 Calibration

To check the calibration, and to find out if the forecasts are biased, the PIT values for each forecast were calculated and then plotted in a PIT histogram for the wind speed (Figure 6.19) and power forecast (Figure 6.20). This analysis was carried out for the lidar-based forecast and for the benchmark forecast method persistence.

For the PIT histograms of the offshore campaign, the forecasts were filtered. Only values were taken into account when the AV04 was not affected by the wakes of its neighbouring turbines. As the wakes represent a decrease in wind speed behind the rotor due to the energy extraction, and this effect is not covered in the forecast model, the wake situations would corrupt the forecasts.

The calibration of the lidar-based offshore forecast actually differ from the onshore forecast for the first forecasting horizon. The PIT distribution of the wind speed and power forecast is over-dispersed (and not under-dispersed as in Stötten). This means that statistically, the wind speed and power observations are not equally spread over the corresponding predictive CDFs, but tend to be between the 0.3 and 0.7 quantile of the predictive distribution. Hence the forecasted distribution has too much spread (or dispersion). What it also means however, is that the observed wind speed and power is most of the time within the forecasted confidence intervals for the 0-10min forecast horizons. For the onshore campaign, the forecast of this first horizon tended to be outside of the forecasted intervals.

For the forecast horizon 10-20min and 20-30min, the calibration of the lidar forecast is biased towards high PIT values, just as it was for the onshore campaign. This means the actual observation tends to be at the upper end or even outside of the forecasted CDF distribution - or in other words, the forecasted wind speed and power is biased and has a tendency to be too low. This bias is more pronounced for the 20-30min horizon than the 10-20min horizon.

To conclude, in contrast to the onshore forecast, the calibration of the first horizon of the lidar-based offshore forecast shows that the forecast and the observations match well, but the predictive distribution tends to have too much spread. The calibration of the lidar-based forecasts of the second and third horizon, shows the same limitations of the method as has been observed onshore: using Taylor's advection model, mostly slow wind speeds that are measured in far range gates reach the turbine with a delay of 10 minutes and more. This leads to a bias in the forecast, and the forecast is too low.

The calibration of the persistence method is the same offshore as it was onshore: all forecasting horizons are fairly well calibrated and show a slight trend towards under-dispersion: the predictive CDFs do not have enough spread and the actual observation tend to be at the lower and upper end of the CDF.




Figure 6.19: PIT diagrams for offshore lidar wind speed forecasts (left) and persistence wind speed forecasts (right) for different forecasting horizons. The ideal frequency is marked as horizontal line. As the probabilities are binned in 10 bins, the ideal frequency for each bin is 0.1.




Figure 6.20: PIT diagrams for offshore lidar power forecasts (left) and persistence wind speed forecasts (right) for different forecasting horizons. The ideal frequency is marked as horizontal line. As the probabilities are binned in 10 bins, the ideal frequency for each bin is 0.1.

6.3.4 Sharpness and skill


To analyse statistically the sharpness and skill (and thus the error) of the forecasts, a boxplot of the CRPS was created of both the wind speed and power forecasts (Figure 6.21). In addition, histograms are calculated of the wind speed (Figure 6.22) and power forecasts (Figure 6.23) to analyse the distribution of the respective data sets. The forecasts are analysed for different forecast horizons, and compared to the benchmark forecast method. As for the calibration, the data are filtered such that only forecast are taken into account, when the AV04 turbine is not affected by the wake of its neighbors.

For the first 0-10min forecast horizon, lidar-based and persistence wind speed forecasts result in a very similar CRPS distribution, with most of the CRPS below 1 m s^{-1} . The median of the lidar-based forecast is only 0.04 m s^{-1} higher (Table 6.6), but the spread of the CRPS is smaller. The lidar-based wind speed forecast errors increase significantly for the second and third forecast horizons. It should be noted however, that there are less data available in those horizons, as the CRPS distributions show in Figure 6.22. The persistence error in contrast increases only slightly.


	0-10min	10-20min	20-30min
Lidar-based	0.34	0.69	1.53
Persistence	0.30	0.38	0.44

Table 6.6: Median of CRPS in m/s of lidar-based and persistence wind speed forecasts.

The CRPS for the power forecasts show a similar trend, only that the error for the lidar-based forecast does not increase as much as for the wind speed forecast with increasing forecast horizon. This phenomenon is against the assumption that an increased error in the wind speed forecast, results in a proportionally increased error by the third power for the power forecast. However, this is not the case, and the explanation is found in the power curve. For wind speeds lower than cut-in wind speed (as they are mostly forecast for the third forecast horizon), the corresponding power is zero. Therefore the forecasted power values in this forecast horizon do not vary but the distribution is small. In comparison the forecasted wind speed values vary less. This results in a predictive distribution which is very steep for the power forecast, but not for the wind speed forecast (Figure 6.24). In such a case, the corresponding CRPS of the power forecast does not increase for the third forecast horizon, as it does for the wind speed forecast, because the area between the CDF and the step function of the corresponding observation is small. Therefore, the power forecast error is small, although the wind speed forecast error is high.



(a) Wind speed forecast



(b) Power forecast

Figure 6.21: CRPS of lidar and persistence forecast from the offshore campaign for different forecast horizons. Central mark indicates the median; bottom and top edges of box indicate $q_1 = 25th$ and $q_3 = 75th$ percentiles; whiskers extend to extreme data points not considered outliers; points are considered outliers if they are greater than $q_3 + 1.5 \cdot (q_3 - q_1)$ or less than $q_1 - 1.5 \cdot (q_3 - q_1)$; outliers are plotted individually.

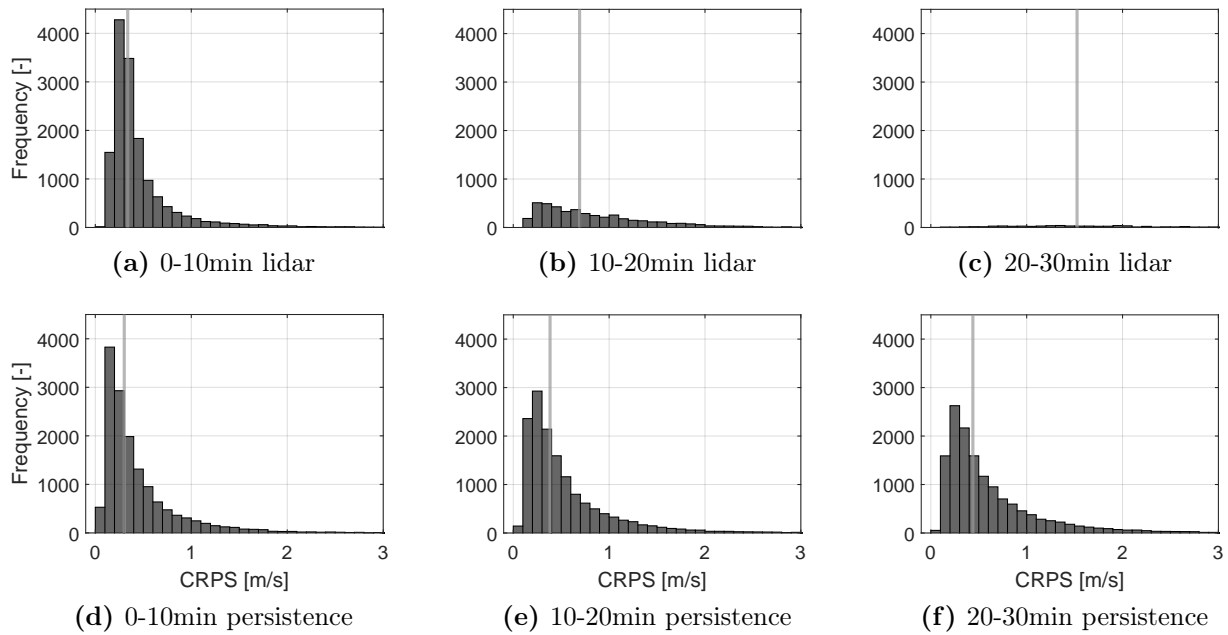


Figure 6.22: Histograms of CRPS of wind speed forecast for lidar (top) and persistence (bottom) forecasts from the offshore campaign for different forecasting horizons. Median of data is marked as grey line. Note: x-axes are limited to 3 m/s and do not comprise all outliers.




Figure 6.23: Histograms of CRPS of power forecast for lidar (top) and persistence (bottom) forecasts from the offshore campaign for different forecasting horizons. CRPS given in percentage of rated power. Median of data is marked as grey line. Note: x-axes are limited to 40% and do not comprise all outliers.




Figure 6.24: Example of Cumulative Distribution Functions (CDF) for three different forecast horizons. Left: wind speed forecast. Right: corresponding power forecast. Forecast horizons: 0-10min (black), 10-20min (dark grey) and 20-30min (light grey). Observation of 20-30min horizon is given as dashed vertical line.

To conclude, the results are in line with the findings of the calibration analysis: the lidar-based forecast method is working well for forecasts in the horizon up to 10 minutes. In fact, compared to the persistence method, the forecasts show less spread in the CRPS values. For forecast horizons above 10 minutes, however persistence performs better.

To confirm the conclusion of the calibration analysis, that the lidar-based forecast above 10 minutes are too low, because only small wind speeds arrive with such a delay, the deterministic forecast error is analysed as well. The advantage of the deterministic error compared to the CRPS is that the sign of the error indicates an over- or under-prediction. To calculate the error, all forecasted wind speed and power vectors in each forecast horizon bin were averaged, and the error between the average and the observation was calculated. Indeed, for forecasts over 10 minutes, the deterministic wind speed and power error of the lidar-based method is negative and underestimates the wind speed and power (Figure 6.25). It can be concluded that the lidar-based method as applied in this thesis, with a limited measurement range, is not suitable for forecast horizons over 10 minutes. All further analysis will concentrate on the forecast horizon of 0-10min.

Apart from the general analysis different forecast horizons and the comparison of the lidar-based forecast with persistence, the forecast accuracy is analysed for different parameter in the following section.




Figure 6.25: Histograms of error of deterministic lidar-based wind speed (top) and power forecast (bottom) of the offshore campaign for different forecasting horizons. Median of data is marked as grey line.

Influence of the wind direction

The results up to now were analysed for periods where the AV04 turbine was not affected by the wakes of other turbines. These wakes have an influence on the CRPS (Figure 6.26). The CRPS from wind speed and power forecasts increase significantly for wind directions from 0° to 180° where the AV04 is affected by the wakes of its neighbors. In the power forecast, this effect is more pronounced.

The lidar-based forecast model as presented in this thesis does not consider the decrease of wind speed due to the energy extraction of each turbine in the wind farm. In future forecast models, a wake model should be implemented and then forecasts for wind directions where the turbine is affected by wakes should be possible [52]. Wake modelling was out of scope for this thesis.




Figure 6.26: Polar plots of CRPS of the 0-10min forecast horizon of the AV04 depending on wind direction for wind speed (left) and power forecast (right). CRPS of wind speed forecast is given in m/s; CRPS of power forecast is normalised with rated power and given in %.

Influence of the wind speed

The CRPS of the wind speed forecast increases for increasing wind speeds (Figure 6.27). This effect is also noticeable and even more pronounced in the power forecast, but only up to rated wind speed. In the partial load region of the power curve, an error in the wind speed forecast leads to an error of the power forecast proportional to the third power. Above rated, the power is constant and therefore wind speed prediction errors do not affect the power prediction.

The reason for the increase in forecast error for increasing wind speeds can be found in the lidar measurement setup. The lidar is supposed to measure on a horizontal plane in front of the turbine. As wind speed increases, the thrust acting on the turbine increases as well until rated wind speed is reached, and the turbine structure tilts backwards. This tilt angle γ is small and in the range up to 1° on the nacelle, but it also affects the lidar measurements. The lidar tilts backwards along with the nacelle and therefore the measurement height is not constant anymore, but increases with the measurement range (Figure 6.28). The height deviation Δh is calculated as follows:

$$\Delta h = \sin(\gamma) \cdot \text{range}. \quad (6.4)$$

To analyse the influence of the measurement height deviation on the measured wind speed, a turbine tilt angle of 1° was assumed and the wind speed deviation Δv between the actual measurement height and the intended measurement height (hub height) was calculated using




Figure 6.27: CRPS of the offshore forecast depending on the turbine wind speed for the 0-10min forecast horizon.




Figure 6.28: Increase of measurement height against the measurement range of the lidar due to a turbine tilt of 1° .

the wind profile power law relationship:

$$\Delta v = v_{Hub} - v_{Lidar} = v_{Lidar} \left(\frac{h_{Hub}}{h_{Lidar}} \right)^\alpha - v_{Lidar} \quad (6.5)$$

where v_{Lidar} is the horizontal wind speed measured by the lidar, v_{Hub} the horizontal wind speed at hub height, and h_{Hub} and h_{Lidar} the respective measurement heights. The wind speed deviation was calculated for three different the power law coefficients α and three different lidar wind speeds (Figure 6.29).

For lower wind speeds, the wind speed deviation is below 1 m s^{-1} . For wind speeds of 18 m s^{-1} the wind speed deviation is around 2 m s^{-1} for a typical measurement range of 6 km and a strong shear. This deviation influences the forecasted wind speed and thus the predicted power.

To correct the measured wind speed deviation resulting from the turbine tilt, the turbine tilt would have to be known. With this information, the lidar trajectory could either be corrected online during the measurement campaign, so it measures horizontally in spite of the turbine tilt. Another approach would be to use the tilt information to correct the wind speed in post-processing using assumptions such as a certain power law coefficient, as shown above.

For this thesis, the information about the turbine tilt is not available, and therefore it is ignored. It is recommended however, to investigate the matter thoroughly in future research projects.




Figure 6.29: Deviation of wind speed between lidar measurement and hub height wind speed for different lidar wind speeds and power law coefficients α . The calculation is based on a turbine tilt angle of 1 deg.

Influence of the measurement range

All available wind speeds from all lidar measurement range gates were taken into account up to now, to calculate the wind speed and power forecast. To analyse, if the measurement range that the wind speed vectors are measured at has an influence on the forecast accuracy, the measurement range of 10 km was split into 10 bins of 1 km range, and then forecasts were calculated with wind speed vectors only from the respective bin. Then the CRPS was analysed (Figure 6.30).

The measurement range does not have a significant influence on the forecast accuracy. Forecasts with wind speeds measured up to 2 km in front of the turbine have a slightly smaller CRPS than forecasts with wind speeds measured up to 5 km.

The CRPS deteriorates for forecasts where only wind speeds from a measurement range of >5 km were taken into account. However, for these ranges, also the number of available measurements is low.

An interpretation of this CRPS analysis is that the Taylor-based propagation model works well up to 5 km. Below that range, there is no optimal measurement range where wind speeds for power forecasts should be measured to gain the highest forecast accuracy. Above that range, there are only few measured wind speed vectors available and those vectors produce forecasts with low accuracy.


Influence of the number of wind speed vectors

Depending on the measurement range of the lidar, a different number of wind speed vectors are available for the forecast. To analyse the effect of the number of wind speed vectors on the forecast accuracy, the number was recorded for each 0-10min forecast and compared with the corresponding CRPS of the wind speed and power forecast (Figure 6.31).


The effect of the number of the available wind speed vectors is clear for the wind speed forecast: the more vectors are available, the smaller the CRPS and the better the forecast. For high numbers of available vectors of around 100 the spread of CRPS of the forecasts is lower and there are fewer outliers.

The effect on the power forecast is not so clear: for a low availability of less than 40 wind speed vectors, the power forecasts tend to have a high CRPS. The average power CRPS is lowest for an available number of vectors of around 60. The CRPS average increases again if more vectors are available, but for 200 and more wind speed vectors the CRPS decreases again. Also the spread of CRPS decreases significantly.

The conclusion for this analysis is that forecasts should not be carried out if there are less



(a) Wind speed forecast



(b) Power forecast

Figure 6.30: CRPS of the offshore forecast for different measurement ranges that were taken into account to calculate the forecast.

than 40 wind speed vectors available. Forecasts for these cases tend to have a low accuracy. Another conclusion is that the measurement range should be as high as possible: the more wind speed vectors are available for the forecast, the better is the forecast accuracy in general and the less spread in CRPS is to be expected.

Influence of atmospheric stability

Atmospheric stability is a measure of the thermal stratification of the atmosphere and determines the fluctuation of the air. Stability describes whether air parcels will rise, sink or stay neutral and can be classified as unstable, stable or neutral [53]. The stability has a strong influence on the wind shear: unstable conditions enhance vertical mixing within the boundary layer and therefore the wind profile does not change as much with increasing height compared to neutral or stable conditions. Stable conditions lead to a strong shear and the wind speed can change significantly with the height [54].

In this section it is analysed whether the stability influences the forecast accuracy. The Speed Richardson number Ri_S is used as a measure of stability [55]:

$$Ri_S = \frac{g}{\bar{\Theta}_v} \frac{d\bar{\Theta}_v/dz}{(\bar{d}U/dz)^2}, \quad (6.6)$$


where g is the gravitational constant, $\bar{\Theta}_v$ an average potential temperature, $d\bar{\Theta}_v$ the potential temperature, $\bar{d}U$ the wind speed difference between two heights and dz the vertical distance between these heights.

The Speed Richardson number was calculated with FINO1 data with temperature and humidity from 101 m and 34 m height, pressure data from 92 m and 21 m height, and wind speed data from 91 m and 34 m height. Stability was classified using the ranges in Table 6.7:


Stability classes	Ri_S
Unstable	<-0.01
Neutral	$ Ri_S \leq 0.01$
Stable	>0.01

Table 6.7: Stability classes using Ri_S

The analysis shows that the stability does not have a significant influence on the forecast quality (Figure 6.3.4). The CRPS for stable and unstable conditions is similar for the wind speed and power forecasts. Only for neutral conditions, the CRPS is increased significantly.



(a) Wind speed forecast



(b) Power forecast

Figure 6.31: CRPS of the offshore forecast depending on the number of measured wind speeds vectors that were available for the forecast. Data is binned and for each bin an average CRPS is calculated (grey marker).




Figure 6.32: CRPS of the wind speed (left) and power forecast (right) depending on stability classes.

However that there are only 412 forecasts available for the stable stability class, whereas for unstable conditions there are 2072, and for stable conditions there are 2851 forecasts available. The increase in CRPS for neutral conditions could be a biased, as not enough forecasts are available to compare them to the stable and unstable conditions.

6.4 Impact of wind ramps on the forecast accuracy

A wind ramp is a fast and strong increase or decrease of wind power generation of a wind turbine or wind farm. An increase in power is considered an upward ramp, and a decrease in power is a downward ramp [14]. Ramp events can be critical to electrical grid stability and therefore the forecast of such events has become more important.

An example of a critical scenario is a severe drop in power generation of a large offshore wind farms. This drop in power could lead to a restriction of the demand in the electrical grid. A wind ramp that is not forecasted well and affects a large cluster of offshore wind farms, can also lead to a significant impact on the cost of balancing the power system on a national level, as a case study from the Thames Estuary has shown [56].

Ramps are rare events that result from a variety of underlying meteorological conditions. There are two main atmospheric processes that can cause ramps [57]: large horizontal processes in the atmosphere, such as weather systems (fronts) or mesoscale circulations (e.g. sea-breeze or mountain valley winds), can affect whole regions and several wind farms. Local vertical processes caused by convection - such as thunderstorms or low level jets - are short-cycle phenomena that are sensitive to local conditions and therefore harder to predict.

Another reason for a (downward) ramp is the exceeding of the turbine's cut-off wind speed. When wind speeds is above cut-off wind speed, turbines shut down and stop producing power. Typically this is at around 25 m s^{-1} and is intended to protect the structure from extreme loads. This type of event is very difficult to forecast, as small errors in wind speed forecast lead to a big error in the power forecast.

A ramp is defined by the following parameter [14]:

- Magnitude ΔP_r : the variation of power
- Duration Δt_r : the time during which the variation of power occurs
- Ramp rate $\Delta P_r / \Delta t_r$: a metric to quantify the intensity of a ramp
- Direction: states whether it is an up-ramp or a down-ramp.

Up to now, there is no consistent definition of what the magnitude and duration of the variation in power should be for an event to be considered as a ramp. For Δt_r the definitions range from 10 min to 5 h. The associated ΔP_r are partially defined as absolute values, e.g. $\Delta P_r = 150 \text{ MW}$ or as relative deviations from the rated power of the considered wind turbine or wind farm. Extreme ramp events, are for example ramps with a magnitude of $\Delta P_r \geq 50 \%$ over a duration of 30 minutes, or ramps with a magnitude of $\Delta P_r \geq 70 \%$ over two hours (Figure 6.33).

In this thesis, all ramps with a duration of 30 minutes and longer and magnitudes of $\Delta P_r \geq 30 \%$ will be considered.

6.4.1 Detecting wind ramps

The wind ramps in this thesis are analysed for the offshore site alpha ventus, using the power data of the AV04 turbine. To detect wind ramps, the so called *minimum-maximum method* [58] is applied to the power data of the AV04¹. This method finds the maximum magnitude change of power ΔP within a moving window with the length Δt , where P_{min} and P_{max} are the minimum and maximum power value within this window. If $\Delta P \geq \Delta P_r$, and the magnitude of the power change is greater or equal to the predefined ramp threshold ΔP_r , the event is considered a ramp. If multiple pairs of points meet the threshold criteria, only the shortest time Δt is used.

The ramp duration Δt_r is determined by the times t_{min} and t_{max} that correspond to P_{min} and P_{max} . If $t_{min} < t_{max}$ the event is an upward ramp, and vice versa.

¹This section is based on the Master thesis "Lidar-based Wind Ramp Forecasting in Complex Terrain" carried out by Svenja Ellinghaus, which was handed in at SWE in 2018 and supervised by the author of this PhD thesis.




Figure 6.33: Schematic diagram of a ramp matrix. The ramp duration is denoted WL, the ramp magnitude is denoted Δp_{RD} . Extreme ramps are in the top-left corner, and low-amplitude ramps of longer duration are in the bottom-right corner. Reproduced without modifications from [58] ©American Meteorological Society. Used with permission.

Up- and down ramps were counted at the AV04 from March 1st, 2020 to December 31, 2020. Four different time windows and five different power thresholds were used [58]. The ramp-count matrix (Figure 6.34) counts the cumulative numbers of ramps with magnitudes ΔP equal to and above this threshold. This means that a wind ramp of $\Delta P \geq 70\%$ over 1 h fulfills this threshold but also the threshold $\Delta P \geq 30\%$ over the same duration and will therefore be taken into account in the matrix.

The number of upward ramps and downward ramps are very similar for the AV04. Small ramps with a magnitude of $\Delta P \geq 30\%$ over 30 minutes occur frequently. Extreme ramps (up- or downward) with a magnitude of $\Delta P \geq 70\%$ over 30 minutes occur on average around twice per month. Figure 6.35 gives an example for both cases.




Figure 6.34: Number of upward (left) and downward ramps (right) for different ramp durations Δt and ramp magnitudes ΔP at the AV04 turbine. Axis limits and color scale are the same on both plots.




Figure 6.35: Timeline of power generation of the AV04 turbine with wind ramps marked with a magnitude of $\Delta P \geq 30\%$ over 30 minutes (top) and wind ramps marked with a magnitude of $\Delta P \geq 70\%$ over 30 minutes (bottom).

6.4.2 Assessing the ramp impact on the forecast accuracy

Since there were several ramp events occurring during the offshore forecast campaign, the question is what impact the ramps have on the forecast accuracy. It is assessed, whether the lidar-based forecast method is able to forecast the variation in power and whether it is able to forecast the variation more precisely than the persistence method.

To that end, the probabilistic error metric CRPS is compared for 0-10min forecasts of the AV04 of periods where a ramp event was taking place and periods where no ramps event was taking place (Figure 6.36). All periods were marked as ramp events that were found in the previous section. This analysis is only carried out for the offshore site, as the availability of forecasts is significantly higher, and therefore the statistical analysis is more meaningful.

For periods where no ramps occur, the lidar-based forecast and persistence have almost the same accuracy. However, the error of the lidar-based forecast is significantly smaller than the error of the persistence forecast during ramp events, both for the wind speed and the power forecast. During the ramp events, the lidar-based forecast demonstrates its strength: using the preview information from the wind speed measured a few kilometers in front of the turbine, it is able to forecast the changes of power. Persistence only uses old information from the past and therefore cannot forecast large changes in the future.

6.4.3 Case study of a failed ramp forecast in the onshore campaign

In the previous section the impact of ramp events on the forecast was statistically analysed for the offshore campaign. The same analysis is not conducted for the onshore campaign, because in the analysis process it was found that the lidar-based forecast failed in several cases to forecast ramp events in Stötten. In order to determine the cause, one ramp event is analysed in this section as a case study.

In July 21, 2016 the measured wind speed at the onshore turbine increased within 20 minutes from 4.74 m s^{-1} at 02:40 to 13.73 m s^{-1} at 03:00 (Figure 6.37). Accordingly an upward wind ramp in the turbine power was detected with a magnitude of $\Delta P \geq 86\%$ over 20 minutes. The lidar-based forecast however was not able to detect this ramp and resulted in a CRPS error of 4.27 m s^{-1} for the wind speed forecast and a CRPS error of 39.8% for the power forecast at 03:00.

To analyse the cause, the measured lidar wind speed and the wind speed that is actually taken into account for the forecast is analysed (Figure 6.38). At 02:52 the wind direction was around 270° and the lidar measured straight into the wind direction. The maximum measurement range was around 6 km and the wind speeds that were taken into account for the forecast from





Figure 6.36: CRPS of offshore forecast for periods with and without ramp event taking place for the lidar-based and persistence forecast.




Figure 6.37: Timeline of a ramp event at the onshore site in Stöetten. The ramp is seen in the wind speed (top) and power data (second from top) of the turbine. The ramp event with a magnitude $\Delta P \geq 86\%$ over 20min is marked in the power data. Timeline of the lidar-based power forecast (third from top) and wind direction reconstructed from lidar data given for the same time period (bottom).

the lidar scan were low. With the next scan, the wind direction changed to 300° and the lidar started to measure high wind speeds in the distance at 02:54. At this point, the wind ramp started to pass the area, and the lidar-based forecast would have had the chance to forecast it. However, the measured wind speeds were not taken into account for the forecast, because the wind direction had changed while the scan direction stayed fix and the lidar did not measure into the wind direction anymore. In this configuration the measured wind vectors never actually arrived at the turbine.

In the following minutes, the wind ramp passed through the area quickly, and the lidar measured high wind speeds. The forecast however, continued to use old information with low wind speeds vectors from the time, when scan direction and wind direction were aligned. The result was that the wind ramp was not forecasted, and the error in the forecast increased.

Two lessons can be learned from this case study. First, it is crucial that the measurement direction and the wind direction are aligned. If the lidar does not measure into the wind, it misses the information for the forecast, and wind ramps cannot be forecasted. Therefore, it is advised to mount the lidar on top of the nacelle, so it rotates automatically along with the wind direction. If the lidar is installed on a fixed platform, the other option is to automatically adjust the scan direction of the lidar according to the measured wind direction. The downside of this second option is that for the scan direction adjustment, old information from a last scan has to be used (which might not be valid anymore).

The second lesson to be learned from this example is that the lidar scan trajectory should be kept as narrow as possible. In the case of the campaign in Stötten, the range of azimuth angle of the scan trajectory was set to 40° . However, - at the cost of additional scan time - only a fraction of the measured information was used for the forecast. Since time is crucial to forecast wind ramps, the scan trajectory and its duration should be optimized, in order to measure only the necessary wind speed information as quickly as possible.

6.5 Lessons learned

Inflow measurements from lidars contain preview information of the stochastic variation of the wind field, which the wind turbine experiences after a forecast is issued. With the help of the turbine's power curve, the power forecast is obtained from the wind speed forecast. The variation of wind speed results in a variation of forecasted power. Therefore probabilistic forecast methods are necessary to quantify the uncertainty in the forecast for a specific forecast horizon.

Probability Density Functions (PDF) and Cumulative Distribution Functions (CDF) quantify the forecasted wind speed and power variation for each forecast horizon. Using these tools,




Figure 6.38: Scan plots of wind speed propagation from four consecutive scans of measurements during a ramp event at the onshore site in Stötten. Wind direction changes with each scan. Wind speed vectors in grey miss the turbine (marked red) and fade out with time, wind speed vectors in blue are used for the power forecast and do not fade out.

confidence intervals can be calculated, that indicate the likelihood of a wind speed or power forecast value.

To evaluate probabilistic forecasts, the calibration of the forecast needs to be checked and the forecast error needs to be quantified, by comparing the forecast to the actual wind speed or power measurement of the turbine. Appropriate tools need to be chosen, such as the Probability Integral Transform (PIT) histogram for the calibration, or the Continuous Ranked Probability Score (CRPS) for the forecast error.

It should be ensured that the wind speed and power measurement which is used for reference is well calibrated and free of outliers. Good quality observational data is essential to evaluate a forecast.

To analyse a new forecast method, the forecast horizon should be chosen according to the time resolution of the available observation from the wind turbine, which is used for the evaluation of the forecast.

The results of the calibration of the onshore and offshore campaign show that forecast horizons above 10 minutes result in biased lidar-based wind speed and power forecasts. For these forecast horizons, only small wind speeds are used which were measured in farther distances and therefore take longer to reach the turbine. Therefore the forecasts underestimate the actual wind speed and power. A forecast of more than 10 minutes with the lidar-based method proposed in this thesis, and with the limited measurement range of the long-range lidar used, is therefore not advised.

The availability of the forecast for the nacelle-based lidar setup offshore is significantly higher, than for the fixed setup onshore. A forecast was not possible only in 1.8% of the cases when a forecast could have been carried out (because the turbine was running and in power production). As the lidar was rotating along with the wind direction, a forecast was not possible only when the availability of the lidar signal, e.g. due to fog, was too low. At the onshore campaign, the forecast was not carried out, although it could have been, in 28.1% of the cases. This high number is due to the fact that a forecast was only carried out when the scan direction matched the wind direction. It is concluded that it is necessary to install a lidar on top of the nacelle for minute-scale forecasting, or to automatically adjust the scan direction according to the wind direction.

The forecast error of the lidar-based forecast in the first forecast horizon up to 10 minutes of the onshore campaign is higher than the error of the benchmark forecast method persistence. It can be concluded that the simple Taylor-based propagation model used in this thesis, does not reflect the actual propagation at the complex onshore site well enough, and the lidar-based forecast therefore has a higher error. To use lidar-measurements to forecast in complex terrain,

a more sophisticated wind evolution model will be necessary.

The forecast error of the lidar-based forecast in the first forecast horizon up to 10 minutes of the offshore campaign is lower than the error of the benchmark forecast method persistence. Only data where the wind turbine was not affected by the wake of the surrounding turbines was taken into account for this analysis. It can be concluded that lidar-based minute-scale forecast holds value over the benchmark persistence for the first 10 minutes forecast horizon.

Analysing the error of the lidar-based forecast further shows that the error depends on the wind speed. This is most likely related to the tilt of the turbine and the wind shear. Firstly, as wind speeds increase the thrust acting on the turbine increases and so the turbine tilts back. At rated wind speeds the thrust reaches its peak. As a result the lidar does not measure horizontally, but with an upward tilt. Therefore the measurement height of the lidar increases with the measurement range and is highest at rated wind speed. The actual measured wind speed then deviates from the hub-height wind speed. Depending on the wind shear, this leads to a measured wind speed that is higher than the hub height wind speed. Secondly, wind shear usually increases as wind speed increases [59]. This means that the effect of small height differences leads to bigger deviations in lidar derived wind speed at higher wind speeds, than at lower wind speeds. For both reasons the wind speed forecast overestimates the actual wind speed for higher wind speeds. It is noted however, that this does not necessarily affect the power forecast, as the error is most significant for above rated wind speeds, when the power is constant rated power.

The analysis of the lidar-based offshore forecast during wind ramps shows that the power variations are correctly predicted. Compared to the benchmark persistence, the error of the lidar-based forecast is significantly lower during these events (maximum CRPS of 11 % of rated power for the lidar forecast versus 17 % for persistence). During ramp events, the lidar-based forecast demonstrates its strength: using preview information from the wind speed measured a few kilometers in front of the turbine, it is able to forecast changes of power. Persistence on the other hand only uses old information from the past and therefore cannot forecast any future changes. It can be concluded that the added value of using a lidar for minute-scale forecast lies in forecasting significant changes of power. As wind ramps are potentially critical to the grid stability, or can affect the cost of balancing the power system if they are not forecast well, using a lidar to improve the forecast can be beneficial and valuable.

At the onshore site, it was noted that the lidar failed several times to forecast a wind ramp. This is due to the fact that the lidar measured the increase in wind speed, but the wind vectors never arrived at the turbine because scan direction and wind direction were not aligned. Again, it is concluded, that the scan direction should be optimized and aligned with the wind direction at all times. Also the opening angle of the scan trajectory should be as narrow as possible, and

only necessary wind speed information should be measured, to reduce the scan duration - as time is critical to forecast wind ramps.

Challenges for the implementation of lidar-based minute-scale forecasting

In the previous chapters, the forecasting chain for minute-scale forecasting using lidar data was introduced and results from two measurement campaigns were discussed. It was shown that the forecast accuracy can be increased - especially for wind ramp events - when using the preview information that a lidar provides of the inflowing wind field.

In this chapter, the challenges for the implementation of lidar-based minute-scale forecasting are discussed. One challenge is the availability of the measurement itself. Therefore the influence of the measurement range on the forecasts and the factors impacting the measurement range are analysed in Section 7.1. Finally in Section 7.2 other barriers to the adoption of lidar-based forecasting and possible solutions are explored.

7.1 How far do we see?

The measurement range of a lidar is not constant but varies because it depends on the content of aerosols in the air. The aerosols are responsible for the backscattering of the laser light, which in turn allow the lidar to measure the speed of the aerosols. If there is too much, or too little aerosol content, the lidar does not measure the correct wind speeds and is “blind”.

The challenges for minute-scale forecasting are a) a reliable forecast should be available at

all times and b) that the measurement range influences the forecast horizon. Therefore, this Section seeks to answer to the question: how far do we actually see with the long-range lidar used in the measurement campaign and what is the influence of the measurement range on the forecast horizon.

7.1.1 Correlating the measurement range and the forecasting horizon

The possible time horizon for short term forecasts depends on two factors. One is the wind speed itself - at very high wind speeds the wind field moves faster and the forecast time is correspondingly shorter. Another factor influencing the forecast horizon is the measurement range. If, for example, the wind can actually be measured at a distance of 10 km, the forecast time is twice as long as when measured at a distance of 5 km, provided that the parameters remain otherwise constant. The lidar measurement range is, due to its measurement principle, strongly dependent on the prevailing environmental conditions [60] and prone to fluctuations. In unfavourable conditions, the measuring range is strongly limited. As a result, the forecast horizon of the wind speed or power forecast is also influenced by environmental conditions and not constant.

Figure 7.1 shows the relationship between the measurement range, the wind speed and the forecasting horizon based on Taylor's frozen turbulence hypothesis (cf. Chapter 5.1). For a 10-minute-ahead forecast at 5 m s^{-1} , the measurement range needs to be over 3 km. A 20-minutes-ahead forecast is only possible for wind speed less than 8 m s^{-1} , because only low wind speeds take that long to reach the turbine from a distance of 10 km. And if a critical shut down of the wind turbine due to wind speeds above the cut-out wind speeds of 25 m s^{-1} should be forecast with a forecast horizon of 5 minutes, the measurement range needs to be at least 7 km.

7.1.2 Analysing the measurement range

The goal of this analysis is to investigate the possible range of the lidar measurement depending on different factors. Data from the onshore site in Stötten was used for this analysis measurement. In particular, the effects of environmental influences on the achieved measuring ranges were investigated. Met mast data was used to correlate environmental conditions to the lidar measuring range. The met mast was located in the immediate vicinity of the lidar measurement and could therefore provide data on meteorological parameters at the same time (c.f. Chapter 4.1). The data evaluation is done with the help of statistical methods.

For the statistical evaluation of the measurement range, data from different measurement trajectories were analysed (c.f. Table 4.1 in Chapter 4.1). For each trajectory, the number




Figure 7.1: Forecast horizon calculated based on Taylor for different wind speeds and measurement ranges; the horizon is given in minutes. The wind speed ins shown only for wind turbine operational range. Reproduced without modifications from Würth et al. [16] with permission.

of range gates is given in the data sheet. This number indicates how many parts the measuring range is divided into. Together with the length of the range gates the range for each measurement point can be calculated.

The first question is, what is the maximum measurement range per ray. The measurement data that remain after filtering do not allow a clear distinction of the maximum range, as valid and invalid data are mixed in farther measurement distances. However, in order to determine the maximum measurement range, the limit of valid data for each measurement ray must be known. Therefore, a method had to be developed to determine the range of a measurement. To determine the range for each individual beam, the methods *Sum Range* and *Weighted Range* were developed¹. With the *Sum Range* method, the measurement distances where valid data are available are summed up for each beam.

$$\text{Range}_{\text{Sum Range}} = \sum \text{Measurement distance of valid data.} \quad (7.1)$$

The second method *Weighted Range* takes into account the gaps of data along a ray and weights sections with adjacent areas of valid measurement data, so called blocks, higher compared to sections with a lot of data gaps. This means the number of valid data as well as the length of a block with valid data is taken into account:

¹The methods were developed by Alex Brenner in a study thesis supervised by the author of this document [61]. The methods were also discussed in a paper, written by the author of this thesis and co-authored by Alex Brenner [62]




Figure 7.2: Schematic visualization of the results of two methods to determine the maximum range of a lidar measurement of an example data set of 10 range gates (first row) and maximum measurement range (second row). Filled points mark valid data. Reproduced without modifications from Würth et al. [62] with permission.

$$\text{Range}_{\text{Weighted Range}} = \frac{\sum(\text{block end} \cdot \text{block length})}{\sum \text{block length}}. \quad (7.2)$$

The end of the block represents the maximum range of the measurement within a block of valid data. The block length determined by the product between the number of range gates and the range gate length within a data block.

The results of both methods are schematically visualized in Figure 7.2. The comparison shows that the gaps in the data lead to a lower range using the weighted range method.

In order to decide which of the range determination methods should be used, the results were compared for different data sets. Figure 7.3 shows a data set where the range varies greatly over a period of several hours. Comparing the calculated ranges of the Sum Range and Weighted Range methods, it is noticeable that for periods of time with a greater range, the two methods match well. However, for periods of time in which the lidar device measures almost no valid data, the Weighted Range method still calculates a range of several hundred meters. The reason is that measurement data that have been erroneously marked as valid by the filter algorithm are more important for this method. If there are few data available, the calculated range is therefore erroneously extended by these outliers for the Weighted Range method. One can conclude that the Weighted Range method does not work robustly and reliably, especially for periods of time with a very short range. For this reason, further investigations into the influence on the range were carried out using the Sum Range method.

Once it was clear which method could be used to determine the range, research was carried out to determine which factors influence the range. First the maximum measurement range of the onshore data set was calculated (Figure 7.4). During the campaign, the measurement range varied with time, alternating between ranges close to 0 m and ranges up to almost 10,000 m. Three sources were identified and investigated that could cause a variation in the measurement range:




Figure 7.3: Comparison of the Sum Range and Weighted Range method applied to a filtered data set. Reproduced without modifications from Würth et al. [62] with permission.




Figure 7.4: Maximum measurement range over time for the onshore measurement campaign in Stötten averaged over 10 minutes.

1. Data processing method. During the analysis of the data, it was found that the choice of filter algorithm significantly impacts the amount of valid data in farther ranges, and therefore has an impact on the measurement range.
2. Lidar measurement setting. When starting a lidar measurement, the number of pulses emitted per ray must be chosen by the user. This setting influences the quality of the backscattered signal, and therefore also the measurement range.
3. Environmental conditions. The lidar measurement principle is based on the backscatter of the laser pulses from aerosols. Therefore the quality of the signal (and the measurement range), depend on the amount of these aerosols in the air. The amount of aerosols in turn depend on environmental conditions at the measurement site.

The influence of each of these factors on the measurement range is discussed in the following.

Filter algorithm

As described in chapter 4.2.1, the measured lidar data must be checked for invalid data points. Depending on the filter algorithm, different data will be marked as valid or invalid. A very conservative algorithm filters out all invalid data but also marks valid data as invalid. This in turn has an effect on the measurement range, since data is lost, especially at greater measurement distances (Figure 4.8). A robust filter algorithm that detects only invalid data is therefore essential for achieving the highest possible measurement range. In this thesis, the so-called “range filter” was developed and used for this purpose.

Number of pulses

The number of pulses emitted by the lidar device is an important, device-specific factor influencing the measuring range. The number of pulses determines how many pulses are emitted which are then reflected by aerosols. The radial velocity is then determined from the backscatter. The more pulses are emitted, the more laser energy is emitted and the higher the probability of a good backscatter signal. The number of pulses can be selected by the user. During the measurement campaign in Stötten four different configurations were tested. The number of pulses changed from 10,000 to 60,000 and each trajectory was executed over several weeks.

Figure 7.5 shows the progression of the proportions of valid data over the measurement range for different pulse counts. All curves show that the proportion of valid data decreases with increasing range, with the largest decrease being observed at about 4500 m. At higher ranges, the proportion of valid data for all pulse numbers approaches zero. In this graph the data of the first three range gates show no physically meaningful values. The reason for this is the length of the emitted pulse. This pulse cannot be infinitely short and so the backscatter




Figure 7.5: Comparison of the valid measurement points over the measurement range for different number of pulses. The legend also indicates the number of days the measurements were conducted.

signal from the first three range gates would reach the lidar before the pulse is completely transmitted.

If one looks at the figure with regard to the different numbers of pulses, it is noticeable that from 10,000 to 40,000 pulses the measurement distance with the highest loss of data is shifted to higher ranges. Thus it could be assumed that the range improves on average with the number of pulses. This can be explained by the fact that a higher number of pulses increases the probability of encountering aerosols which generate a detectable backscatter signal. However, increasing the number further to 45,000 pulses does not lead to a higher measurement range. An explanation might be found if one compares the amount of days during which the respective measurements were carried out (c.f. Table 7.1). The scan with 45,000 pulses was running four times as long as the scan with 40,000 pulses, which leads to the conclusion that apart from the data filter, and number of pulses, also the season and environmental conditions influence the measurement range. This conclusion is confirmed, as increasing the number of pulses to 60,000 does not increase the range significantly. It should be noted that in order to analyse the impact of the number of pulses independently of other influences, a measurement campaign should be carried out with several lidars measuring simultaneously with different numbers of pulses. However, this was outside of the scope of this thesis.

There is high variability in the measurement range, independent of the number of pulses (Figure 7.6). Also, it is clearly shown that the average measurement range which can be achieved does not come close to the measurement range of 10,000 m from the product specification. In

ID	Period	No. pulses	No. days
1	28.04.2016 00:00 26.05.2016 23:40	10,000	29
2	07.06.2016 00:00 02.09.2016 12:20	30,000	87.5
3	02.09.2016 14:00 26.09.2016 08:20	40,000	23.8
5/6	05.09.2016 08:20 25.04.2017 23:50	60,000	120.1
7	26.04.2017 14:20 01.08.2017 00:00	45,000	96

Table 7.1: Overview of measurement periods with different number of pulses, cf. also Table 4.1.

fact, depending of the number of pulses, the median of the measurement range is close to 5000 m. Only rarely and only with high number of pulses of 60,000 the measurement range extends to 10,000 m. The histograms in Figure 7.7 confirms this and also show that frequently the measurement range is close to 0 m for all number of pulses. This bi-modal distribution of the range with two peaks is typical for long-range lidar measurements. Either they measure several kilometers, or they measure close to 0 m. This variability of measurement cannot be explained with the lidar settings but is caused by external factors.

To conclude, an increase of the number of pulses does not necessarily improve the measuring range proportionally. A doubled number of pulses however doubles the time required for the respective measurement. During this time the wind field within a scan can change and it becomes difficult to get a coherent picture of the wind speeds. The choice of the number of pulses should therefore be made according to the intended use of the measurements or it could be made adaptive according to wind speed and environmental conditions.

Environmental conditions

Aerosol content and composition have a decisive influence on the measuring range of the lidar measurement. The aerosols in turn are changed by environmental influences [63]. For this reason, the correlations between environmental influences and the resulting measuring range of the lidar measuring device have been analysed. The measurement data of the environmental parameters for this purpose were obtained from the measuring mast which recorded data simultaneously to the lidar measurements at the beginning of the measurement campaign.

Figure 7.8 shows an overview of the correlation between the maximum measurement range




Figure 7.6: Boxplot of maximum measurement range for different number of pulses. The central line indicates the median, the bottom and top edges of box indicate 25th/75th percentiles, whiskers extend to extreme data points not considered outliers. Points are considered outliers if they are greater than $q_3 + 1.5 \cdot (q_3 - q_1)$ or less than $q_1 - 1.5 \cdot (q_3 - q_1)$. Outliers are plotted individually.




Figure 7.7: Histograms of maximum range for different number of pulses. Median of data marked as grey line.

and several environmental variables of the site: the relative humidity, rain, air temperature, atmospheric pressure, air density, wind speed, wind direction and turbulence intensity. In Figure 7.9 the correlation between humidity and rain is given. For the figures only scans with 30,000 pulses, measured in summer conditions, were considered and the data were averaged over 10 min.

The relative humidity is an important indicator for fog or cloud formation. Figure 7.8a shows the relative humidity from 5 m height in correlation with the measurement range of the lidar. High humidity was frequently observed, when the lidar range was close to zero. When lidar measured only up to 2000 m, the observed humidity was at least 85%. When the lidar measurement range was higher than 2000 m, there is no correlation observed between the range and the humidity.

To investigate which environmental variable has the strongest correlation to the measurement range, the Maximum Information Coefficient (MIC) between the maximum range and the environmental variables is calculated (Figure 7.10). MIC is a tool to measure the strength of the linear or non-linear association between two variables X and Y. It belongs to the maximal information-based non parametric exploration (MINE) class of statistics [64]. Similar to the coefficient of determination R^2 for linear regressions, the MIC score takes values between [0 1], with 1 showing a complete dependency between the two variables and vice versa for 0. While all investigated variables do not have a very strong dependency with the range with MICs below 0.3, humidity scores the highest, followed by temperature, pressure and wind direction. Wind speed and turbulence intensity do not have a big impact on the measurement range.

With the help of a detailed observation for one day on June, 8th 2016 the correlation between humidity and measurement range is investigated further. Figure 7.11 shows the trend of the measurement range together with the trend of the relative humidity. The selected day shows large fluctuations in the range of over 5000 m. In the morning, the range increases to over 6000 m while the humidity decreases. A first sharp decline of the range can then be observed from 11:00 on. The range decreases from above 5000 m to about 3000 m. At the same time the humidity begins to rise. A further decrease in range at 14:00 is accompanied by an increase in relative humidity. The measurement range drops to few hundreds of meters. For the rest of the day the range remains very low and the humidity high, only interrupted by a short period of higher range at 18:00.

In the time periods with very high humidity it is assumed that fog or precipitation occurs and thus the range is greatly reduced. To show the correlation between range and fog, the same time period of one day is chosen and the range is compared with webcam images from the measurement site (Figure 7.12). When cross checking significant points during the day with high or low measurement range with the cloud formation at the site, it becomes clear that the




Figure 7.8: Maximum measurement range versus different environmental conditions. Data are averaged over 10 minutes for measurement period from June 7th - August 8th 2016 where lidar measured with 30,000 pulses.




Figure 7.9: Correlation of humidity and rain sensor data. Data are averaged over 10 minutes for measurement period from June 7th - August 8th 2016 where lidar measured with 30,000 pulses.




Figure 7.10: Maximum information coefficient (MIC) between maximum range and different environmental variables.




Figure 7.11: Maximum measurement range (black line) on June, 8th 2016 with humidity (grey line) and rain (blue bars).




Figure 7.12: Webcam pictures of the measurement site on June 8th, 2016.

maximum measurement range is directly related to the visibility. At 09:00 in the morning, the high visibility leads to a high measurement range. Due to the cloud front that moves through at noon, the range drops. When rain starts around 15:00 in the afternoon, the measurement range collapses further. Only when it brightens up at around 18:00 in the evening, the range increases again.

In summary, it can be said that the range of the lidar measurement, and thus also the possible prediction horizon, is strongly dependent on the humidity, the number of pulses and also the filter used. Although the investigations in this sections were conducted only for the onshore campaign, this conclusion is also valid for offshore measurements. The physical measurement principle is the same offshore, but other environmental factors, such as the amount of salt particles in the air might influence the measurement range. However, this investigation of offshore conditions was outside of the scope of this thesis.

For the location on the Swabian Alb the average measurement range of the lidar was around 5 km and with a pulse number of 40,000 pulses, with a high variability between 0 km and 8 km. The measurement range of 10 km documented in the product data sheet, was not reached.

7.2 Overcoming barriers to adoption with the help of the community

On June 12/13 2018 a collaborative workshop between the IEA Wind Tasks 32 and 36 on "Very short-term forecasting of wind power" took place in Roskilde, Denmark. It was organised by the author of this thesis organized together with colleagues from the University of Oldenburg and DTU². The goal of the workshop was to use the international platform provided by IEA Wind to bring together experts from different countries; to enable an exchange of experience; and to discuss the barriers and solutions to the adoption of minute-scale forecasting of wind power. A tangible outcome from the workshop was an overview paper of methods for minute-scale forecasting, led by the author of this thesis [12]. .

Several methods for minute-scale forecasting were discussed during the workshop and in the paper. Using preview data from lidar data was categorized as a remote-sensing-based propagation model method which brings unique advantages compared to historic time-series based models or NWP models.

Only remote sensing data offers high resolution information of the wind field which has not yet reached the turbine or wind farm and is therefore able to forecast not only trends, but future critical events.

The discussion also identified barriers which stood in the way of implementing lidar as a state-of-the art tool for minute scale forecasting.

- Availability of measurement: The measurement range fluctuates and measurements and thus forecasts are not always available.
- Reliability and pricing of lidar devices: The hardware can break and loose functionality and long-range lidars are very costly. A favourable cost-benefit ratio has not been proven.
- Need for common tools and standards: There are no common tools or standards available for measurement campaign setup, maintenance scheduling or data processing.

In this section, those barriers will be discussed in light of the results and insights of this thesis.

²Information including minutes and a link to a video recording of the workshop can be found at <https://community.ieawind.org/task32/events/event-information/workshop-09>

7.2.1 Availability of measurements

The measurement range of the lidar fluctuates due to device settings but also due to environmental conditions, which cannot be influenced. The measurement range is site specific and depends on (for example) rain showers or fog. Physics cannot be changed and the measurement principle dictates the measurement range and availability of data. The dilemma is that the measurement range also determines the forecast horizon, and online forecasts are not possible if the lidar cannot obtain data.

For this thesis, an analysis of the measurement range of a scanning-lidar was carried out for the first time (cf. Chapter 7.1.2). Thanks to the measurement period over several months, it was possible to thoroughly study the availability of data for different conditions. On average the StreamLine XR measured at the onshore site with a measurement range of around 5 km, which corresponds to a forecasting horizon ranging from 20 min for a wind speed of 4 m s^{-1} (cut-in) to 7 min for a wind speed of 12 m s^{-1} (rated) (Figure 7.1). Humidity was identified as the strongest influence on the measurement range and it was shown that the range corresponded with the visibility at the site. It was also shown that the range was frequently down to few hundred meters and thus power forecasts were not possible.

A solution to this dilemma that was discussed in the workshop is to find out what the optimum conditions for a good measurement range actually are. A thorough study in this area is lacking. A conclusion from this thesis is that it should be investigated what the conditions are that influence the visibility. Aerosol contents need to be studied, and it will be interesting to investigate, which sites (onshore/ offshore, complex/ non-complex) will come out as measurement range winners.

A second solution to the range dilemma that was discussed, is the need for “Plan B” - a fallback solution which is able to provide minute-scale forecasts in case the lidar does not deliver. Other sensors, e.g. radars or drones should be investigated as supplements. Also assimilating the measurement data into NWP models could make the forecasts more robust. A simple solution could also be to fall back on persistence.

7.2.2 Reliability and pricing

Just as every other electronic device, the lidar can break down and lose functionality. Scanning wind lidars have only been on the market for less than 10 years and especially new generations of lidar devices still need to be optimized for robustness. Since there are only a handful of manufacturers on the market, the competition has not led to a drop in prices of the devices. Prices of more than 200k Euros are standard. Therefore the idea to have a second lidar as a

backup, redundant system for minute-scale forecasting is financially unattractive.

To overcome the reliability and pricing barrier of lidar devices, the focus of the ongoing research should be on evaluating the added value of the lidar-based minute-scale forecasting. End users of the forecasts, such as wind farm operators, or transition system operators need to be included in future research projects, to find out what the cost-benefit balance of lidar-based forecasting is.

Another trend in research that could help overcome the pricing barrier, is conducted in the area of lidar hardware. The initiative OpenLidar strives towards an open source collaboration on wind lidar hardware and software [65]. Its goal is to encourage collaboration around wind lidar by developing a modular wind lidar architecture and providing a framework for cooperation. Modular lidars have the benefit, that malfunctioning parts can be easily exchanged, which drives down the cost of maintenance. Also production lines for manufacturing can be set up more cost effective, if the lidar modules are standardised in the future.

Another research activity that is of interest in the area of hardware development is the project ANWIND. In the BMWi³-funded ANWIND project, the University of Stuttgart and a lidar manufacturer are developing a robust and compact lidar scanner. Both initiatives, OpenLidar and ANWIND, show clearly the trend: the future devices will have a lower cost and at the same time increased robustness and reliability.

7.2.3 Need for standards and common tools

When planning the lidar measurement campaign for this work, there was no previous work or experience that could be built on which explained how to set up such a campaign for the application of minute-scale forecasting.

The campaigns that were carried out in this work are just a first attempt and there are more open questions about what an ideal campaign should look like. Possible questions are

- how many lidars are necessary,
- if they are best installed on the nacelle of a wind turbine, or on the transition piece,
- and also what an ideal measurement strategy should be: only horizontal measurements or mixing them up with vertical scans.

The solution to this lack of standards is to continue the work and set up measurement campaigns for different forecasting use cases, such as for onshore/offshore sites and small/very large wind farms. The workshops participants agreed that the next step should be the writing of

³Federal Ministry for Economic Affairs and Energy

Recommended Practices which gather the experience and knowledge from a range of measurement campaigns and experts. The IEA Wind community offers the perfect platform for such a document. These Recommended Practices could lead to a set of standards, that give guidance for the setup of the measurement campaigns, but also for practical matters such as lidar maintenance schedules or data quality criteria to be assessed.

Another barrier is the lack of a common tool set to process lidar data. This barrier is not unique to the application minute-scale forecasting but is relevant for all applications that use long-range scanning lidar data. These devices only provide radial wind speed measurements that need to be filtered and further processed before use. Significant amount of work in this thesis had to go into the data processing (cf. Chapter 4). Up to now, there is no open source toolbox available to process data, and every end user writes their own code. The solution lies in creating such a toolbox that allows the exchange of data processing methods and code. Also lidar suppliers should support this community effort and adapt to the needs of their customers better.

7.3 Lessons learned

There are three challenges that need to be overcome in order to establish lidar-based forecasting as a state of the art minute-scale forecasting method:

- the availability of the measurements,
- the robustness of the device and its cost,
- and the lack of standards.

The availability of the lidar measurement depends on the aerosol content in the air and so is not guaranteed. Therefore the maximum measurement range varies and is especially impacted by environmental conditions. Humidity has the strongest influence. If the lidar is used for minute-scale forecasting, a fallback solution needs to be implemented that still guarantees a forecast, even if there is no lidar data available.

Apart from environmental conditions, the maximum measurement range of a lidar is also impacted by the trajectory setting and the data processing, meaning the number of pulses per measurement and the filtering method used to find outliers. The maximum ranges were achieved with 40,000 pulses and a filter that uses only the wind velocity data achieved the highest ranges.

The forecasting horizon of lidar-based forecasting depends on the wind speed but also on the maximum measurement range of the lidar. A measurement range between 2 km and 7 km is needed, for forecasting horizons of 10 minutes. For forecasting horizons of more than twenty

minutes, the measurement range needs to be at least 5 km for low wind speeds of 4 m s^{-1} and over 10 km for wind speeds of 8 m s^{-1} and more.

Scanning-lidars are still a relatively new technology. The robustness of the devices and their cost are perceived as a challenge. However, the lidar used in this thesis, worked reliably in measurement campaigns of several months on- and offshore. A maintenance before and after the campaign at the manufacturer is still recommended strongly. Although the current costs of scanning lidar are high, it is expected that the costs will drop in the next years. Initiatives like OpenLidar, striving for open, modular lidar systems, might help this development.

The lack of standards for using lidar for an application such as minute-scale forecasting can be overcome with the help of the international community gathered around the IEA Wind initiative. Task 32 on wind lidar and Task 36 on forecasting could work together to enable a community driven exchange of experience that results in the writing of Recommended Practices.

Summary and conclusions

This chapter concludes this thesis. First a summary of the major findings are given in Section 8.1, then the future work is outlined in Section 8.2 and the final conclusions are drawn in Section 8.3.

8.1 Summary

In this thesis, the use of a long-range lidar to forecast wind power on the minute-scale was investigated. The objective of the thesis, as stated in the introduction, was to answer four fundamental questions about lidar-based minute-scale forecasting. In the following section the answers that were found are summarised.

8.1.1 How should the lidar data information be processed to gain a power forecast for a wind turbine?

When the work for this thesis was begun, no prior knowledge on how to process lidar data to gain a power forecast for a wind turbine was available. A process of a forecasting chain was established for this purpose. The chain describes the transformation of measured radial wind speed data to horizontal wind speed and wind direction as step number one; the propagation of the measured wind speed information to the wind turbine, with the help of an advection model as step number two; and finally the transformation of the forecasted wind speed at the

turbine into forecasted power using the power curve as step number three.

Along this forecasting chain methods that were previously applied for other applications, were adapted to the need of lidar-based forecasting. In step number one, a new filter method to detect wind speed outliers in the lidar data was optimized for the use of long-range lidar data. To transform the radial wind speed measurements into horizontal wind speed and wind direction information, a new global-local wind field reconstruction algorithm was introduced. The result was the wind field information upstream of the turbine that was needed for the forecast.

In step number two, a simple Taylor-based advection model was implemented. This model sets a linear relationship between the distance of each lidar measurement point, the measured wind speed value at the point, and the time when the measured wind speed vector will arrive at the turbine. With this information, a wind speed forecast at the turbine became possible.

In step number three, the forecasted wind speed information was processed. Each forecasted wind speed was transformed into a forecasted power with the power curve of the wind turbine. The power curves were generated with measured wind speed and power data. To account for the uncertainty in the wind speed and power forecast, probabilistic methods were applied to quantify the uncertainty. The probabilistic forecasts were then evaluated with the help of the actual wind speed and power measurements.

The forecast chain that was established in this thesis is valid for the forecast of a single wind turbine. However, to use it for the power forecast of a whole wind farm, it can be extended to take wake effects into account. It is implemented as a modular process and if other methods along the forecasting chain should be tested, existing methods can easily be swapped.

8.1.2 How does the measurement setup and measurement site influence the forecast?

To answer this question, two different measurement campaigns were carried out over several months with the same lidar. The two campaigns had a different measurement setup and were conducted at different locations.

During the first campaign, the lidar was installed fixed on top of a radio tower, next to the wind turbine that the forecast was made for. Here the lidar measured horizontally with a fixed azimuth angle into the same direction. The measurement direction was aligned into the main wind direction. During the second campaign, the lidar was installed on top of the nacelle of a wind turbine. Again, the lidar measured horizontally with a fixed azimuth angle, but this time the lidar rotated along with the nacelle and therefore the measurement direction was automatically aligned with the wind direction.

The different measurement setup resulted in a different availability of the forecasts. In the first campaign, lidar-based forecasts were only possible, when the wind direction was aligned with the measurement direction. This resulted in a low forecast availability. In contrast, the forecast availability was much improved in the second campaign, where the lidar rotated along with the wind direction. It is therefore concluded that the lidar should be installed either on top of the nacelle of a wind turbine, or the lidar measurement direction should automatically be adapted to the wind direction if the lidar is installed on a fixed platform.

In terms of location, the difference was that the first campaign was conducted onshore, in complex terrain and the forecast was carried out for a free standing turbine. The second campaign was carried out offshore at the alpha ventus wind farm and the lidar was installed on top of the AV04 turbine, which is surrounded by other turbines of the wind farm.

The location had an effect on the forecast accuracy. The forecast error was significantly increased at the onshore site. It was deduced that the simple Taylor-based propagation model used in this thesis, does not model the actual propagation at the complex onshore site well enough. To use lidar-measurements to forecast in complex terrain, a more sophisticated model will be necessary.

8.1.3 What is the forecast horizon of lidar-based forecasts and what influences the forecast horizon? How does the lidar-based forecast perform in comparison to state-of-the-art statistical methods and what are its benefits?

These last two questions can only be answered together. In principle the forecast horizon depends on the measurement range of the lidar and the measured wind speed. The further the lidar measures and the lower the wind speed, the longer the possible forecast horizon can be. In this thesis three horizons were tested: 0-10min, 10-20min, and 20-30min. The reason for this choice was the time resolution of the available power data that was used to validate the forecasts. Only ten minute-averages were available, and the forecast horizon was aligned with this resolution.

At the onshore and offshore campaign, forecasts were available in all forecast horizons up to 30 minutes. However, the number of available forecasts decreased significantly in the horizons 10-20min and 20-30min. The reason can be found in the measurement range of the lidar. The range was found to be around 5 km on average at the onshore site, which limited the forecast horizon. For a wind speed of 7 m s^{-1} this would mean a forecast horizon of no more than 12 minutes. For a longer forecast horizon, a longer lidar measurement range is necessary.

Also the forecast accuracy decreased along with the forecast horizon. In fact, the lidar-based

forecast outperformed the persistence model only in the first 0-10min horizon offshore. In the horizons from 10 minutes to 30 minutes ahead, the persistence model performed better. The reason is that only small wind speeds arrived with a delay of more than 10 minutes and therefore the forecast was biased. At the onshore campaign, the persistence model outperformed the lidar-based forecast in all forecast horizons.

The lidar-based forecast was able to forecast changes in power significantly better than the persistence model (11 % maximum error versus 17 %). Here the new method shows its strengths: the preview information of the upstream wind speed measurements allow forecasting of rapid changes in wind speed and power. During wind ramps, the lidar-based forecast error was therefore significantly smaller than the error of persistence. As persistence only uses old power data, future changes cannot be predicted. As wind ramps are potentially critical to the grid stability, or can affect the cost of balancing the power system if they are not forecast well, using lidars at wind farms to improve the power forecast is advised.

8.2 Discussion and future work

The forecasting chain presented in this thesis for lidar-based minute-scale forecasting has a significant benefit: the modularity of the process chain allows for improvements. The methods in each chain link can be adapted, optimized or swapped for better methods to improve the forecast. This section goes through each link and suggests such improvements.

In the first forecasting link, the lidar data processing, the following improvements are suggested:

- **Data filtering:** when the work for this thesis began, no tools were available to filter long range lidar data adequately. A new method was developed, to find outliers in the radial wind speed data. The result was a simple method that worked well, but not all wind speed outliers were identified. By now more work has been done by others in the wind lidar community on optimizing filtering methods. A new, dynamic filter approach that is based on the assumption of self-similarity showed promising results [66]. Another approach, using spatial information and a clustering algorithm was able to retrieve 38 % more data compared to the CNR filtering [67]. In the future, the different filter approaches should be compared to find the best solution for minute-scale forecasting.
- **Measurement trajectory:** at numerous points in this thesis, it was pointed out that it is necessary to align the measurement direction into the wind direction, to achieve a high availability of the forecast. It is therefore advised to install the lidar on top of the nacelle of a wind turbine. If this (for some reason) cannot be done, and the lidar is fixed on a platform next to or in the wind farm, the measurement direction should be adapted

automatically into the wind direction. The lidar used in the measurement campaign does not offer such a feature up to now, but this could be developed. In case the lidar is mounted on top of the nacelle of a turbine, an algorithm should be developed that ensures that the lidar trajectory automatically measures horizontally taking into account the turbine tilt

In the second link in the forecasting chain, the advection of the wind field using propagation model, the following improvements are suggested:

- Advection model: This thesis has shown that the simple advection model based on Taylor's hypothesis does not work in complex terrain. In the future, an optimized propagation model should be tested, which accounts for the complex and turbulent flow in such terrain. Such models could be for example WASP or other flow models [68] or a simple model such as the open source tool Wind Ninja [69].
- Coupling to NWP model: another approach to optimize the propagation model is to couple the measurements to a high resolution NWP model such as WRF, and to assimilate the lidar data into the model. In the current ParkCast project, which also provided the offshore lidar data for this thesis, this approach is tested. The idea is to use the lidar data to correct the wind speed error of the NWP model, thus optimizing the minute-scale forecast.

In the last link of the forecasting chain, the transformation from forecasted wind speed to forecasted power, the following improvement are suggested:

- Extend forecast horizon: the forecast horizon in this thesis was limited to 10 minutes. Forecasts of more than 10 minutes were biased and underestimated the wind power generation. The reason, as explained, was the measurement range of the lidar. In order to extend the forecast horizon, new lidar developments are necessary. This is the task of lidar manufacturers. By increasing the measurement range well over 10 km, while at the same time keeping a compact, lightweight shape of the lidar devices, minute-scale forecasting using lidar would be significantly improved.
- Adapt to forecast power of wind farms: this thesis optimized power forecast of a single wind turbine using lidar data. The next step is to adapt this approach to forecast the power output of a whole wind farm. First works by colleagues in this direction have been carried out already. Power output of seven wind turbines from the offshore wind farm Global Tech 1 has been successfully forecasted using a long range-lidar [70]. In Australia the power output of a whole wind farm was forecasted by converting forward-propagated lidar wind fields to turbine generation using in situ power curves [71]. In order to account for wake effects between the turbines, wake models need to be implemented. The wake

models are needed to predict the wind speeds at downwind turbines. Such a wind farm model that accounts for the wakes is for example FLORIS [72].

8.3 Conclusions

This work shows that lidar-based forecasts have great potential to improve forecasts in the minute-scale. Compared to the benchmark method (persistence), offshore forecasts up to 10 minutes were more accurate. The method is not perfect: on land, forecasts on all time horizons were worse, due to the effects of complex terrain on wind field propagation. Also, offshore forecasts for more than 10 minutes ahead were less accurate than persistence.

Also other barriers to adoption remain. As the availability of the measurements and therefore the forecasts is not guaranteed, a fallback solution needs to be implemented. The lidar hardware needs to be improved to allow for an increased measurement range, which in turn would lead to an extended forecast horizon. The cost of lidar hardware is still high and expert knowledge is necessary to handle the measurements. Standards are not available yet.

However, wind lidar data coupled with propagation models and power curves has fundamental advantages for minute-scale wind power forecasting. It is the only method that processes preview information of the inflow of the wind turbine or wind farm and measures changes in wind speed and power ahead of time. Although this thesis has shown that current approaches may not be perfect, the rapid pace of wind lidar technology development, the increasing number of users, the fact that recommended practices are just being written that include the use of lidar, and the growing network of third party service providers, suggests that wind lidar can play a role in the future of minute-scale wind power forecasting.

A

Appendix

A.1 Lidar Data Sheet

Parameter	Value
Eye safety	Class 1M
Wavelength	1.5 μm
Laser pulse energy	$\approx 100 \mu\text{J}$
Laser pulse width	150 ns
Pulse rate	15 kHz
Nyquist velocity	19.4 m s^{-1}
Unambiguous range	10 km
Aperture	75 mm
Volume approximately	0.5 m^3
Power consumption	<300 W
Temporal resolution	selectable from 0.1 to 30 s
Range gate size	18 to 60 m
Velocity precision	<20 cm s^{-1} for SNR > -17 dB
Minimum range	<100 m, typically 75 m
Scanning	Step-stare, full upper hemisphere
Enclosure	Weatherproof, temperature stabilized

Table A.1: Data sheet for Stream Line XR lidar.

A.2 Met mast sensor equipment




Figure A.1: Sensor equipment of the met mast in Stötten, effective 13.10.2015

Sensor	Heights [m]
Wind speed (cup anemometer)	34, 41, 51, 61, 71, 81, 91
Wind direction (wind vane)	34, 51, 71, 91
Wind speed and direction (sonic anemometer)	42, 62, 82
Air pressure	21, 92
Air temperature	34, 42, 52, 72, 101
Precipitation	24, 101
Relative humidity	34, 42, 52, 72, 101

Table A.2: Sensors on the met mast of the offshore research platform FINO1.

Bibliography

- [1] Wikipedia contributors, “Weather forecasting — Wikipedia, the free encyclopedia,” 2021, [Online; accessed 21-June-2021]. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Weather_forecasting&oldid=1029123018
- [2] G. Giebel and G. Kariniotakis, “3 – wind power forecasting –a review of the state of the art,” in *Renewable Energy Forecasting*, ser. Woodhead Publishing Series in Energy, G. Kariniotakis, Ed. Woodhead Publishing, 2017, pp. 59 – 109. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/B9780081005040000032>
- [3] X. Wang, P. Guo, and X. Huang, “A review of wind power forecasting models,” *Energy Procedia*, vol. 12, pp. 770–778, 2011, the Proceedings of International Conference on Smart Grid and Clean Energy Technologies (ICSGCE 2011). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1876610211019291>
- [4] A. Clifton, N. Vasiljevic, I. Wuerth, S. Raach, F. Haizmann, and H. Fürst, “The OpenLidar Initiative for collaboration on wind lidar hardware and software,” Sep. 2019, This version is a first review version and may be subject to revision. [Online]. Available: <https://doi.org/10.5281/zenodo.3414197>
- [5] D. Schlipf, A. Rettenmeier, F. Haizmann, M. Hofsäß M. Courtney, and P. W. Cheng, “Model based wind vector field reconstruction from lidar data,” *German Wind Energy Conference DEWEK 2012, Bremen, Germany*, 2012. [Online]. Available: <http://dx.doi.org/10.18419/opus-8136>
- [6] Y. Zhang, J. Wang, and X. Wang, “Review on probabilistic forecasting of wind power generation,” *Renewable and Sustainable Energy Reviews*, vol. 32, pp. 255 – 270, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1364032114000446>
- [7] E. Simon, “Minute-scale wind forecasting using lidar inflow measurements,” Ph.D. dissertation, Technical University of Denmark, Denmark, 2019.
- [8] L. Valldecabres, A. Peña, M. Courtney, L. von Bremen, and M. Kühn, “Very short-term

- forecast of near-coastal flow using scanning lidars,” *Wind Energy Science*, vol. 3, no. 1, pp. 313–327, 2018. [Online]. Available: <https://wes.copernicus.org/articles/3/313/2018/>
- [9] L. Valdecabres, N. G. Nygaard, L. Vera-Tudela, L. Von Bremen, and M. Kühn, “On the use of dual-Doppler radar measurements for very short-term wind power forecasts,” *Remote Sensing*, vol. 10, no. 11, 2018. [Online]. Available: <https://www.mdpi.com/2072-4292/10/11/1701>
- [10] L. Valdecabres, N. Nygaard, L. von Bremen, and M. Kühn, “Very short-term probabilistic forecasting of wind power based on dual-Doppler radar measurements in the North Sea,” *Journal of Physics: Conference Series*, vol. 1037, p. 052010, jun 2018. [Online]. Available: <https://doi.org/10.1088/1742-6596/1037/5/052010>
- [11] L. Valdecabres, L. von Bremen, and M. Kühn, “Minute-scale detection and probabilistic prediction of offshore wind turbine power ramps using dual-Doppler radar,” *Wind Energy*, vol. 23, no. 12, pp. 2202–2224, 2020. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.2553>
- [12] I. Würth, L. Valdecabres, E. Simon, C. Möhrlen, B. Uzunoglu, C. Gilbert, G. Giebel, D. Schlipf, and A. Kaifel, “Minute-scale forecasting of wind power – results from the collaborative workshop of IEA wind task 32 and 36,” *Energies*, vol. 12, no. 4, 2019. [Online]. Available: <https://www.mdpi.com/1996-1073/12/4/712>
- [13] C. L. Vincent and P.-J. Trombe, “8 - forecasting intrahourly variability of wind generation,” in *Renewable Energy Forecasting*, ser. Woodhead Publishing Series in Energy, G. Kariniotakis, Ed. Woodhead Publishing, 2017, pp. 219 – 233. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/B9780081005040000081>
- [14] C. Gallego-Castillo, A. Cuerva-Tejero, and O. Lopez-Garcia, “A review on the recent history of wind power ramp forecasting,” *Renewable and Sustainable Energy Reviews*, vol. 52, pp. 1148 – 1157, 2015.
- [15] A. Truewind, “AWS Truewind’s final report for the Alberta forecasting pilot project,” AWS Truewind, Tech. Rep., 2008.
- [16] I. Würth, S. Ellinghaus, M. Wigger, M. J. Niemeier, A. Clifton, and P. W. Cheng, “Forecasting wind ramps: can long-range lidar increase accuracy?” *Journal of Physics: Conference Series*, vol. 1102, p. 012013, oct 2018. [Online]. Available: <https://doi.org/10.1088%2F1742-6596%2F1102%2F1%2F012013>
- [17] G. Pearson, F. Davies, and C. Collier, “An analysis of the performance of the UFAM pulsed Doppler lidar for observing the boundary layer,” *Journal of Atmospheric and Oceanic Technology*, vol. 26, no. 2, pp. 240–250, publisher: American Meteorological

- Society Section: Journal of Atmospheric and Oceanic Technology. [Online]. Available: https://journals.ametsoc.org/view/journals/atot/26/2/2008jtecha1128_1.xml
- [18] A. A. Kokhanovsky, *Aerosol optics: light absorption and scattering by particles in the atmosphere*. Berlin: Springer, 2008. [Online]. Available: <http://swbplus.bsz-bw.de/bsz277453429cov.htm>
- [19] Andrew Clifton, Michael Courtney, and Dennis Elliott, “Ground-based vertically-profiling remote sensing for wind resource assessment,” IEA Wind, Tech. Rep. Technical Report RP 15.
- [20] D. Schlipf, “Lidar assisted control of wind turbines,” Ph.D. dissertation, Stuttgart Wind Energy at Institute of Aircraft Design, University of Stuttgart, 2015.
- [21] R. Wagner, B. Cañadillas, A. Clifton, S. Feeney, N. Nygaard, M. Poodt, C. S. Martin, E. Tüxen, and J. W. Wagenaar, “Rotor equivalent wind speed for power curve measurement – comparative exercise for IEA wind annex 32,” *Journal of Physics: Conference Series*, vol. 524, p. 012108, jun 2014. [Online]. Available: <https://doi.org/10.1088/1742-6596/524/1/012108>
- [22] P. J. M. Clive, “7.16_Clive: Offshore wind turbine power curve tests using lidar,” Aug. 2019. [Online]. Available: <https://doi.org/10.5281/zenodo.3378685>
- [23] L. Landberg, G. Giebel, H. A. Nielsen, T. Nielsen, and H. Madsen, “Short-term prediction – an overview,” *Wind Energy*, vol. 6, no. 3, pp. 273–280, 2003. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.96>
- [24] M. Felder, A. Kaifel, B. Matthiss, K. Ohnmeiß, L. Schröder, F. Sehnke, I. Würth, M. Wigger, and P. W. Cheng, “Optimierung der Auslegung und Betriebsführung von Kombikraftwerken und Speichertechnologien mittels Kurzestfristvorhersagen der Wind- und PV-Leistung (VORKAST) : Abschlussbericht : Laufzeit des Vorhabens: 01.09.2014 bis 31.10.2017,” Zentrum für Sonnenenergie- und Wasserstoff-Forschung Baden-Württemberg (ZSW), Tech. Rep., 2018. [Online]. Available: <https://doi.org/10.2314/GBV:1028384602>
- [25] N. Vasiljevia, J. Palma, N. Angelou, J. Matos, R. Menke, G. Lea, J. Mann, M. Courtney, L. Ribeiro, and V. Gomes, “Perdigão 2015: Methodology for atmospheric multi-Doppler lidar experiments,” *Atmospheric Measurement Techniques*, vol. 10, no. 9, pp. 3463–3483, 2017.
- [26] M. Hofsäβ, A. Clifton, and P. W. Cheng, “Reducing the uncertainty of lidar measurements in complex terrain using a linear model approach,” *Remote Sensing*, vol. 10, no. 9, 2018.
- [27] W. Commons, “File:germany (+districts +municipalities) location map current.svg —

- wikimedia commons, the free media repository,” 2020, [Online; accessed 12-July-2021]. [Online]. Available: [https://commons.wikimedia.org/w/index.php?title=File:Germany_\(%2Bdistricts_%2Bmunicipalities\)_location_map_current.svg&oldid=460638089](https://commons.wikimedia.org/w/index.php?title=File:Germany_(%2Bdistricts_%2Bmunicipalities)_location_map_current.svg&oldid=460638089)
- [28] Federal Maritime and Hydrographic A and Federal Ministry for the Environment, Eds., *Ecological Research at the Offshore Windfarm alpha ventus*. Springer Fachmedien Wiesbaden. [Online]. Available: <http://link.springer.com/10.1007/978-3-658-02462-8>
- [29] V. Pettas, M. Kretschmer, A. Clifton, and P. W. Cheng, “On the effects of inter-farm interactions at the offshore wind farm alpha ventus,” *Wind Energy Science Discussions*, vol. 2021, pp. 1–25, 2021. [Online]. Available: <https://wes.copernicus.org/preprints/wes-2021-50/>
- [30] Lencer. (2009) Lage und Detailkarte des Windparks alpha ventus (mit Umspannplattform und FINO 1) in der Nordsee. Published under the Creative Commons Attribution-Share Alike 3.0 Unported license. [Online]. Available: https://commons.wikimedia.org/wiki/File:Windpark_alpha_ventus_Lagekarte.png
- [31] M. D. (Chumwa). (2020) Map of offshore wind farms and connecting power cables in the German Bight. Published under the Creative Commons Attribution-Share Alike 3.0 Unported license. [Online]. Available: https://commons.wikimedia.org/wiki/File:Map_of_the_offshore_wind_power_farms_in_the_German_Bight.png
- [32] D. Muñoz Esparza, B. Canadillas, T. Neumann, and J. Beeck, “Turbulent fluxes, stability and shear in the offshore environment: Mesoscale modelling and field observations at FINO1,” *Journal of Renewable and Sustainable Energy*, vol. 4.
- [33] A. Westerhellweg, T. Neumann, and V. Riedel, “FINO1 mast correction,” *DEWI-Magazin*, vol. 21. [Online]. Available: <https://www.osti.gov/etdeweb/biblio/21558009>
- [34] R. Jain, R. Kasturi, and B. G. Schunck, *Machine Vision*, 01 1995.
- [35] P. U. I. 7854. (2020) Stuttgart fernsehturm. Published under the Pixabay license, with free commercial use and no image credit necessary. [Online]. Available: <https://pixabay.com/de/photos/stuttgart-fernsehturm-turm-69069/>
- [36] D. Schlipf, “Lidar-assisted control concepts for wind turbines,” accepted: 2016-06-24T12:35:42Z. [Online]. Available: <http://elib.uni-stuttgart.de/handle/11682/8813>
- [37] G. I. Taylor, “The spectrum of turbulence,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 164, no. 919, pp. 476–490, publisher: Royal Society. [Online]. Available: <https://royalsocietypublishing.org/doi/10.1098/rspa.1938.0032>
- [38] R. B. Stull, *An Introduction to Boundary Layer Meteorology*, ser. Atmospheric

- and Oceanographic Sciences Library. Springer Netherlands. [Online]. Available: <https://www.springer.com/de/book/9789027727688>
- [39] J. Done, C. A. Davis, and M. Weisman, “The next generation of NWP: explicit forecasts of convection using the weather research and forecasting (WRF) model,” *Atmospheric Science Letters*, vol. 5, no. 6, pp. 110–117, 2004. [Online]. Available: <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/asl.72>
- [40] T. N. Palmer, “The economic value of ensemble forecasts as a tool for risk assessment: From days to decades,” *Quarterly Journal of the Royal Meteorological Society*, vol. 128, no. 581, pp. 747–774, 2002. [Online]. Available: <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1256/0035900021643593>
- [41] A. Timmermann, “Density forecasting in economics and finance,” *Journal of Forecasting*, vol. 19, no. 4, pp. 231–234, 2000. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/1099-131X%28200007%2919%3A4%3C231%3A%3AAID-FOR771%3E3.0.CO%3B2-%23>
- [42] L. Alkema, A. E. Raftery, and S. J. Clark, “Probabilistic projections of HIV prevalence using bayesian melding,” *Annals of Applied Statistics*, vol. 1, no. 1, pp. 229–248, publisher: Institute of Mathematical Statistics. [Online]. Available: <https://projecteuclid.org/euclid.aoas/1183143737>
- [43] J. L. S. Charles M. Grinstead, “Conditional probability - discrete conditional,” in *Grinstead & Snell’s Introduction to Probability*, ser. Orange Grove Texts. [Online]. Available: https://www.dartmouth.edu/~chance/teaching_aids/books_articles/probability_book/Chapter4.pdf
- [44] J. V. Seguro and T. W. Lambert, “Modern estimation of the parameters of the weibull wind speed distribution for wind energy analysis,” *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 85, no. 1, pp. 75–84. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167610599001221>
- [45] M. C. Jones, J. S. Marron, and S. J. Sheather, “A brief survey of bandwidth selection for density estimation,” *Journal of the American Statistical Association*, vol. 91, no. 433, pp. 401–407, 1996. [Online]. Available: <https://www.tandfonline.com/doi/abs/10.1080/01621459.1996.10476701>
- [46] A. W. Bowman and A. Azzalini, *Applied Smoothing Techniques for Data Analysis: The Kernel Approach with S-Plus Illustrations*, ser. Oxford Statistical Science Series. Oxford University Press.
- [47] P. Pinson, H. A. Nielsen, J. K. Møller, H. Madsen, and G. N. Kariniotakis,

- “Non-parametric probabilistic forecasts of wind power: required properties and evaluation,” *Wind Energy*, vol. 10, no. 6, pp. 497–516, 2007. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.230>
- [48] T. Gneiting, F. Balabdaoui, and A. E. Raftery, “Probabilistic forecasts, calibration and sharpness,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 69, no. 2, pp. 243–268, 2007. [Online]. Available: <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9868.2007.00587.x>
- [49] S. Vannitsem, D. S. Wilks, and J. Messner, *Statistical Postprocessing of Ensemble Forecasts*. Elsevier, google-Books-ID: yKNPDwAAQBAJ.
- [50] A. Jordan, F. KrÄ4ger, and S. Lerch, “Evaluating probabilistic forecasts with scoringRules,” *Journal of Statistical Software*, vol. 90, no. 1, pp. 1–37, number: 1. [Online]. Available: <https://www.jstatsoft.org/index.php/jss/article/view/v090i12>
- [51] W.-Y. Chang, “A literature review of wind forecasting methods,” *Journal of Power and Energy Engineering*, vol. 2, no. 4, pp. 161–168. [Online]. Available: <http://www.scirp.org/Journal/Paperabs.aspx?paperid=44881>
- [52] L. Vermeer, J. Sørensen, and A. Crespo, “Wind turbine wake aerodynamics,” *Progress in Aerospace Sciences*, vol. 39, no. 6, pp. 467–510, 2003. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0376042103000782>
- [53] V. Spiridonov and M. Ćurić, *Atmospheric Stability*. Cham: Springer International Publishing, 2021, pp. 115–121. [Online]. Available: https://doi.org/10.1007/978-3-030-52655-9_9
- [54] M. C. Holtslag, W. A. A. M. Bierbooms, and G. J. W. v. Bussel, “Estimating atmospheric stability from observations and correcting wind shear models accordingly,” *Journal of Physics: Conference Series*, vol. 555, p. 012052. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1742-6596/555/1/012052>
- [55] A. Clifton, S. Schreck, G. Scott, N. Kelley, and J. Lundquist, “Turbine inflow characterization at the national wind technology center,” *Journal of Solar Energy Engineering*, vol. 135.
- [56] D. R. Drew, D. J. Cannon, J. F. Barlow, P. J. Coker, and T. H. Frame, “The importance of forecasting regional wind power ramping: A case study for the UK,” *Renewable Energy*, vol. 114, pp. 1201–1208, 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0960148117306973>
- [57] J. W. Zack, “Optimization of wind power production forecast performance during critical periods for grid management,” accessed: 31-5-2021. [Online]. Available: https://aws-dewi.ul.com/assets/AWEA_Windpower_2007_Forecasting.pdf

- [58] L. Bianco, I. V. Djalalova, J. M. Wilczak, J. Cline, S. Calvert, E. Konopleva-Akish, C. Finley, and J. Freedman, “A wind energy ramp tool and metric for measuring the skill of numerical weather prediction models,” *Weather and Forecasting*, vol. 31, no. 4, pp. 1137–1156, publisher: American Meteorological Society Section: Weather and Forecasting. [Online]. Available: https://journals.ametsoc.org/view/journals/wefo/31/4/waf-d-15-0144_1.xml
- [59] M. Türk, K. Grigutsch, and S. Emeis, “The wind profile above the sea-investigations basing on four years of FINO 1 data,” *DEWI Magazine*, vol. 33, pp. 12–16, 01 2008.
- [60] S. Davoust, A. Jehu, M. Bouillet, M. Bardon, and B. Vercherin, “Assessment and optimization of lidar measurement availability for wind turbine control,” *EWEA, Barcelona*, 2014.
- [61] A. Brenner, “Statistische Auswertung der Reichweite von long-range Lidardaten unter Einbeziehung von Umgebungsbedingungen,” Master’s thesis, University of Stuttgart, Stuttgart, 2017.
- [62] I. Würth, A. Brenner, M. Wigger, and P. W. Cheng, “How far do we see? Analysis of the measurement range of long-range lidar data for wind power forecasting,” accepted: 2018-03-15T17:48:48Z. [Online]. Available: <http://elib.uni-stuttgart.de/handle/11682/9710>
- [63] A. A. Kokhanovsky, *Aerosol optics: light absorption and scattering by particles in the atmosphere*. Berlin: Springer, 2008. [Online]. Available: <http://swbplus.bsz-bw.de/bsz277453429cov.htm>
- [64] D. N. Reshef, Y. A. Reshef, H. K. Finucane, S. R. Grossman, G. McVean, P. J. Turnbaugh, E. S. Lander, M. Mitzenmacher, and P. C. Sabeti, “Detecting novel associations in large data sets,” *Science*, vol. 334, no. 6062, pp. 1518–1524, 2011. [Online]. Available: <https://science.sciencemag.org/content/334/6062/1518>
- [65] A. Clifton, N. Vasiljevic, I. Wuerth, S. Raach, F. Haizmann, and H. Fürst, “The OpenLidar Initiative for collaboration on wind lidar hardware and software,” Sep. 2019, This version has been published to obtain feedback and may be subject to revision. Readers are encouraged to check for more recent versions of this document. [Online]. Available: <https://doi.org/10.5281/zenodo.3414197>
- [66] H. Beck and M. Kühn, “Dynamic data filtering of long-range Doppler lidar wind speed measurements,” *Remote Sensing*, vol. 9, no. 6, 2017. [Online]. Available: <https://www.mdpi.com/2072-4292/9/6/561>
- [67] L. Alcayaga, “Filtering of pulsed lidar data using spatial information and a clustering algorithm,” *Atmospheric Measurement Techniques*, vol. 13, no. 11, pp. 6237–6254, 2020.

- [Online]. Available: <https://amt.copernicus.org/articles/13/6237/2020/>
- [68] E. Berge, A. Gravdahl, J. Schelling, L. Tallhaug, O. Undheim, K. Vindteknikk, A. Vector, A. Hydro, and O. Energy, “Wind in complex terrain. a comparison of WAsP and two CFD-models.” *European Wind Energy Conference and Exhibition 2006, EWEC 2006*, vol. 2, 01 2006.
- [69] N. S. Wagenbrenner, J. M. Forthofer, B. K. Lamb, K. S. Shannon, and B. W. Butler, “Downscaling surface wind predictions from numerical weather prediction models in complex terrain with WindNinja,” *Atmospheric Chemistry and Physics*, vol. 16, no. 8, pp. 5229–5241, 2016. [Online]. Available: <https://acp.copernicus.org/articles/16/5229/2016/>
- [70] F. Theuer, M. F. van Dooren, L. von Bremen, and M. Kühn, “Minute-scale power forecast of offshore wind turbines using long-range single-Doppler lidar measurements,” *Wind Energy Science*, vol. 5, no. 4, pp. 1449–1468, 2020. [Online]. Available: <https://wes.copernicus.org/articles/5/1449/2020/>
- [71] M. Pichault, C. Vincent, G. Skidmore, and J. Monty, “Short-term wind power forecasting at the wind farm scale using long-range Doppler lidar,” *Energies*, vol. 14, no. 9, 2021. [Online]. Available: <https://www.mdpi.com/1996-1073/14/9/2663>
- [72] 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*, vol. 19, no. 1, pp. 95–114, 2016. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1822>

Curriculum Vitae

Ines Würth was born on 25 March 1986 in Stuttgart, Germany.

She started studying aerospace engineering at the University of Stuttgart in October 2005. After the first half of her studies, she completed a 6-month internship at the wind turbine manufacture ENERCON in Aurich, Germany and found her passion for wind energy. In the second half of her studies, she then specialised in wind energy and aerodynamics and wrote her study thesis on lidar-based wind field measurements at Stuttgart Wind Energy (SWE) at the Institute of Aircraft Design. At the end of her studies, she worked for 6 months at the blade manufacturer LM



Wind Power in Kolding, Denmark for her diploma thesis on the spacing of vortex generators on wind turbine blades. She completed her thesis in March 2012.

Ines started working as a research assistant at SWE in April 2012. She joined the test and measurement group and worked on several national and international research projects in the field of lidar applications for wind energy. In 2015 she became the team lead of the group. She was responsible for the strategic development of the group, for project acquisition, and for the team development. In parallel she initiated and led projects on the use of lidar for minute-scale forecasting of wind power, which developed into her dissertation topic. In 2019, she stepped back as team lead of the test and measurement group and since then has focused on her dissertation.

Ines is part of several national and international committees which enable the deployment of lidar for wind energy. In the RAVE (Research At Alpha Ventus) consortium she represents and manages SWE's offshore lidar measurements on the first German offshore wind farm alpha ventus. As an advisory board member of IEA Wind Task 32, she introduced the use of lidar for minute-scale forecasting and organised several international workshops.

She has campaigned for more gender balance in engineering studies. In 2013 she took over the organisation of TryScience for aerospace engineering, an initiative to interest more girls in studying STEM subjects at the University of Stuttgart.

In August 2020 she co-founded the start-up enviConnect, a spin-off of the University of Stuttgart, together with team member of the test and measurement team. enviConnect offers technology transfer, rapid R&D, and research project management for wind energy.