

A topographic map of the North Sea region, showing the United Kingdom to the west and the European continent to the east. The map uses a color gradient where green represents lower elevations and brown/orange represents higher elevations. The sea is depicted in shades of blue. Overlaid on the map are three black silhouettes of wind turbines, positioned in the central North Sea. A semi-transparent light blue rounded rectangle is placed in the upper left quadrant, containing the title text.

# CHARACTERISATION OF OFFSHORE WINDS FOR ENERGY APPLICATIONS

PETER C. KALVERLA





# Propositions

1. A complete wind climatology must include thermal effects. (*this thesis*)
2. Wind energy meteorology is mostly engineering, not science. (*this thesis*)
3. Good coding principles also make good governance principles.
4. (Inter)national science foundations should take the lead in establishing next-generation publishing platforms.
5. A PhD without philosophy is like a wind turbine without wind.
6. Where quantity reigns, quality suffers.
7. Sustainable policies strive for balance instead of growth.
8. Retirement should be phased out in favour of life-long part-time careers.

Propositions belonging to the thesis, entitled:

**“Characterisation of offshore winds for energy applications”**

Peter C. Kalverla

Wageningen, 13 November, 2019



# **Characterisation of offshore winds for energy applications**

Peter C. Kalverla



## **Thesis committee**

### **Promotor:**

Prof. Dr A. A. M. Holtslag  
Professor of Meteorology  
Wageningen University & Research

### **Co-promotors:**

Dr G. J. Steeneveld  
Associate Professor, Meteorology and Air Quality  
Wageningen University & Research

Dr R. J. Ronda  
Researcher, Meteorology and Air Quality  
Wageningen University & Research

### **Other members:**

Prof. Dr R. Uijlenhoet, Wageningen University & Research  
Dr A. N. Hahmann, Technical University Denmark  
Prof. Dr B. Koren, Technical University Eindhoven  
Dr A. van Delden, Institute for Marine and Atmospheric Research, Utrecht University

This research was conducted under the auspices of the Graduate School for Socio-economic and Natural Sciences of the Environment (SENSE)

# Characterisation of offshore winds for energy applications

Peter C. Kalverla

## **Thesis**

submitted in fulfilment of the requirements for the degree of doctor  
at Wageningen University  
by the authority of the Rector Magnificus  
Prof. Dr A.P.J. Mol,  
in the presence of the  
Thesis Committee appointed by the Academic Board  
to be defended in public  
on Wednesday 13 November 2019  
at 11 a.m. in the Aula.

Peter C. Kalverla

Characterisation of offshore winds for energy applications,  
190 pages.

PhD thesis, Wageningen University, Wageningen, the Netherlands (2019)  
With references, with summary in English

ISBN: 978-94-6395-092-3

DOI: <https://doi.org/10.18174/498797>



## *A Sunlight Sonata*

*The sunlight shines upon earth's silent skin,  
and wakes the wind by tickling heaven's toes.*

*Disturbed, the air to colder quarters goes;  
it must abate the troubled state it's in.*

*A breeze at first, but strong the current grows.*

*Rotation brings about a soaring spin.*

*Destructive energy appears within  
what, once so peaceful, now with fury blows...*

*... till ragged claws of cold and distant earth,  
immune to sunbeams, grab at heaven's feet.*

*Majestic storms to smaller whirls give birth*

*as tumbling forward they release their heat.*

*A tranquil moment lasts, for what it's worth,  
for soon ensuing dawn is guaranteed.*

*Peter Kalverla, 2019*



# Contents

<b>English summary</b>	<b>xi</b>
<b>Nederlandse samenvatting</b>	<b>xv</b>
<b>Preface</b>	<b>xix</b>
<b>Chapter 1: General introduction</b>	<b>1</b>
1.1 A quest for sustainability . . . . .	3
1.2 Tapping from earth's energy cycle . . . . .	4
1.3 The evolution of turbinus ventus . . . . .	5
1.4 Wind energy meteorology . . . . .	6
1.4.1 Five intertwined branches . . . . .	6
1.4.2 Engineering models: from power formula to digital twins . . . . .	8
1.4.3 Meteorological data sources: from local observations to global simulations (and back) . . . . .	10
1.4.4 Introducing uncertainty . . . . .	11
1.5 Problem statement . . . . .	14
1.6 Outline of the thesis . . . . .	15
<b>Chapter 2: An observational climatology of anomalous wind events over the North Sea</b>	<b>19</b>
2.1 Introduction . . . . .	21
2.2 Background . . . . .	23
2.3 Methodology . . . . .	25
2.3.1 Description of the dataset . . . . .	25
2.3.2 Definition of anomalous wind events . . . . .	27
2.3.3 Weather conditions . . . . .	29
2.4 Results . . . . .	30
2.4.1 Background meteorological conditions . . . . .	30
2.4.2 Characterization of anomalous wind events . . . . .	33
2.4.3 (In)dependence of wind events . . . . .	40
2.5 Discussion . . . . .	41
2.6 Conclusions . . . . .	42
<b>Chapter 3: Evaluation of the IFS, WRF, and HARMONIE weather models with observations over the North Sea</b>	<b>45</b>
3.1 Introduction . . . . .	47
3.2 Methodology . . . . .	48
3.2.1 Description of the models . . . . .	48
3.2.2 Observations and model data . . . . .	50
3.2.3 Case selection: the $uvs \cdot t_2$ method . . . . .	50
3.2.4 Statistical analysis: error diagrams . . . . .	51
3.3 Results and discussion . . . . .	52
3.3.1 Case selection . . . . .	52



3.3.2	Overall performance . . . . .	54
3.3.3	Role of model resolution . . . . .	57
3.3.4	Boundary-layer structure . . . . .	57
3.3.5	A case against the average . . . . .	59
3.4	Summary and conclusions . . . . .	61

## Appendices

3.A	WRF set-up . . . . .	62
3.B	Selected cases . . . . .	63

## Chapter 4: Low-level jets over the North Sea in ERA5 and observations: together they do better 67

4.1	Introduction . . . . .	69
4.2	A brief description of both datasets and their shortcomings . . . . .	71
4.3	Jet detection: a precarious procedure . . . . .	73
4.4	Vertical range affects perceived jet morphology . . . . .	75
4.5	Datasets agree: most jets in spring and summer . . . . .	76
4.6	Simple scalings for the seasonal cycle . . . . .	77
4.7	Other jet characteristics and their scaling potential . . . . .	79
4.7.1	Diurnal cycle and stability . . . . .	79
4.7.2	Weather types and the spatial distribution of low-level jets . . . . .	81
4.8	Combining multiple predictors to extend observations . . . . .	82
4.9	Discussion . . . . .	85

## Appendices

4.A	Lidar observations . . . . .	87
-----	------------------------------	----

## Chapter 5: Understanding the physics of coastal winds and low-level jets 91

5.1	Introduction . . . . .	93
5.2	The kinship between sea breezes coastal jets . . . . .	93
5.2.1	On the evolution of the sea breeze . . . . .	94
5.2.2	Baroclinity and thermal wind . . . . .	96
5.2.3	Interaction with background wind . . . . .	97
5.3	A survey of explicit low-level jet studies . . . . .	99
5.3.1	Blackadar's inertial oscillation . . . . .	99
5.3.2	Holton's baroclinic mode . . . . .	100
5.3.3	Orographic effects . . . . .	101
5.4	Exploratory analysis of low-level jets in the ERA5 re-analysis data . . . . .	103
5.4.1	Intro – An animation is worth a thousand pictures . . . . .	104
5.4.2	From jet profile to jet event . . . . .	105
5.4.3	Clustering similar jet events together . . . . .	106
5.4.4	Back to basics: the hypothetico-deductive approach . . . . .	109
5.5	Summary . . . . .	114

## Chapter 6: Quality of wind characteristics in recent reanalyses over the North Sea 117

6.1	Introduction . . . . .	119
6.2	Description of the datasets . . . . .	120
6.3	Evaluation of wind speed and the role of atmospheric stability . . . . .	121

6.4	Evaluation of wind direction . . . . .	125
6.5	Anomalous wind events in reanalysis products . . . . .	126
6.5.1	Low-level jets . . . . .	126
6.5.2	Quantification of model performance for anomalous events . . . . .	128
6.5.3	Wind ramps . . . . .	129
6.5.4	Extreme shear . . . . .	133
6.5.5	Wind extremes . . . . .	134
6.6	A spatial low-level jet climatology based on the Dutch Offshore Wind Atlas . . .	136
6.7	Summary and conclusions . . . . .	137
<b>Chapter 7: General discussion: from idealized inflow fields to embedded emulations</b>		<b>139</b>
7.1	Introduction: a bird's-eye view . . . . .	141
7.2	Building a digital wind farm . . . . .	141
7.3	Incentives for a paradigm shift . . . . .	142
7.4	Towards data-driven case selection . . . . .	144
7.5	A note on model performance . . . . .	145
7.6	Establishing the status futurus . . . . .	146
7.7	Towards an energetic odyssey . . . . .	147
<b>Bibliography</b>		<b>149</b>
<b>Acknowledgements (Dankwoord)</b>		<b>161</b>
<b>About the author</b>		<b>165</b>
<b>List of publications</b>		<b>165</b>
<b>Education certificate</b>		<b>166</b>





# English summary

Offshore wind energy is regarded as one of the main ingredients of a sustainable energy system. A popular scenario for 2050 describes a stunning fleet of twenty-five thousand wind turbines that is to be installed in the North Sea. Obviously, the realisation of such a grand vision requires realistic estimates of the expected energy production, optimal wind farm layouts, and lifetime of various turbine designs. A wide range of computer models is used to simulate, for example, the interaction between the wind and the turbines. As inputs, they require (climatological) information about the wind field.

An inherent property of model simulations and associated predictions is that they come with uncertainties. There is always a chance that reality will deviate from the anticipated course. Since offshore wind projects require large investments, these uncertainties naturally transform into financial risks. Therefore, a growing interest emerges in uncertainty quantification and reduction.

This thesis addresses uncertainties that originate from the representation of the offshore wind field in the aforementioned engineering models. This goal is partly achieved through cross-validation between available observations, output from meteorological models, and combined products of observations and meteorological model simulations that are known as reanalysis datasets. Additionally, the thesis reflects on the impact of (over)simplifications that are commonly used to reduce computational demands.

Standard inflow specifications for engineering models have been built around highly idealized descriptions of the wind, such as the Weibull distribution, the logarithmic wind profile, and the ‘wind rose’. With the growth of individual wind turbines, offshore wind projects, and the corresponding investments, these simplified descriptions are becoming inadequate. Moreover, contemporary design standards consider conditions of extreme winds, extreme shear, wind gusts, etcetera in an abstract statistical fashion that does not reflect the physical nature of these events. **Chapter 1** further elaborates on the different models, data sources, assumptions and simplifications, and uncertainties.

The concept of anomalous wind events, introduced in **chapter 2**, provides a framework to define, identify and characterize wind events or conditions that are exceptional and are therefore usually not or poorly represented. Four years of measurement data up to 300 m height from a prospective wind farm site at the North Sea (met mast IJmuiden; MMIJ) are used to identify and characterize several such events. Low-level jets, wind speed maxima near the surface, stand out as an important phenomenon and also as an illustrative example. They occur predominantly in spring and early summer, with wind directions typically oriented parallel to the coast.

Other anomalous events are rapid fluctuations in wind speed or direction called wind ramps, and strong changes in wind speed or direction with height, called (extreme/anomalous) wind shear and wind veer. While these are not new, the chapter reflects on the way in which they are typically measured. Alternative methods are proposed and a first step is made towards a more in-depth characterization. Not surprisingly, atmospheric density stratification appears to be a relevant parameter.

**Chapter 3** presents an evaluation of three mainstream meteorological models. These models are not only used for weather forecasts and, in extension, power forecasts; they are also used to construct climatological datasets that provide inflow fields for engineering models. The observations from met mast IJmuiden provide a unique dataset to evaluate not only the wind speed near the surface, but also the representation of the vertical structure of the atmospheric boundary layer.

An innovative case-selection strategy is employed to select 30 days that span a wide range of weather conditions. It is shown that in terms of wind speed, this set is already quite representative for the climatology at the site. Most importantly, the evaluation demonstrates that the performance of the three models is quite comparable, with a typical bias smaller than 0.5 m/s and a root mean square error up to 2 m/s. The difference in model performance between the 30 cases is substantially larger, and it is concluded that stably stratified conditions still provide the largest challenge for weather models, especially concerning the vertical structure of the atmosphere.

A further characterization of low-level jets is presented in **chapter 4**. Additional observations and ERA5 reanalysis data are included to generalize the results for a larger area. Some key strengths and weaknesses of both types of datasets are identified along the way, and methods to overcome these limitations are demonstrated. Importantly, it is found that jet detection is very sensitive to the observation heights. This considerably influences the resulting climatology. Low-level jets in the ERA5 reanalysis data appear ‘smeared out’ and vertically displaced.

While the reanalysis data thus suffer from model deficiencies, the observations are mainly limited by their spatial and temporal extent. Even a four-year observation dataset still provides an erratic climatology of low-level jets, due to inter-annual variability. A smoother signal, more representative of the long-term climatology, is derived by combining the two data sources.

To better understand the climatology of low-level jets so obtained, a synthesis of the relevant literature on (coastal) atmospheric dynamics is provided in **chapter 5**. While there are several mechanisms that can produce jet-shaped profiles, most studies relate coastal jets to the thermal contrast between land and sea. As such, the coastal jet is closely related to the formation of the sea breeze, thermal lows, and the process of geostrophic adjustment. Other mechanisms include orographic effects, and an inertial oscillation similar to the forcing responsible for the well-described nocturnal jet over land.

Hypotheses taken from this literature survey are evaluated on the ERA5 reanalysis data. It turns out that indeed, the thermal contrast between land and sea provides an important mechanism for the formation of coastal low-level jets over the North Sea. Additionally, evidence is presented for contributions of inertial and orographic effects. While chapter 5 focuses on low-level jets, it also serves to illustrate the complexity of meteorological processes in coastal areas.

In the meantime, the weather models evaluated in chapter 3 have all been used to construct reanalysis datasets. **Chapter 6** presents an evaluation of these reanalysis data, focusing on the representation of wind speed and direction in general, but also on the representation of anomalous events. The state-of-the-art is that the Dutch Offshore Wind Atlas (DOWA) provides the best representation of the North Sea wind climate, with a typical root mean square error of 1.5 m/s and almost unbiased wind speed. The downscaling strategy employed in the New European Wind Atlas (NEWA) proved less fruitful, at least for the evaluated measurement site.

Wind direction is still not adequately represented in all datasets, and the error generally increases with height. Anomalous events are typically under-represented in the ERA5 data. The other reanalysis datasets perform better in a climatological sense, but the one-to-one correspondence with observations is still not superb, and is probably aggravated by the double-penalty problem.

The question of how to deal with anomalous events in practice is addressed in the general discussion in **chapter 7**, where a number of trends are identified that point to a much more data-driven approach for future wind farm/turbine simulations and towards more sophisticated models that can represent much more details of the flow. For this approach to become the new standard, a thorough evaluation of the relevant models and datasets is required. This thesis partly fills that void by presenting an evaluation at the top of the envisioned model chain, employing innovative methods to evaluate beyond standard summary statistics.

In terms of increased understanding, this thesis contributes by illustrating the complexity of the North Sea wind field. Compared to other ‘low-level jet hotspots’, the situation over the North Sea is not as clear-cut: the thermal forcing is less strong due to the relatively high latitude, and the area is characterized by complex interactions between the various surrounding coast-lines. It is argued that wind energy applications should properly account for these thermal and orographic effects in order to reduce uncertainties due to idealized inflow fields.



# Nederlandse samenvatting

*Wind op zee* als bron van hernieuwbare energie wordt gezien als belangrijk onderdeel van een toekomstig, duurzaam energiesysteem. Een populaire visie voor 2050 schetst een beeld van de Noordzee als centrale spil in de energievoorziening van de omliggende landen. Daartoe zouden maar liefst vijftientigduizend windturbines moeten worden geïnstalleerd. Zo'n omvangrijke onderneming vereist een realistische inschatting van de te verwachten energieopbrengst, de meest gunstige inrichting van de windparken, de levensduur van de turbines en alle bijbehorende kosten en opbrengsten. Die inschattingen leunen zwaar op geavanceerde modelsimulaties, waarin de interactie tussen de wind en de turbines wordt nagebootst.

Zelfs het meest gedetailleerde model is echter slechts een benadering van de werkelijkheid. Bovendien, al zou zo'n model de interactie tussen de wind en de turbines perfect kunnen beschrijven, dan is er altijd nog de onzekerheid die voortkomt uit de *input*: een klimatologische beschrijving van het windveld zelf. Deze onzekerheden brengen risico's met zich mee die niet kunnen worden verwaarloosd, gezien de grote investeringen die nodig zijn voor 'wind op zee'. Daarom is er een toenemende behoefte om niet alleen de modellen te verbeteren, maar vooral ook informatie toe te voegen over de (on)zekerheid van de simulaties.

Dit proefschrift voorziet in die behoefte door het onderzoeken van onzekerheden die voortkomen uit de (klimatologische) beschrijving van het windveld (de *input*). Aan de ene kant is de kennis over wind op zee beperkt omdat metingen niet altijd en overal beschikbaar zijn. Om de gaten op te vullen worden weermodellen gebruikt, maar ook die berusten op aannames en versimpelingen<sup>1</sup>. Daarom bestaat een groot deel van dit proefschrift uit vergelijkingen tussen beschikbare observaties, uitvoer van weermodellen en afgeleide klimatologische datasets.

Er is echter nog een bron van onzekerheid die samenhangt met de beschrijving van het windveld. In de praktijk wordt namelijk vaak een sterk versimpelde vorm van het windveld gebruikt. Zo wordt bijvoorbeeld aangenomen dat het windveld zich conformeert aan geïdealiseerde beschrijvingen, zoals het logaritmisch windprofiel dat de verandering van de wind met de hoogte dicteert. Ook wordt aangenomen dat het vóórkomen van die profielen goed beschreven kan worden met een statistische verdeling (de zogenaamde Weibull-verdeling). Het gebruikt daarvan heeft deels een historische oorsprong, omdat primitieve modellen niet toegerust waren voor een realistischere weergave van het windveld. Tegenwoordig dient het echter vooral nog om het aantal simulaties terug te dringen dat nodig is voor een volledige doorrekening van de klimatologie. **Hoofdstuk 1** gaat dieper in op deze aannames en de verschillende modellen.

Om na te gaan of die veelgebruikte aannames nog steeds acceptabel zijn of, anders gezegd, wat voor onzekerheid ze met zich meebrengen, wordt in **hoofdstuk 2** het concept 'afwijkende windverschijnselen' geïntroduceerd. Aan de hand van een vierjarige dataset van meetmast IJmuiden, bestaande uit observaties tot op driehonderd meter hoogte, wordt een aantal van dit soort abnormale gevallen in kaart gebracht voor de Noordzee. Hierbij valt met name het grote aantal 'low-level jets' op. Die situaties worden gekenmerkt doordat de windsnelheid afneemt

---

<sup>1</sup>Weermodellen zijn in principe vergelijkbaar met 'windparkmodellen', maar ze worden gebruikt op veel grotere tijd- en ruimteschalen en kunnen daardoor veel minder details (zoals windturbines) meenemen. Om verwarring te voorkomen zal ik het woord 'model' alleen gebruiken voor de kleinschalige modellen die voor windparken of individuele turbines worden gebruikt en het woord 'weermodel' voor, nouja, weermodellen.

met de hoogte, nadat het nabij het oppervlak een maximum bereikt heeft. Zulke omstandigheden komen regelmatig voor, met name in het voorjaar, en kunnen een aanzienlijke invloed hebben op zowel de belasting als de opbrengst van windturbines.

Andere afwijkende verschijnselen zijn snelle veranderingen in de windsnelheid of -richting in de tijd, of sterke veranderingen van de wind met de hoogte. De manier waarop deze extremen tot nu toe worden beschreven in verschillende standaarden is onrealistisch. Met nieuwe methodes om afwijkende omstandigheden in kaart te brengen zet dit proefschrift een belangrijke stap in de richting van een meer fysische benadering. Niet geheel verrassend blijken omstandigheden waarbij de opbouw van de atmosfeer nabij het oppervlak zeer gelaagd is een bepalende rol te spelen bij het vóórkomen van sterke windschering en draaiing van de wind met de hoogte.

**Hoofdstuk 3** bestaat uit een beoordeling van drie verschillende weermodellen. Deze modellen worden gebruikt voor het voorspellen van de wind en, in het verlengde daarvan, de energieopbrengst, maar ook bijvoorbeeld voor het samenstellen van windatlassen. De metingen van meetmast IJmuiden vormen een unieke dataset om te onderzoeken hoe goed de weermodellen de wind, maar bovendien ook de verticale opbouw van de atmosfeer nabij het oppervlak weergeven voor het Noordzeegebied. De modelevaluatie is gebaseerd op 30 verschillende dagen, gekozen aan de hand van een vernieuwende methode. Zodoende komen vele weersomstandigheden aan bod die, in ieder geval qua windsnelheid, een redelijk compleet beeld geven van de klimatologie.

De studie laat zien dat het gedrag van de drie modellen onderling erg vergelijkbaar is. Gemiddeld wijken de modellen minder dan  $0.5 \text{ m/s}$  af van de observaties; een typische fout is  $2 \text{ m/s}$ . Er zijn echter aanzienlijke verschillen tussen de uiteenlopende weersomstandigheden. Met name de momenten waarin de opbouw van de atmosfeer zeer gelaagd is vormen (nog steeds) een uitdaging voor de weermodellen. Juist deze omstandigheden zijn van groot belang voor wind-energietoepassingen.

**Hoofdstuk 4** gaat verder in op het verschijnsel van low-level jets, ofwel: laaggelegen windmaxima. Aanvullende waarnemingen van zogenaamde (wind)lidars en de klimatologische dataset ERA5 breiden de resultaten van hoofdstuk 2 uit voor een groter gebied. Beide datasets hebben sterke en zwakkere kanten en ik laat een aantal manieren zien om daarmee om te gaan. Een belangrijk resultaat is dat de detectie van windmaxima, en daarmee de klimatologie, erg sterk afhangt van de hoogte tot waarop gemeten wordt. In ERA5 lijken de windmaxima te zijn ‘uitgesmeerd’ en liggen ze hoger boven het oppervlak.

De klimatologische dataset ERA5 is een combinatie van observaties en modelsimulaties; gebreken komen vooral voort uit de beperkingen van het weermodel. De (onafhankelijke) observaties zijn juist met name beperkt in hun ruimtelijke dekking en tijdsperiode. Zelfs een dataset van vier jaar levert nog steeds een grillige klimatologie op van windmaxima, doordat het vóórkomen van windmaxima sterk kan variëren van jaar tot jaar. Door de twee gegevensbronnen te combineren blijkt het mogelijk een eenduidigere, representatievere klimatologie af te leiden.

Om beter te begrijpen hoe die laaggelegen windmaxima ontstaan geeft **hoofdstuk 5** een overzicht van de relevante literatuur. Hoewel er verschillende mechanismen bestaan die tot een windmaximum kunnen leiden, koppelen de meeste studies het verschijnsel aan het temperatuurcontrast tussen land en zee dat ontstaat op zonnige dagen wanneer de zee nog relatief koud is. Daarmee is het windmaximum sterk gerelateerd aan de zeewind, zogenaamde thermische lagedrukgebieden en een proces dat bekend staat als ‘geostrofe aanpassing’. Andere mecha-

nismen die windmaxima kunnen vormen hebben te maken met het reliëf of met zogenaamde inertiaaloscillaties. Die laatste worden ook vaak gekoppeld aan het 'nachtelijk windmaximum' boven land.

Vervolgens heb ik een aantal hypothesen geformuleerd om, aan de hand van de ERA5 data, te testen welke mechanismen verantwoordelijk kunnen zijn voor de windmaxima op de Noordzee. Het blijkt dat 'het zeewindmechanisme' inderdaad een belangrijke rol speelt. Daarnaast is er ook bewijs te vinden voor invloeden van het reliëf en inertiaaloscillaties. Hoewel hoofdstuk 5 in principe is toegespitst op het verschijnsel van windmaxima, dient het ook ter illustratie van de complexiteit van meteorologische processen in kustgebieden.

In de tussentijd zijn alle modellen uit hoofdstuk 3 gebruikt om klimatologische datasets (waaronder ERA5) samen te stellen. **Hoofdstuk 6** bestaat uit een beoordeling van deze drie datasets op basis van vier jaar aan data, in plaats van de 30 dagen in hoofdstuk 3. De nieuwe windatlas van het KNMI (DOWA) blijkt uitzonderlijk goed, met een typische fout van 1.5 m/s en bijna geen systematische afwijking in de windsnelheid. De modelleerstrategie die gebruikt is voor nieuwe Europese windatlas (NEWA) blijkt minder succesvol, in ieder geval in vergelijking met de metingen van meetmast IJmuiden.

Alle datasets hebben echter nog steeds grote moeite met de representatie van windrichting. Over het algemeen neemt de fout toe met de hoogte. Bovendien legt hoofdstuk 6 extra nadruk op afwijkende gevallen. In ERA5 worden die over het algemeen onderschat. De andere klimatologische datasets doen het beter, hoewel de timing niet altijd overeenkomt met de waarnemingen.

In de algemene discussie in **hoofdstuk 7** ga ik in op de vraag hoe om te gaan met afwijkende gevallen in windenergietoepassingen waarbij een volledige afspiegeling van de klimatologie belangrijk is. Verschillende ontwikkelingen duiden op een nieuwe status quo waarin data en kennis van de onzekerheid een veel grotere rol spelen bij het selecteren van de benodigde simulaties. De beoogde modelketen begint bij klimatologische datasets die steeds verder worden verfijnd. Dit proefschrift draagt bij aan de verwezenlijking van die visie, door uitgebreide en vernieuwende validatie van die klimatologische datasets, met uitgebreide aandacht voor afwijkende gevallen.

Qua begrip draagt het proefschrift bij doordat het de complexiteit van wind op zee toelicht. De situatie op de Noordzee is niet zo eenduidig als sommige andere kustgebieden waar veel windmaxima voorkomen: de temperatuurcontrasten zijn minder sterk en de interactie tussen de verschillende kustlijnen maakt dat de stroming vrij complex is. Ik pleit voor een betere representatie van de effecten van het temperatuurcontrast en het reliëf in windenergietoepassingen. Daardoor kunnen onzekerheden worden teruggedrongen die voortkomen uit een versimpelde weergave van de meteorologie.





# Preface

Up-to-date weather forecasts do not exist. Weather models depart from the most recent known state of the atmosphere. New observations keep coming in as the simulation progresses and consequently, any forecast is outdated before it completes.

A similar paradox presents itself to researchers in a rapidly evolving field such as energy meteorology. Ideally, new studies depart from and build upon the existing knowledge or prevailing theory; presumably, the scientist is aware of all the relevant previous work or at least of a synthesis thereof. In practice, the academic literature is so vast and so rapidly expanding that it is virtually impossible to keep up. Moreover, the researcher's perspective is largely determined by prior knowledge and experience, which continually evolves as (s)he spends more time on this planet; even in the process of writing itself. Thus, any research paper is outdated before it's even published...

The research described in this thesis crosses the borders between meteorology and wind energy engineering, and has been conducted and written from the perspective of a young graduate student in atmospheric sciences. Had an engineer written this thesis, it would have been completely different. Indeed, had I started this project knowing what I know now, it wouldn't be the same. And you as a reader, I'm sure you would bring yet other insights to the table.

The main body of this little book consists of academic journal publications – brief updates from my perspective on the subject matter. In a way, science journals are like the Twitter-feeds of academic in-crowds; publications are nothing but a quick update on researchers' latest insights for potentially interested colleagues. From time to time, these findings are collected and summarized in comprehensive review papers or scientific monographs – to stick with online vocabulary: Wikipedia gets updated. But even these texts are nothing but snapshots. Knowledge and understanding evolve in both individuals and the scientific community as a whole.

Online services such as Twitter and Wikipedia offer a suitable format for a continuous flow of information. Additionally, these platforms are very accessible for a broad audience. In a time when scepticism towards academia is voiced louder than ever, I think the academic publishing institutions can take inspiration from these popular online platforms. The accessibility of science in the broadest sense should have top priority among publishers. Visual elements play an important role in the communication of science. Consequently, I have rendered my figures in an unorthodox style that carries the suggestion of a hand-drawn sketch. I expect to encounter some resistance to this format. Indeed, I myself have experienced some hesitation as well. But over the years I've grown more and more fond of these figures, and I've gathered my thoughts on their use. At this point, it seems appropriate to briefly explain my motivation.

Most people will agree that these figures are visually attractive; they make the thesis pleasant to read, and accessible to a broad audience. More importantly, however, they convey a sense of uncertainty which, in my opinion, is more than justified: at various places I demonstrate that the results shown are sensitive to subjective choices in the data collection or analysis. Besides, although the lines are a bit wiggly, the precision that is lost is actually very small. I would even go so far as to turn the argument around and say that a very rigid figure style conveys a notion of precision that is often unjustified. Finally, this style reminds the reader that the figures were

made by a human being. If practising research is fundamentally subjective, I think it's better to make that perspective explicit.

Consequently, the last point I want to make is about the motivation for this study. As Yuval Harari put it, "*most scientific studies are funded because somebody believes they can help attain some political, economic, or religious goal.*" This bold statement seems to apply well to the present work: through the Topsector Knowledge and Innovation (TKI) offshore wind program, the Dutch government strategically invests in research (a.o.) that is expected to accelerate the *energy transition*. The ultimate goal of this particular project is to make offshore wind projects more cost-effective. My personal motivation is first and foremost that I want my work to contribute. A PhD in the renewable energy sector seemed to accommodate this wish. This does not, however, imply that I am biased. The discussion on the pros and cons of an energy transition is one I would gladly have over a beer, but as far as this thesis is concerned, I just hope that my findings support better decision making.

To wrap up, I want to stress that by walking the beaten track, we actively take part in maintaining the status quo. But the rules of academic practice are not set in stone, and continuous reflection (and perhaps some rebellion) is essential to move forward. With that in mind, I wish you a pleasant read, and I hope you'll be inspired.





# CHAPTER 1





## GENERAL INTRODUCTION

---

FOR THE TWENTY-FIRST CENTURY A FAR BIGGER GOAL IS NEEDED: MEETING THE HUMAN RIGHTS OF EVERY PERSON WITHIN THE MEANS OF OUR LIFE-GIVING PLANET.

- KATE RAWORTH, AUTHOR OF DOUGHNUT ECONOMICS

This thesis is part of a larger project which aims to identify, quantify, and, where possible, reduce uncertainties that hamper offshore wind farm development. One important source of uncertainty is the prescription of inflow fields for engineering models that are used to estimate turbine lifetime and annual energy production. Those uncertainties are addressed in this thesis.

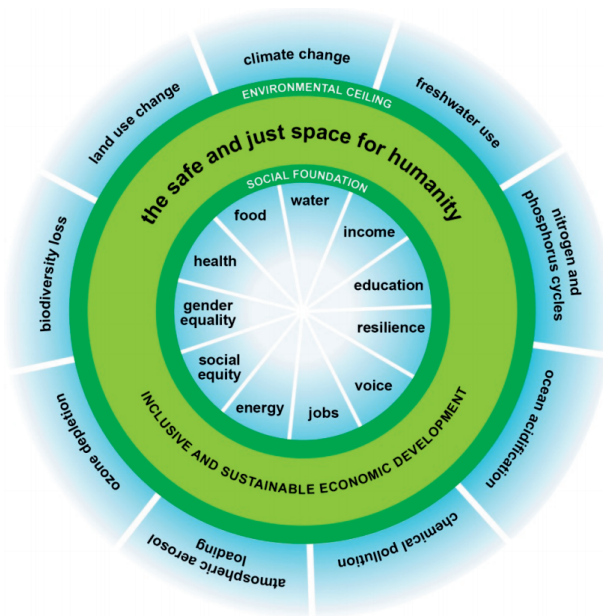
In the general introduction, I first describe the broader context of sustainability and renewable energy within which the current endeavour is rooted. After a historical note on recent developments in the field of wind energy, I provide a brief introduction to wind energy meteorology. This provides the basis for an examination of the aforementioned uncertainties and a formulation of the general problem statement that is to be addressed in subsequent chapters.

## 1.1 A quest for sustainability

The turn of the millennium has seen an increasing quest for sustainable alternatives to wasteful products and detrimental human activities. The motivation for this endeavour is nicely synthesized in figure 1.1: humanity is faced with the challenge of meeting the basic needs of a growing world population, while preserving the thriving ecosystems that provide for these basic needs in the first place. Structural exceedance of the planetary boundaries that constitute the environmental ceiling in figure 1.1 destabilizes the earth's ecosystems, possibly beyond repair, and is therefore unsustainable. For example, agricultural activities supply food and resources, but excessive exploitation of arable soil may cause land degradation. The initial boost in productivity pales into insignificance when compared to the hazardous infertility in the long run. Similar feedbacks apply to various aspects of our energy system.

Following the introduction of the planetary boundaries concept (Rockström *et al.*, 2009), research has focused on identifying and redefining these boundaries, assessing current status, projecting this framework onto more regional scales, and developing techniques and policies to help achieve that ultimate objective (Dearing *et al.*, 2014; Steffen *et al.*, 2015; Dorninger *et al.*, 2017; Butz *et al.*, 2018; O'Neill *et al.*, 2018). Alarmingly, multiple boundaries (biosphere integrity, climate change, biogeochemical flows, and land-system change) are already exceeded today (Steffen *et al.*, 2015).

An important notion is that these planetary boundaries should not be viewed in isolation as in fact, they are tightly linked (Rockström *et al.*, 2009). For example, traditional combustion of fossil fuels increases atmospheric CO<sub>2</sub> levels, leading to global warming and ocean acidification – and associated climate change and biodiversity loss. By-products of fuel combustion,



**Figure 1.1:** Basic human needs and planetary boundaries. Source: Raworth, A Safe and Just Space for Humanity. Oxfam discussion paper, February 2012.



especially in traffic, cause air pollution exactly in those places where most people suffer from it. Concerning the social foundations on the inside of the doughnut, (access to) natural fuel reserves have long been a source of political conflict and social inequality.

Renewable energy sources are envisioned to mitigate (some of) the adverse effects of traditional fossil fuels. Yet, while environmental concerns provide an important incentive for the energy transition today, it is worth noting that the Dutch electricity sector has undergone continuous change since the 1960s, and has been influenced by various technological and political developments (Verbong and Geels, 2007).

## 1.2 Tapping from earth's energy cycle

Every second, the earth intercepts about  $1.7 \times 10^{17}$  J of solar energy, roughly the equivalent of a Tsar Bomba – the largest nuclear weapon ever tested. Some of this radiation ( $\sim 30\%$ ) is immediately reflected back into space, but most of it is absorbed. The earth in turn also emits radiation. In fact, the outgoing energy at the top of the atmosphere must balance the incoming energy, or otherwise the planet would heat up quickly. Yet, before taking off to space again, the incoming solar radiation runs through a cascade of energy conversions and transport mechanisms.


Absorbed solar radiation is quickly converted to thermal energy, most of which is used to heat the adjacent air, or evaporate water into it. This heating and evaporation is not distributed uniformly over the globe. There are differences between day and night, summer and winter, pole and equator, ocean and land, clear and cloudy skies, mountains and valleys, urban and rural areas, and so on. Consequently, density and pressure gradients are created, resulting in the flow of air, which acts to dissolve the very gradients that drive it. If the sun were to shut down at once, the kinetic energy of the wind would cascade down to smaller scales, eventually dissipating into the molecular scales where the aggregated motion of molecules is no longer described as wind but rather as thermal energy. Back to the start, but with the important difference that the energy is now distributed uniformly over the globe<sup>1</sup>.

Recognizing the enormous amount of energy that is continuously pumped through our atmosphere, people set out to harvest some of this high-grade wind (and solar) energy to put it to use. On average, the entire human population consumes approximately  $1.7 \times 10^{13}$  J of energy per second, i.e. one ten-thousandth of all solar energy we receive. This puts the potential of wind and solar energy in perspective.

It is important to realise that energy 'consumption' does not actually decrease the total amount of energy. The energy is just converted to different forms, and in the process it becomes more dispersed and therefore less useful. A light bulb converts electric energy to light and heat, both of which quickly radiate away. Similarly, combustion releases energy stored in chemical bonds in the form of heat, which quickly dissipates. The fuel seems to disappear, but in fact it just transforms from a dense material into volatile gases that quickly find their way into the open sky<sup>2</sup>.

<sup>1</sup>For simplicity, I have left out other conversions, such as those related to clouds and rain.

<sup>2</sup>An alternative quantity that *can* be consumed is *exergy*. Exergy is the maximum amount of energy in a system that can be 'put to work'.



**Figure 1.2:** Evolution of wind turbine sizes and capacity since the 19th century. Source: Liebreich, M. "Breaking Clean". Bloomberg New Energy Finance Future of Energy Summit, London, 2017. Presentation available at <https://about.bnef.com/blog/michael-liebreich-state-industry-keynote-bnef-emea-summit-2017/>, accessed 6 June, 2019.

Thus, harvesting wind energy is nothing more than tapping into the natural flow of energy and diverting it through our devices, just like a water wheel extracts momentum from the downhill stream of water.

### 1.3 The evolution of *turbinus ventus*

Intriguing historical accounts of the origin of wind turbines can be found in Fleming and Probert (1984) and Dodge (2015). A review more focused on technological development can be found in Ackermann and Söder (2002). Trends in offshore wind power development are described in Snyder and Kaiser (2009), Rodrigues *et al.* (2015), and Vieira *et al.* (2019). A more technical review of its current status and challenges in Colmenar-Santos *et al.* (2016). An illustration of the recent evolution of wind turbines is shown in figure 1.2. What follows is a very brief summary to put the current endeavour in perspective.

Most likely, the first windmills appeared around 200 BC in Persia, but since then it took another fourteen centuries before records appeared of windmills in Western Europe and China around the 12th century AD. During the late medieval period windmills became widespread in Western Europe, culminating in the traditional Dutch windmills that were used to drain the Beemster Polder in the early 17th century. Shortly after, wind turbines were introduced in North America and in 1759 John Smeaton laid the scientific foundation for today's aerodynamic theory of wind turbines. In the centuries that followed, however, the number of traditional windmills declined due to the introduction of coal-fired steam engines, while smaller turbines with many (steel) blades appeared in the USA and Australia.

The twentieth century saw the emergence and advance of electricity-generating wind turbines, but progress was intermittent due to war, fluctuating oil, coal and gas prices, and the advent of nuclear energy. In the 1970s, concerns about the safety of nuclear energy and two subsequent oil crises gave rise to the first large-scale penetration of wind energy. In the last decades of the twentieth century the development of modern wind turbines continued. The now familiar three-blade design became popular, turbines increased in size and power, and costs decreased

– although government funding schemes were still needed to make wind energy competitive with traditional fuels.

In 1991 the world's first *offshore* wind farm, Vindeby, was installed. After a decade of building up experience, larger projects were undertaken in the 21st century. Offshore wind energy has many benefits compared to onshore, such as (generally) higher wind resources, more available space and less visual and acoustic impact, but it also comes with challenges such as the more difficult construction at sea and connection to the mainland. Still, the offshore wind turbine fleet has seen a steady increase and in 2018, the first offshore wind farms without subsidy were announced.

Over the past decades, wind energy has evolved very rapidly (figure 1.2). Turbines have grown, offshore wind farms have expanded, investments have risen, experience has been gained, technology has improved, and scientific underpinning has advanced. This trend is expected to continue, with even some serious exploratory studies seeking to provide the whole of Europe with offshore wind energy (Hajer and Pelzer, 2018; Sørensen and Larsen, 2018). These developments provide two important motivations for the present study:

1. Because investments have grown so large, a very small accuracy gain in estimates of power production, turbine lifetime, etc., can substantially reduce the total costs and financial risks of offshore wind projects.
2. Assumptions and simplifications that were previously justified may no longer be valid and need to be revisited.

Some of these assumptions are discussed in the next section.

## 1.4 A brief introduction to wind energy meteorology

The objective of this thesis in the broadest sense is to address uncertainties that enter (offshore) wind power projects through imperfect knowledge or representation of the 'external conditions', that is, predominantly the wind field<sup>3</sup>. Further specification of this objective requires an answer to the questions of where and how meteorological expertise finds its way into energy applications. Therefore, this section is concerned with a brief introduction to wind energy meteorology. I'll first introduce five 'branches' of wind energy meteorology. After a brief note on engineering models used for these applications, I'll describe the sources of their input data: meteorological observations and models. All this is necessary to finally come to a discussion of the uncertainty from external conditions. This will then naturally lead to the problem statement in the next section.

### 1.4.1 Five intertwined branches

Wind energy meteorology is a rapidly evolving branch of science, and an account of its status quo at any given moment is outdated as soon as written down. Below, I describe five 'branches' of energy meteorology which, in my opinion, are useful in a discussion of the subject. However, this distinction is rather subjective and the boundaries between these subdomains are not

---

<sup>3</sup>Another factor that is frequently referred to as external condition is the wave field. While there is obviously a connection between the two, this thesis focuses only on wind.

clearly defined; it's rather a continuum. Advances in computer power and progress in engineering may render some boundaries redundant within the next decade.

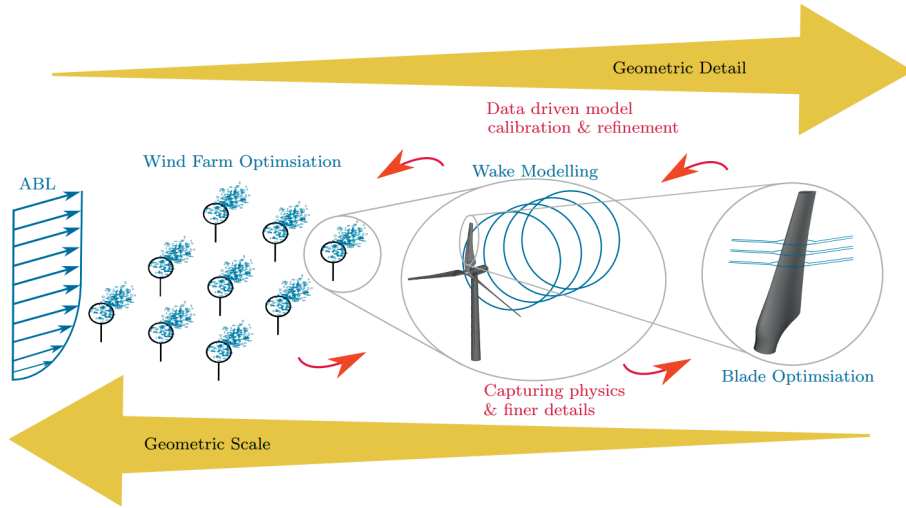
**Forecasting** is perhaps the most visible and broadly recognized occupation of (some) meteorologists, and it certainly also has its applications in wind power (Jung and Broadwater, 2014). Weather forecasts are important during construction and maintenance of offshore wind farms, for these activities require prolonged periods of calm weather (so-called weather windows). Moreover, in the abiding absence of efficient energy storage solutions, successful exploitation of existing wind power plants relies on accurate power predictions. Grid operators are faced with the challenge of balancing power supply and demand. Underestimated production may necessitate grid curtailment (some turbines have to be shut down to prevent overload), while a power shortage due to overestimation must be compensated by (high-tariff) back-up generators.

**Resource assessment.** A good characterization of a potential wind energy site is very important for the viability of such a project. Knowledge of the wind climate can provide an answer to questions such as: *What is the expected annual energy production?* and *Which turbines are most suitable for this location?* The process of estimating site parameters is known as *resource assessment*. The wind climate comprises the *long-term* characteristics of the wind field, and since year-to-year variability can be substantial, it requires exceedingly long data records. A complete site description includes information on wind speed and direction, turbulence, and atmospheric stability. Turbulence is the term used for small or rapid fluctuations in space and time, which are the result of whirls or eddies. Atmospheric stability is a related concept. In an unstable atmosphere, turbulent mixing is enhanced by buoyancy effects, similar to boiling water or a lava-lamp. In a stable atmosphere, density stratification suppresses turbulent mixing, just like oil and water always tend to return to a layered configuration after being stirred up.

**Load assessment.** Design and certification of wind turbines is based on the expected (wind) loads on the structure. The IEC guidelines (IEC, 2005, 2009; Burton *et al.*, 2011) distinguish two types of structural loads: ultimate loads due to extreme events leading to instant failure, and fatigue loads due to smaller but repetitive impacts on the structure, adding up over the lifetime of a turbine and eventually leading to cracks and other forms of damage. Several load cases are defined for both types of loads. Each load case describes a wind inflow and a certain turbine state (e.g. in operation, stalled, ...). The wind inflow is described in terms of wind speed, wind shear, wind veer, and turbulence intensity. For some special cases, an (extreme) wind gust is prescribed.

**Wind farm optimization.** When multiple turbines are placed in a grid configuration, they influence each other by reducing the wind speed in their wakes. Since the distance to the 'next' turbine is dependent on the wind direction, a suitable wind farm layout for a given wind climate can substantially reduce wake losses. The exercise of finding the layout with minimal wake loss is often called *wind farm optimization*. In addition to information on wind speed and turbulence intensity (or atmospheric stability), which both influence the length of the wakes, wind farm optimization also accounts for wind direction.

**Feedbacks on the atmosphere.** An important, but often overlooked aspect of large-scale wind power implementation is that it may also feed back on the wind field itself, either directly



**Figure 1.3:** Illustration of the different levels of detail and scales of engineering models used for wind energy applications. Source: Hewitt *et al.* (2018).

or indirectly (Miller *et al.*, 2011; Vautard *et al.*, 2014). Recent studies have focused on wake effects influencing downstream wind farms (Nygaard, 2014; Lundquist *et al.*, 2019). The Dutch national weather service is currently implementing a wind farm parametrization to account for the effects of wind farms in their weather model<sup>4</sup>.

## 1.4.2 Engineering models: from power formula to digital twins

All the applications listed above rely on models to convert wind information to power, structural loads, optimal wind farm lay-out, or vice versa. To this end, a suite of engineering models has been, and is being, developed, ranging from simple and idealized models to very complex and computationally expensive simulations (figure 1.3). An overview of various methods to describe or model the interaction between the wind and the turbines is given in Sanderse *et al.* (2011), and a worthwhile future perspective in Hewitt *et al.* (2018). Since this thesis is not concerned with the details of these methods, but rather with the specification of their inflow fields, an elementary introduction suffices.

The simplest approach to convert wind data to power is perhaps the traditional power formula:

$$P = \frac{1}{2} \pi \rho R^2 v^3 c_p$$

where  $P$  is the power,  $\rho$  is the air density,  $R$  is the rotor radius,  $v$  is the wind speed and  $c_p$  is a power coefficient. This equation multiplies the kinetic energy of the wind per unit volume,  $E = \frac{1}{2} \rho v^2$ , with the volume of air that passes through the rotor per unit time,  $\pi R^2 v / t$  (note that  $E = Pt$ ). The power coefficient accounts for the efficiency with which the turbine converts the kinetic energy of the wind to electricity. This formula can be used to construct a simple *power curve*, which shows the generated power as a function of wind speed (figure 1.4)<sup>5</sup>. Notice

<sup>4</sup>“The main innovations in the [DOWA] project are (...) to include wind farm (cluster) wake effects in mesoscale weather model HARMONIE.” <https://www.knmi.nl/research/observations-data-technology/projects/dowa>, accessed 8 July, 2019

<sup>5</sup>In practice, an empirical power curve is often preferred over this theoretical version.

that the power formula assumes that the wind speed is constant in space and time, the density is uniform, and it only accounts for power, not loads.

At the other end of the spectrum, sophisticated CFD (computational fluid dynamics) models explicitly simulate the interaction between the wind and the wind turbine. In addition to power yields, these models can compute the structural loads on various components of the turbines. While these models could in principle account for a wide range of natural variability in the inflow fields, such simulations are severely limited by computational constraints. Therefore, the number of simulations must be reduced to obtain an acceptable balance between accuracy and efficiency.

Figure 1.3 shows that there is a two-way interaction between large-scale models with little detail and small-scale models with high levels of detail. The most obvious direction is wind input propagating down the model cascade. The way up requires some additional explanation. By learning from the interaction between the wind and turbines in the detailed simulations, simplified models can be derived that emulate the basic features of this interaction. These emulators can then be used in the larger-scale models as ‘parametrizations’: models embedded in other models to account for the effects of unresolved processes. Wake models are a great example. Wind farm simulations need to be repeated many times for various grid configurations, and therefore high-resolution simulations are out of the question. Simpler representations of the wake can be validated and tuned with high-resolution simulations, though. Englberger and Dörnbrack (2018), for example, designed a cheap method to derive realistic turbulent inflow fields, from a small number of expensive simulations. Kanev *et al.* (2018) combine a database of structural loads obtained from a relatively limited number of expensive simulations with a cheaper model for wind farm optimization.




Figure 1.4: Example power curve, showing the power generated by a wind turbine as function of the wind speed.

### 1.4.3 Meteorological data sources: from local observations to global simulations (and back)

The engineering models discussed so far take the wind, or an abstraction of it, as inflow condition, and so far I've assumed that this knowledge is simply available. To fully appreciate the source of this information, I shortly discuss atmospheric observations. Thereafter, I describe the some important aspects of meteorological models and their applications.

**The art of observation.** Intuitively, the most reliable way to characterize the wind climate at a given site is to go out there and measure it, but often this is easier said than done. Measuring at sea, for example, is quite challenging and expensive, especially when the measurements are to be taken at the hub height of potential wind turbines (nowadays typically above 100 m). Traditional wind measurements are influenced by the mast on which the instruments are mounted; newer measurements techniques based on lidar can replace observations of wind speed and direction, but when it comes to turbulence, this technique is still in its infancy. After dealing with all these challenges, a measurement is still only valid for a single location. Moreover, the wind *climate* comprises the long-term characteristics of the wind fields, so ideally, the measurements should be carried out over extended periods of time. Traditionally, a period of 30 years has been agreed upon as a standard for reliable characterization of the (wind) climate. This is impractical, to say the least.

**The art of modelling.** Weather models may be used for forecasting, but they can also be used to fill in the gaps between sparse observations. Today's weather forecasts are produced using sophisticated computer codes which are the product of a fascinating scientific success story (Bauer *et al.*, 2015). On a fundamental level, these codes, or *numerical weather prediction models*, implement the equations governing fluid flow on a rotating sphere. Basically, they divide the atmosphere into a large number of grid cells or grid boxes. By comparing the wind, temperature, pressure and humidity between adjacent grid cells, they can predict how these quantities will evolve in the near future. Adding more boxes generally results in a higher accuracy, but it also requires more computations. Some processes that cannot be represented explicitly, such as turbulence and microphysics, are accounted for through parametrizations.

**Data assimilation.** Initialization of a weather model requires information about the present wind, temperature, pressure and humidity in each grid cell. This poses a challenge, because such a dense measurement grid simply does not exist. Only adjusting the grid boxes for which observations are available, or interpolating between these measurements is not an option, because that would violate the laws of physics (e.g. increasing the wind speed in one box without decreasing it somewhere else is equivalent to adding excess energy to the atmosphere). Therefore, modellers use a 'trick' to get the model as close as possible to all the available observations, while maintaining physical consistency. Simply put, they run yesterday's forecast several times in hindsight, compare the outcome with the observations, and choose the best match for the current state of the atmosphere. This trick is called *data assimilation* and there are several sophisticated mathematical techniques to carry it out reliably. The initial state of the model thus obtained is called the *operational analysis*. Together, data-assimilation and the actual forecast form a very demanding task (according to Bauer *et al.* 2015 it is similar to simulating the human brain) which is carried out several times per day on some of the worlds largest supercomputers.

**From forecast to reanalysis.** As outlined above, data-assimilation is a useful technique to infer the state of the atmosphere at locations where no measurements are available. Thus, an archive of historical operational analyses constitutes the best-guessed state of the atmosphere as it evolved over time. These archives can be used to extract site parameters at locations where measurements are not available. Since weather models and data-assimilation techniques are continuously developed and improved, the operational archive is quite inconsistent. Therefore, several institutes have used recent model versions to construct a consistent ‘archive’ in retrospect. This procedure is called reanalysis. The ERA-interim dataset (Dee *et al.*, 2011), developed by the European Centre for Medium-range Weather Forecasts (ECMWF), represented the state-of-the-art for many years, and has been widely used in wind energy applications. Its successor, ERA5 (Olauson, 2018), became available in the midst of this research project.

**The weather model cascade.** Similar to figure 1.3, there is a two-way interaction between coarse, large-scale weather models and detailed, small-scale simulations. ECMWF nowadays produces global forecasts on a horizontal grid cell size of approximately 15 km. For ERA-interim, the horizontal grid spacing was about 80 km. Enhanced computer power enabled an increase in resolution to about 30 km for ERA5, but still, this is too large to represent some small-scale features (let alone turbulent fluctuations). Therefore, (national) weather services run similar models with smaller domains, enabling finer grids and a more detailed representation of the flow. Traditionally, the models used for this first refinement step are called limited-area models or mesoscale models. These go up to a horizontal grid spacing of about 1 km. Further refinement can be achieved with large-eddy simulation (LES) models, which resolve most of the turbulence explicitly. The model domains that are feasible with LES are large enough to simulate a wind farm, but way too small for a comprehensive wind atlas. The engineering models discussed previously can be interpreted as an extension of the meteorological model chain, and LES is at the interface. Currently, wind turbine parametrizations are also finding their way into mesoscale weather models (Fitch *et al.*, 2012; Volker *et al.*, 2015).


#### 1.4.4 Introducing uncertainty

There are two main sources of uncertainty concerning the prescription of wind fields in wind energy applications: (1) uncertainties in the data sources themselves and (2) uncertainties that are introduced by the way in which these data are summarized, represented, extrapolated, or otherwise employed.

Both observations and model data have their limitations. A well-known uncertainty is the accuracy of measurement instruments. More difficult to estimate is the uncertainty due to the representativeness of a measurement, which is inherently valid at a single point (or a bounded volume in the case of lidar, for example). Model data suffer from inaccuracies due to the limited grid resolution, and the imperfect representation of physical processes. The non-linear behaviour of the atmosphere also sets a fundamental limit on our knowledge of the exact state of the atmosphere at any given place and time (Lorenz, 1963). Quantification of these uncertainties is a never-ending exercise, as meteorological models and measurement methods continuously evolve and improve.

The second type of uncertainty is related to the way in which the meteorological data is employed. Working with raw, multidimensional data is often impractical or undesirable. In resource assessment, quick comparison between different sites is easiest if the wind climate can





**Figure 1.5:** Illustration of summary statistics. A full rendering of the time/height variability of four years of wind speed observations is shown in Panel B. Panel A shows the mean wind speed profile that is obtained by taking the time-average of B. Two commonly used wind profile functions are fitted to the mean profile. Panel D is a time series of wind speed observations at a single height (115 m, approximately hub-height). Panel C is a histogram of those wind speed observations. A Weibull distribution is fitted to the data.

be described in terms of one or maybe a few summary statistics. This is illustrated in figure 1.5. Four years of wind speed observations up to a height of 315 m are plotted in panel B. Although very condensed, this figure encapsulates a full four-year climatology of wind speed. However, it is not easy to process this information. There are two straightforward ways to simplify the data. The first is time averaging. Panel A shows the mean wind speed observations at each measurement height. The average wind profile can reasonably be described by two common parametrizations that are often used in practice. The first is the empirical wind profile power-law,  $u(z) = u_{ref}(z/z_{ref})^\alpha$ , and the second is the (neutral) logarithmic wind profile,  $u(z) = u_* / \kappa \ln(z/z_0)$ , which is founded on physical considerations. Both parametrizations are based on the premise that the variation of the wind speed with height follows a characteristic shape that can be described by a simple mathematical function: the wind speed ( $u$ ) increases monotonically with height ( $z$ ), more strongly so near the surface and approaching an asymptotic limit aloft. The rate with which the wind shear decreases is dictated by the empirical shear exponent ( $\alpha$ ) in the power-law, and by the surface roughness ( $z_0$ ) in the logarithmic wind profile. The asymptotic limit scales with reference wind speed  $u_{ref}$ , or with  $u_*$ , the friction velocity.  $\kappa$  is the Von Karman constant.


The other simplification is to take advantage of the high degree of correlation in the vertical direction, and consider only a single measurement height. The resulting time-series though, visualized in figure 1.5D, is still quite incomprehensible. Further simplification is achieved by rendering the data in the form of a histogram (figure 1.5C). It turns out that the shape of this histogram can be described quite well by the two-parameter Weibull distribution, which is characterized by a scale parameter and a shape parameter. Together with the characteristic surface roughness length of the logarithmic wind profile, or a typical shear exponent of the wind profile power-law, the two parameters of the Weibull distribution basically encapsulate

the entire variability of the wind field at a given site. Therefore, they have become ubiquitous in resource assessment studies.

The summary statistics described above can in principle be converted to the real quantities of interest, such as expected annual energy production or estimated turbine lifetime. This is where the engineering models come in. The simplest estimate of annual energy production is obtained by combining the Weibull distribution with the power curve illustrated in figure 1.4. This is very efficient, but any information on the vertical structure of the wind speed profile is discarded in the process. This was acceptable for early wind turbines, but with modern turbines reaching rotor diameters of over 100 m this simplification is no longer appropriate. Figure 1.5A shows that the difference over a typical rotor span of 100 m *on average* may exceed  $2 \text{ m s}^{-1}$ , which can make a big difference in the total energy production (see figure 1.4; the impact is large in the steep part of the power curve) and also enhances the structural loads.

More accurate load (and power) assessment thus requires detailed inflow fields including wind shear and also turbulence. To still exploit the efficiency of the Weibull distribution for time extrapolation, the fatigue load cases combine each wind speed bin of the histogram in figure 1.5C with a wind profile parametrization and additionally a mathematical description of turbulence. Illustrated in figure 1.6, this procedure represents a substantial improvement over the assumption of a uniform wind field, but the variability is still restricted to the fixed form of the parametrizations employed. Consequently, the use of idealized inflow fields is an important source of uncertainty.

Wind farm optimization encounters a similar conflict between efficiency and accuracy. Because wind direction is very important for wake effects, the Weibull distribution is replaced by a joint probability of wind speed and direction, frequently referred to as ‘the wind rose’. Often, the wind rose only includes data at a single height. Depending on the complexity of the engineering models employed (e.g. two- or three-dimensional), additional information can be re-introduced using idealized descriptions. Again, by using the wind rose probability distribution for time extrapolation, some of the variability hidden in the distribution is inevitably lost. The question is thus, how much information is discarded, and is this (still) acceptable?




**Figure 1.6:** Conceptual illustration of the inflow prescription for fatigue load cases. One load simulation is performed for each  $2 \text{ m s}^{-1}$  bin of the Weibull distribution. Parametrizations are used to account for vertical wind shear and turbulence.

## 1.5 Problem statement: the limitations of a spherical cow

A BIOLOGIST, A CHEMIST, AND A PHYSICIST ARE TAKING A WALK THROUGH THE COUNTRY WHEN THEY COME UPON A COW. THE BIOLOGIST THINKS FOR A SECOND AND THEN DECLARES, "I KNOW WHAT THAT IS. THAT'S BOS TAURUS." THE CHEMIST LOOKS FOR A SECOND AND THEN SAYS, "IT'S JUST A CARBON-BASED LIFE FORM, APPROXIMATELY 75% WATER, AND THE REMAINING 25% TRACE ELEMENTS." THE PHYSICIST STARES BLANKLY FOR A SECOND AND THEN SAYS, "WELL.... I GUESS WE COULD APPROXIMATE IT AS A SPHERE."

Uncertainties in external conditions enter energy applications through imperfect knowledge of the (wind) climate, and through the way in which the available information is processed. Observations are limited in space and time, and simple methods or sophisticated models are relied upon to infer the wind climate based on the available information. But also sophisticated models are limited in their grid resolution and their representation of physical processes, and even if a very accurate record of the past thirty years exists, one can only assume that the next thirty years will be more or less similar. This last assumption is equivalent to postulating that the wind climate does not change, which, in face of global warming, has become a topic of renewed scientific debate (e.g. Pryor and Barthelmie, 2010, and citing articles).

Summary statistics and idealized parametrizations facilitate spatial comparison for wind resource assessment and computational savings for load simulations, among others. Yet, whilst idealized descriptions adequately capture the average wind conditions, their application to instantaneous wind fields may introduce uncertainties because of oversimplification. This is illustrated in figure 1.7, which clearly shows that some observed wind profiles do not follow the typical shape that is dictated by the logarithmic wind profile. And indeed, some of the underlying assumptions are violated in practice. For example, the logarithmic wind profile assumes that the surface is horizontally homogeneous, but this is clearly not the case near the coasts. The logarithmic approach also assumes that the turbulent stress is independent of height, which may be valid in a shallow layer near the surface, but does certainly not hold above, say, 100 m.



**Figure 1.7:** A. Mean wind profile over four years of observations at an offshore site. Black dots are observations, full and dashed lines are two common parametrizations. B. A selection of instantaneous wind profiles from the same dataset.

Furthermore, it assumes stationary conditions, which is also not justified in the vicinity of a (coastal) discontinuity. And finally, the (neutral) logarithmic wind profile as presented in the previous section neglects atmospheric stability, which can vary substantially in practice.

The objective of this thesis is to illuminate both sources of uncertainties: the (in)accuracy of model performance, the impact of space and time limitations in observations, and the frequency and severity with which common assumptions are violated. This will enable researchers and industrial users to evaluate whether their methods are still acceptable, and provide guidance to account for, or even reduce these uncertainties.

Part of this objective is achieved through cross-validation between models and observations. The other part is achieved through the concept of *anomalous wind events*, which I introduce in chapter 2. Anomalous events is a term I invented to account for those conditions that fall outside the range of variability normally considered, or for which common assumptions are violated. The academic literature and wind atlases are flooded with Weibull parameters and logarithmic wind profiles. Contemporary practice provides a climatology of ‘regular’ wind conditions. In contrast, this thesis focuses on the *not* normal phenomena: a climatology of *anomalous* events.

The overall objective inspires a number of specific research questions that will be addressed in the subsequent chapters:

1. What is the impact of space and time limitations on the representativeness of observations?
2. What is the typical performance of weather models and reanalysis datasets?
3. Which anomalous wind events can we define and how can we identify and characterize them?
4. How well are these anomalous events represented in reanalysis data?
5. What are the mechanisms that govern wind in coastal areas and specifically, anomalous events?
6. How can anomalous events be accounted for in wind energy applications?

It is nearly impossible to compile a complete list of anomalous events, because there will always be assumptions that can be violated. In fact, most of the thesis focuses on a single phenomenon: the low-level jet. This might be the most relevant and appealing event, but for that reason it is also a perfect exemplar. In a sense, this thesis can be seen as a framework, or methodology, to deal with anomalous events in practice. Moreover, it is a perspective, an unorthodox view that hopefully inspires some critical reflection on the ‘external’ conditions that are often taken for granted. And finally, it is an invitation to appreciate the richness and complexity of wind in coastal areas, and how far we’ve come to understand and exploit it.

## 1.6 Outline of the thesis

The contents of the following chapters are schematically illustrated in figure 1.8. The research data primarily concern the (Southern) North Sea, although the methods employed and the mechanisms discussed are of course universally applicable. The thesis follows a clear chro-

nological structure that is intertwined with the rapid developments in the field of wind energy. During the research, several datasets were released, and then used in later stages of the research.

In chapter 2, I further introduce the concept of anomalous events. I consider how these events can best be defined, and continue to characterize them. In retrospect, some of these events are better described as extreme rather than anomalous, but that does not make their characterization less valuable. This chapter is based on the met mast IJmuiden (MMIJ) dataset, which already made its first appearance in figure 1.5. Its location with respect to the Dutch coast and its temporal coverage are illustrated in figure 1.8.

The MMIJ dataset is used again in chapter 3, but the focus is quite different. In this chapter, I evaluate the performance of three numerical weather prediction models: IFS, HARMONIE, and WRF. Each of these models is linked to the production of a brand new reanalysis dataset, but these were not yet available at the time. Thus, chapter 3 can be seen as a first acquaintance with these models. Alongside the model validation, some useful methods for case selection and performance evaluation are introduced in this chapter.

Having assessed the ‘general’ performance of these three models, the next interesting question is whether they are able to truthfully represent the anomalous events introduced in chapter 2. In the meantime, the ERA5 reanalysis dataset (based on the IFS) had become available, as well as additional measurements from various locations (as shown in figure 1.8). This provided a perfect opportunity to extend the results from chapter 2 in a spatial sense. One of the most interesting results from chapter 2 is the abundance of low-level jets at the MMIJ location. To maintain a clear focus, low-level jets are the only anomalous event considered in chapter 4. However, a similar approach can be pursued for the other events. In that sense, chapter 4 can be seen as a methodology to develop a true spatial climatology of anomalous wind events.

So far, all chapters focussed on obtaining a validated climatological characterization of wind in general and anomalous wind events in particular, but the mechanisms that govern these events were not discussed in detail. This question is addressed in chapter 5, where I present a literature survey of relevant physical mechanisms that govern wind in coastal areas. Again, the main focus is on low-level jets, but the discussion is very relevant for, if not tightly linked to, other (anomalous) events. In the second part of chapter 5, I try to attribute the low-level jets identified in the ERA5 data in chapter 4 to a physical governing mechanism.

As mentioned above, new data was released as I was writing the chapters of this thesis. These new datasets provide many opportunities to build upon the previous chapters. Moreover, in those chapters I have zoomed in on low-level jets to maintain a clear focus, but a similar path can be followed to construct or extend the climatologies of other anomalous events. Chapter 6 incorporates many results to evaluate the representation of wind *and* anomalous wind events in all three reanalysis datasets. For brevity, only MMIJ data are used. Chapter 4 shows that this is by far the most suitable dataset for model evaluation.

The general introduction featured a description of contemporary strategies to estimate, among others, annual energy production and structural loads. In the general discussion in chapter 7, finally, I will come back to the start, and address the question *How may the lessons from this thesis contribute to the development of next-generation load- and power assessment procedures?* To answer this question, I will draw upon the lessons learned from the previous chapters, but also from other contributions to the scientific literature, popular science, real-world examples, documentaries, and of course, imagination.

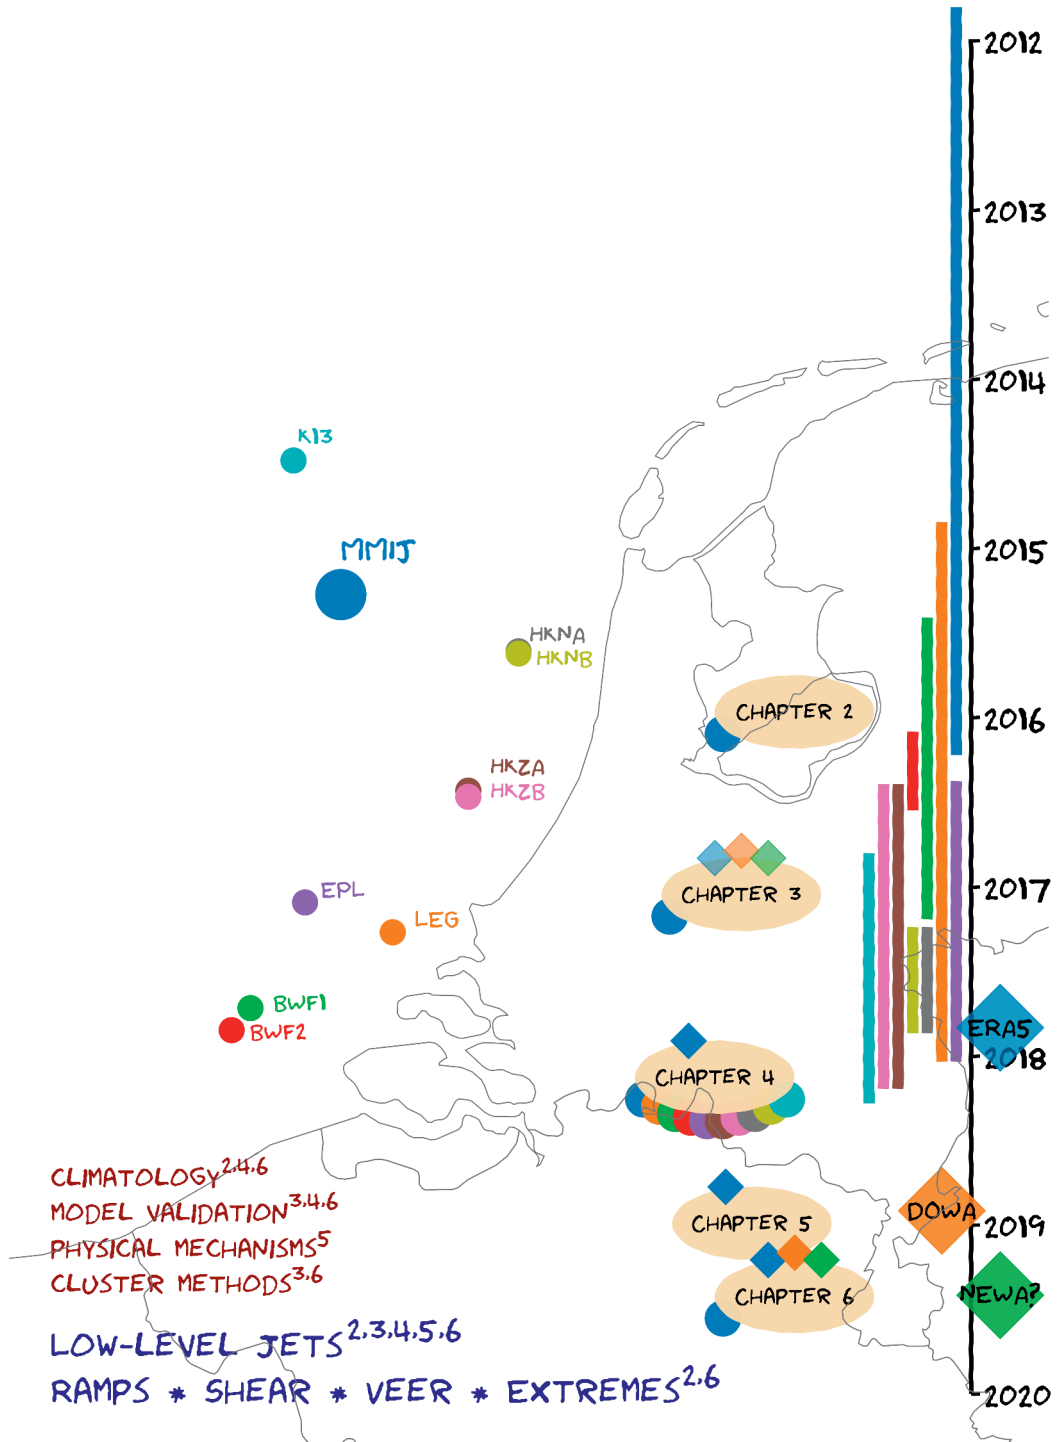



Figure 1.8: Schematic illustration of connection between the chapters in this thesis, and how this relates to the availability of various measurement and reanalysis datasets.

# CHAPTER 2







## AN OBSERVATIONAL CLIMATOLOGY OF ANOMAL- OUS WIND EVENTS OVER THE NORTH SEA

---

THREE STATISTICIANS GO HUNTING FOR DEER. THEY SPOT ONE OFF IN THE DISTANCE.  
THE FIRST ONE SHOOTS ABOUT A METRE TOO HIGH, THE SECOND ONE, ABOUT A METRE  
TOO LOW. THE THIRD ONE YELLS, 'WE GOT IT!'



In this chapter, the concept of anomalous wind events is introduced. Inflow fields for wind energy engineering models are often highly idealized. While summary statistics and standardized wind profiles provide a reasonable climatology of regular wind conditions, they introduce uncertainties with respect to less common or atypical wind fields. Within the framework of anomalous events, these wind fields are systematically identified, characterized and eventually understood.

Four years of high-quality observations, gathered up to 300 m, are analysed to investigate the wind climate at “met mast IJmuiden”. Six anomalous wind events are identified and described: low-level jets, extreme wind speeds, shear, veer, turbulence and wind ramps. In addition, typical weather conditions that favour their formation are identified. This provides the first clues towards increased understanding of these phenomena.

This chapter is published as:


Peter C. Kalverla, Gert-Jan Steeneveld, Reinder J. Ronda and Albert. A. M. Holtslag (2017). An observational climatology of anomalous wind events at offshore metemast IJmuiden (North Sea). *Journal of Wind Engineering and Industrial Aerodynamics*, 165:86–99.

## 2.1 Introduction

Wind energy has been established as an outstanding resource of renewable energy, with over 433 GW of installed capacity worldwide by the end of 2015 (REN21, 2016). Onshore wind power has now become cost-competitive with traditional sources, and offshore wind power follows this track. Offshore wind farms have the advantages of stronger and steadier winds than onshore, more available space, less visual distortion, and less noise pollution, among others. The worldwide installed offshore wind power was estimated at 12 GW by the end of 2015 (REN21, 2016), with the largest share in Europe with 7.5 GW divided over 76 offshore wind farms (Rodrigues *et al.*, 2015). A further increase to 40 GW is imperative to reach the European Union's target of 20 % renewable energy by 2020 and follow-up targets thereafter (EWEA, 2013). Despite recent advances and numerous assets of offshore wind, continued growth is still hampered by relatively high installation and maintenance costs and uncertainties in average wind on short and longer time scales (Breton and Moe, 2009; EWEA, 2013; Burton *et al.*, 2011; Rodrigues *et al.*, 2015; Carroll *et al.*, 2016; Kuik *et al.*, 2016; Mérigaud and Ringwood, 2016). Maintenance of offshore wind turbines is carried out with large, expensive crane vessels that can only be put to sea in favourable weather concerning wind and significant wave height. This can lead to long waiting times and extended periods of power loss (Martini *et al.*, 2017). It is thus of utmost importance that turbine failure and maintenance are minimized, which is challenging since the offshore environment is more demanding than its onshore counterpart due to the stronger winds, additional wave and current loads, and corrosion by salty water.

To minimize turbine failure, prototype turbines are extensively tested with simulation software, in which the wind field is 'translated' to actual forces on the wind turbine (Burton *et al.*, 2011; Vorpahl *et al.*, 2013). Obviously, the input wind field determines the resulting load assessment, but contemporary guidelines (IEC, 2005, 2009; Burton *et al.*, 2011) still rely on idealized extrapolation methods. Figure 2.1a conceptually illustrates this approach. A frequency distribution of wind speed at one height is combined with a wind profile parametrization (e.g. power law, logarithmic wind profile) to obtain a suite of 'representative' wind profiles. Subsequently, the estimated loads resulting from these profiles are extrapolated to the turbine lifetime ( $\sim 20$  yr) using the frequency distribution. Many authors have argued that these methods cannot adequately capture all variability that is observed in the field (Lange *et al.*, 2004; Gryning *et al.*, 2007; Kettle, 2014; Nunalee and Basu, 2014; Optis *et al.*, 2014). For example, they cannot account for atmospheric stability, vertical variations of wind direction, or low-level jets. Therefore, the current methodology leads to substantial uncertainty in the estimated loads.

Apart from the regular wind conditions, the IEC also specifies load cases to account for ultimate loads resulting from extreme wind speeds, extreme gusts, and extreme shear (IEC, 2005; Burton *et al.*, 2011). These wind events are described in a rather idealized fashion, e.g. extreme shear is simulated with a mathematical function that tilts the wind profile with a fixed factor. This approach is physically unrealistic and therefore leads to additional uncertainties in the load estimates. Finally, although the current methodology does account for short-term wind fluctuations, it does not consider temporal transitions in the mean wind profile. In the context of power forecasting, these so-called 'wind ramps' have received considerable attention recently (Gallego *et al.*, 2014; Gallego-Castillo *et al.*, 2015; Kiviluoma *et al.*, 2016), but to our knowledge wind ramps in relation to wind farm site- or load assessment have not been assessed so far.



**Figure 2.1:** Conceptual figure to illustrate the contemporary guidelines for wind representation. A probability distribution of wind speeds is sampled and for each sample a simple profile parametrization is used to extrapolate the wind speed to other altitudes.

Thus, the realism of the wind climate as represented in load assessment studies is open for further improvement, at least for a number of conditions that are currently un- or ill-defined. In this paper, we use a unique observational dataset obtained with a tall tower and lidar observations up to 300 m height, situated 85 km off the Dutch coast, to investigate the wind climate at the North Sea and obtain a reference climatology of these ‘anomalous wind events’ (AWEs, particularly: low-level jets, extreme wind speed, shear, gusts and ramps)<sup>1</sup>. Accordingly, we formulated the following research questions:

1. What are typical values and temporal characteristics of low-level jets, extreme wind speed, shear, gusts and ramps?
2. How are wind, wave and other weather conditions related to these anomalous wind events?
3. Can we identify temporal overlap between the anomalous wind events?

The first question deals with objectively defining each of the anomalous wind events, and is in fact one of the unique aspects of this study, enabling straightforward comparison with other datasets, and model validation. The typical conditions of question 2 will help us to characterize the AWEs. An analysis of potential temporal overlap between events is included in order to obtain reliable statistics considering that multiple extremes may occur simultaneously. The resulting climatology of anomalous wind events characterizes the wind variability at the site, and is particularly well-suited as a reference for future studies. As such, this may serve as a first step towards more accurate description of the anomalous wind events in load assessment studies.

<sup>1</sup>Although it would be possible to define more events, such as thunderstorm downbursts and waterspouts, these are very hard to measure and are therefore not part of this study. A useful reference on the latter may be found in Dotzek *et al.* (2010).

## 2.2 Background

This section provides a brief overview of each of the identified anomalous wind events.

**Low-level jets.** A low-level jet (LLJ) is a maximum in the vertical wind speed profile, relatively close to the surface (Baas *et al.*, 2009; Nunalee and Basu, 2014). When located in the rotor plane, the increased wind speed might boost the power production, but can also lead to enhanced wind loads on the turbines. E.g. Gutierrez *et al.* (2016) used high frequency tower observations up to 200 m and found that the enhanced wind shear during LLJs leads to higher structural loads on turbines. Park *et al.* (2014) performed a similar study using wind input from large-eddy simulations and found comparable results. Therefore, accurate LLJ representation is critical in atmospheric models, in load assessment models, and in site assessment studies. Mirocha *et al.* (2016) recently evaluated the representation of LLJs in the weather research and forecasting model (WRF, Skamarock and Klemp, 2008) and found that “The considerable root mean square error and bias values, even among the ‘best’ performing simulations, underscore the need for improved simulation capabilities for the prediction of near-surface winds during LLJ conditions.”

Several physical mechanisms for the LLJ have been hypothesized (Stull, 1988; Stensrud, 1996). Over land, LLJs can be explained by the evening transitions, when turbulent mixing decays (Baas *et al.*, 2009; Van de Wiel *et al.*, 2010). The reduced friction leads to acceleration of the wind at the top of the ‘decoupled’ layer. This jet follows an inertial oscillation with a period equal to the Coriolis period at a given latitude ( $\sim 15$  h for the Southern North Sea). A jet formed over land and subsequently advected offshore, may well be able to persist there. Alternatively, a similar mechanism can occur when warm air mass is advected over a cold (sea) surface (Smedman *et al.*, 1989; Dörenkämper *et al.*, 2015). The bottom cooling leads to a stably stratified layer which can also ‘decouple’ from the flow aloft. Other authors have linked low-level jets to baroclinity and strong horizontal temperature gradients (Zemba and Friehe, 1987; Burk and Thompson, 1996; Andreas *et al.*, 2000). These jets have a direction more or less parallel to the coast (Stensrud, 1996). Finally, downward mixing of higher wind speeds aloft may lead to short-lived jet-like structures.

**Extreme wind speeds.** Extreme wind speeds are often studied in a statistical sense, e.g. turbine load assessment typically considers the estimated 50-year wind extreme (Burton *et al.*, 2011). For class I turbines, the reference 50-year 10-minute-averaged wind extreme is set to  $50 \text{ m s}^{-1}$  (IEC, 2005). A comprehensive overview of statistical extreme value theory application to wind speeds is given in Palutikof *et al.* (1999). An often encountered challenge is that datasets are too short, since extreme value theory typically selects only one event per year. Therefore, Palutikof *et al.* (1999) discuss several methods to make optimal use of the available data. At the North Sea, wind extremes are always associated with extratropical cyclones. An interesting meteorological treatment of many storm events is given in Lamb and Frydendahl (1991). Recently, Christakos *et al.* (2016) presented a detailed investigation of cyclone Xaver that passed over the North Sea in December 2013. WRF was used to study the wind power density, and was found to perform rather well. Wind speeds exceeded  $25 \text{ m s}^{-1}$  for 30 hours during this extreme event. At the same time, wind power generation reached a new record due to the strong winds associated with the cyclone just before and after this peak wind speed.

**Wind shear and veer.** (Vertical) shear is represented by the change of wind with height. Naturally, the flow must adjust to vanishing wind at the bottom boundary, and therefore most shear is observed close to the surface, where it can lead to considerable differential loading on wind turbines (Burton *et al.*, 2011). It is thus important to correctly account for wind shear in offshore wind systems.

Shear is sometimes expressed as a ratio of the wind speed between two altitudes. Alternatively, a power-law of the form  $u/u_{ref} = (z/z_{ref})^x$  is sometimes used to describe the relation between wind speed and altitude (Farrugia, 2003; Storm and Basu, 2010). In that case, the exponent  $x$  determines the steepness of the wind profile. In the IEC guidelines  $x$  is fixed at a value of 0.14 for offshore conditions, but in fact  $x$  encompasses information about atmospheric stability and flow characteristics and may thus assume a range of values (Irwin, 1979). Obviously, neither of these methods is applicable to the previously described low-level jets. Additionally, both are unable to represent directional changes.

Since wind is a vector, we can distinguish between speed shear and direction shear. Speed shear is the change in wind vector magnitude with height. The change of wind direction with height is called wind veer. Veering with height can occur due to different mechanisms (Brown *et al.*, 2005). Under idealized conditions, Ekman theory predicts 45° wind veering in the atmospheric boundary layer, which in reality reduces to 15–40° depending on stability. Additionally, horizontal temperature gradients lead to a thermal wind that veers (backs) with height if warm (cold) air is advected. The strongest veering is to be expected when the two processes act together. Finally, mesoscale phenomena like the sea breeze may lead to additional rotation of the wind vector with height.

Over 25 years ago, Hollingsworth (1994) already identified serious deficiencies in the ECMWF model representation of backing winds with height during periods of cold air advection. Ten years later, Brown *et al.* (2005) evaluated the representation of wind turning in the ERA interim dataset, and found that the model performance is still quite poor, with biases up to 30° under strong baroclinic conditions. They found similar results for higher resolution ECMWF forecasts and for the MetOffice model.

**Turbulence and wind gusts.** Turbulence is an important atmospheric process characterized by chaotic motions on relatively short length and time scales, covering the range between viscous processes on the molecular scale and the mean flow on the scales of the atmospheric boundary layer (Stull, 1988). It is the main driver for transport of heat, moisture and momentum, and as such turbulence governs the environment in which turbines have to operate. This is relevant for wind energy because the rapid motions lead to differential loads on the turbine, and to power fluctuations. Turbulence is responsible for wind gusts that are superimposed on the mean wind. Therefore, studies focusing on wind energy often report the so-called turbulence intensity, i.e. the standard deviation of the wind speed, normalized by the mean wind speed. Similarly, gust factors can be computed as the ratio of the maximum wind speed to the mean wind. Usually, the gust factor and turbulence intensity are closely related (Bardal and Sætran, 2016).

**Wind ramps.** Wind ramps are rapid changes in the wind speed or direction in time (within one to several hours). In that sense, they are different from the previously discussed AWEs that were all treated as instantaneous profiles. Sudden fluctuations in the wind speed can result

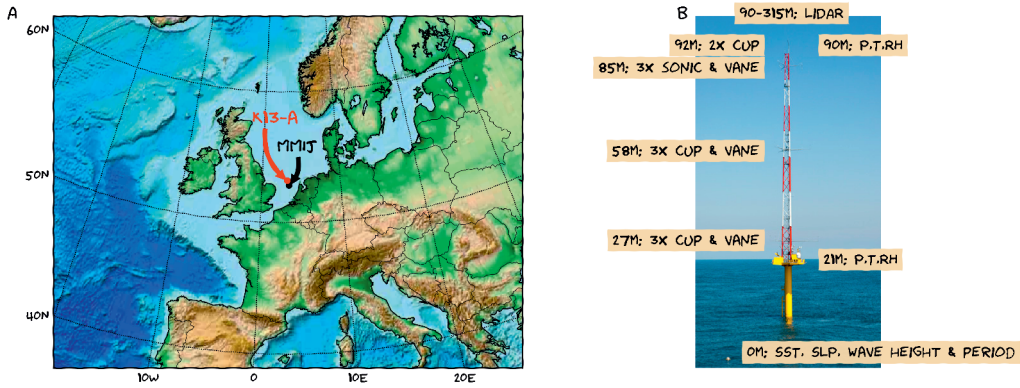
in power fluctuations, especially in the case of large wind farms. Therefore, wind ramps have received considerable attention in the field of forecasting lately (Gallego *et al.*, 2014; Gallego-Castillo *et al.*, 2015; Kiviluoma *et al.*, 2016). Sudden changes in wind direction require adjustment of the rotor orientation. Wind ramps can be governed by different mechanisms, e.g. the passage of fronts, formation of thunderstorms due to moist convection, perturbations in stability leading to downward momentum mixing, or the onset of local circulations like a sea breeze. Since these processes act on various time scales and produce wind fluctuations of different magnitude, many different ramp criteria have been defined (see Gallego-Castillo *et al.*, 2015). The time scales range from 5 minutes to 6 hours, and the power fluctuations vary between 10 % and 75 %, expressed as percentage of the rated power that was considered in each case study. Gallego-Castillo *et al.* (2015) encourage that further research is conducted to explain ramp occurrence and its representation in numerical weather prediction models. In this context, it is most relevant to report temporal wind variations in our climatology.

## 2.3 Methodology

### 2.3.1 Description of the dataset

The observations for this study have been gathered at the meteorological mast IJmuiden (MMIJ; Werkhoven and Verhoef, 2012), situated in the North Sea approximately 85 km off the Dutch coast (52.85°N, 3.44°E, figure 2.2). The water depth at the site is about 28 m. In the four-year period from 2012 to 2015, a 90 m mast has been maintained and equipped with several instruments. On the platform (21 m) and the top of the mast (90 m), standard meteorological quantities pressure, temperature and relative humidity were observed (Vaisala PTB210 and HMP155). At 27 and 58 m, three booms with Thies first class cup anemometers and wind vanes were mounted to the tower. Moreover, three Metek USA-1 sonic anemometers and three more wind vanes were mounted at 85 m. Two additional cup anemometers were mounted at the top of the tower (92 m). A Zephir 300 continuous-wave lidar recorded wind speed and direction at 25 m intervals between 90 and 315 m. Finally, a Triaxis wave buoy recorded wave and current characteristics, as well as water temperature and sea level pressure. Compared to other masts at the North and Baltic Seas, MMIJ stands out because of its measurement height (up to 315 m), its distance to shore, and because it has three booms at each measurement altitude, allowing for the elimination of mast disturbances. For comparison, the Høvsøre tower (Peña *et al.*, 2015) is located at the coastline and is 116 m tall. In addition, the widely used FINO-1,2,3 towers reach up to a height of 100 m. Moreover, FINO-1 and -2 are situated closer to the coast (40 km) and in the vicinity of wind farms. Therefore MMIJ offers the opportunity to investigate undisturbed wind patterns for higher altitudes than before.

Data post-processing has been performed by the Energy Centre of the Netherlands (ECN) before uploading to <http://www.meteomastijmuiden.nl/data/>. All data are provided as 10-minute averages, standard deviations, and min/max values, except for the buoy data, which are available at hourly intervals. Care was taken to minimize the effects of tower distortion, using the differences between sensors at the different booms. This procedure is outlined in the instrumentation report (Werkhoven and Verhoef, 2012). Using the one level of overlap between the lidar and cup-anemometer measurements (90 m), ECN also performed a validation of the lidar data (Maureira Poveda and Wouters, 2015). Good agreement was found between mean wind speed and direction (mean bias of 1 %). However, lidar-based turbulence intensities appear to be



**Figure 2.2:** A: Map of northwest Europe, with the location of the IJmuiden meteomast indicated. Source: ETOPO. B: a photograph of the measurement site, with brief annotations explaining the instrumentation. P, T and RH are pressure, temperature and relative humidity, respectively. SLP and SST are sea-level pressure and sea-surface temperature.

largely overestimated at low wind speeds. Peña *et al.* (2009) also evaluated a Zephir lidar against observation from cup anemometers, and conversely, they find an underestimation of turbulence intensity, which they attribute to attenuation of standard deviation observations from the lidar. They also explain that the Zephir lidar relies on a thermal wind sensor to determine the sign of the wind direction. This sometimes leads to an  $180^\circ$  error in the wind direction. For the two years of data that were analysed by ECN, this occurred for 3.6 % of the records (Maureira Poveda and Wouters, 2015). Indeed, several instances of this  $180^\circ$  deviation were present in the lidar data, as well as some spikes in wind speed. These records were carefully removed from the dataset ( $\sim 4\%$ ).

Several instruments were subject to technical problems during extended periods of time. The wave buoy has been offline between February 2013 and April, 2014. In April, 2015, most instruments mounted to the mast were disconnected. The 85 m booms did not record from January to mid-August, 2013. To obtain wind profiles that are more or less uniformly spaced, we averaged the 85, 90, and 92 m measurements. Performing these adjustments resulted in a data coverage of 84 % concerning all wind observations. Considering hub-height wind speed, data availability appears to be 95 %, but with fewer measurements in spring. Still, this is a satisfactory data availability for our research purpose, especially considering the remote location of the tower. Because much of the buoy data from MMIJ was missing, we used observations from a nearby oil rig, K13-a, for the classification of waves (see figure 2.2). This station is part of the Dutch network of automated weather stations. Sea state observations from K13-a are publicly available via the Dutch ‘Rijkswaterstaat’ via `live.waterbase.nl`. For the episodes with joint data availability, a significant correlation ( $r^2=0.95$ ) was found between K13-a and MMIJ for significant wave height and period. Because of the superior data availability and long standing record, we decided to use K13-a wave data rather than MMIJ observations. Unfortunately, sea surface temperature was not available at K13-a.

Note that Holtslag *et al.* (2015) used one year of the MMIJ dataset (2012) to validate Monin-Obukhov similarity theory for this offshore site and later to assess wind turbine fatigue loads resulting from the joint probability of stability and wind speed (Holtslag *et al.*, 2016). We extend

their work by including all four years of observations, and moreover, we focus on anomalous wind events that are typically not well described by Monin-Obukhov similarity theory.

### 2.3.2 Definition of anomalous wind events

In this section we objectively define the criteria for the anomalous wind events as introduced in the previous section. To identify low-level jets, the criterion by Baas *et al.* (2009) is often used nowadays. They define a LLJ as the lowest maximum that is at least  $2 \text{ m s}^{-1}$  and 25 % stronger (=20 % fall-off) than the subsequent minimum. Earlier studies have used slightly different fall-off criteria, and the choice for any criterion is quite arbitrary. We decided to use Baas' absolute fall-off criterion of  $2 \text{ m s}^{-1}$ , but to perform additional sensitivity analysis to the effects of using a  $0.5 \text{ m s}^{-1}$  higher and lower threshold.

To quantify extreme wind speed, we first apply statistical extreme value analysis to the 115 m wind speed, which corresponds to the hub height of the Leanwind 8 MW reference wind turbine (Desmond *et al.*, 2016). As the dataset is rather short for such a procedure, we apply the method of independent storms, as outlined in Palutikof *et al.* (1999). In this approach, independent storms are separated by periods of low winds (so-called lulls). This guarantees independence of the data, while at the same time allowing for several extreme events per year. The wind speed maxima  $x$  of the independent storms are then sampled, sorted, and plotting positions are calculated according to

$$F(x_m) = \frac{m - 0.44}{N + 12} \quad (2.1)$$


where  $m$  is the sorted index and  $N$  is the total number of events. Following Harris (1999), we can then calculate the (Gumbel) reduced variate as

$$y = -\ln(-r \ln(F(x_m))) \quad (2.2)$$

where  $r$  is the average amount of storms per year. Cook (1985) found that a reliable estimate of the 50-year wind extreme is made for  $r = 10$ , leading us to define a lull threshold of  $21.3 \text{ m s}^{-1}$ . To calculate the standard error of the estimate, we used a Monte Carlo method in which we repeatedly drew random samples from the extreme value distribution defined by the parameters that we first identified. Subsequently, these samples were fed back to the extreme value algorithm to obtain a realistic range of values for the two parameters.

For analyzing wind shear and veer, we assume that the turbine always aligns the rotor axis with the mean wind speed at hub height. Accordingly, we can split the wind measurements, and thus the shear, in streamwise (or: longitudinal) and spanwise (or: lateral/normal) components, the latter resulting from changes in wind direction. Note that it is not possible to directly convert between veering angle and normal wind shear, since the same wind shear vector will have a different veering angle for different wind speeds. We prefer this method, in which we express the directional shear in a normal wind component rather than an angle, because in this way the amount of shear can be compared without considering the absolute wind speed, and moreover, because it is more readily translated to forces on the turbine. Figure 2.3 clearly illustrates that for a veering wind profile, the normal wind component is positive below hub height and negative above, which seems consistent with the clockwise rotation of the rotor. Because the vertical extent of MMIJ is not a universal property, we will calculate and report the accumulated shear





**Figure 2.3:** A. Hodograph of a typical Ekman wind profile that veers with height in the boundary layer. The rotated coordinate system is defined by  $s$ , aligned with the wind at hub height and  $n$ , perpendicular to  $s$  and positive to the left. B. Typical vertical profiles of the rotated wind components. The shading represents three layers spanning most of the rotor plane (see section 2.3).

in six layers of equal depth (50 m) as:

$$S_s = \frac{1}{50} \sum_{z=1}^{50} |s_z - s_{z-1}| \Delta z \quad (2.3)$$

$$S_n = \frac{1}{50} \sum_{z=1}^{50} |n_z - n_{z-1}| \Delta z \quad (2.4)$$

where  $s$  and  $n$  denote the stream-wise (longitudinal) and normal (span-wise/lateral) components of the wind vector, respectively, as in a natural coordinate system. Because the boundaries of these six layers do not coincide with the observation heights, we linearly interpolated between the wind observations (assuming that the wind vanishes at the surface). This appeared to be the best method, considering that for most wind profiles, the layer where a logarithmic wind profile would be appropriate was below the lowest measurement height. Additionally, a logarithmic wind profile is not suitable for anomalous events like LLJs. Using polynomial interpolation we found that more artefacts were introduced than realism was gained. Taking the magnitude of the vertical differences in  $s$  and  $n$  in equation (2.3) ensures that ‘negative shear’ is also accumulated. This is relevant, because otherwise the accumulated shear might equal zero if, for example, a low-level jet would occur in the middle of a layer. To the authors’ best knowledge this approach has not been reported elsewhere so far.

For the classification of extreme turbulence and associated wind gusts, we use the turbulence intensity,  $TI = \sigma_u / \bar{u}$ , as is common practice in the wind energy sector. In a preliminary study we found a very strong relation between  $TI$  and gust factor, similar to Wieringa (1973) and Burton *et al.* (2011). Therefore, and in the absence of higher frequency data that would provide us with more specific information on gusts, we decided to focus only on levels of turbulence as a proxy for wind gusts. Because the lidar values of the variance are not very reliable, we only use the data from the cup anemometers and the sonic. To investigate the impact of dividing by the mean wind speed in the definition of  $TI$ , we will also discuss the results if only the standard deviation of the wind,  $\sigma_u$ , would be used.

Previous studies that focused on identifying (power) ramps used thresholds ranging from 10% to 75% of the rated power of a wind farm that should be exceeded within a specified time frame of 5 min to 6 h (Gallego-Castillo *et al.*, 2015). Note that these studies used time series of produced

wind power rather than time series of wind speed itself, because their focus was on evaluating power forecasts. The diversity in ramp definitions demonstrates that there is no general consensus on the relevance of certain power fluctuations. Alternatively, we focus on fluctuations of the wind itself. Additionally, we do not choose arbitrary fixed thresholds, because in the context of a climatological paper, it makes more sense to convert the criterion again. We evaluate the difference between two consecutive wind profiles and analyse the 5- and 95-percentiles. This yields the 10-minute ramp for each observation height. We take the maximum ramp for each record and use the 95-percentile as threshold for classification as extreme ramp event. Similarly, we define the 30- and 60-minute ramps by taking the difference between observations separated by 30 and 60 minutes of time, respectively. We perform the analysis for both wind speed and direction.

### 2.3.3 Weather conditions

Having defined the anomalous wind events, our objective is to assess their temporal variability and to relate them to other weather variables. This is important to help us understand the mechanisms that control these events. Additionally, if it is possible to relate anomalous wind events to typical weather patterns, this provides very valuable information, for example in forecasting, where the exact manifestation of anomalous wind events may not be predicted, but can be estimated based on other (better-represented or larger-scale) parameters. Likewise, it may provide a starting point for ‘smart sampling’ methods, where one wants to represent most of the variability with as few as possible simulations.

The analysis pointed out that the following five variables (in addition to seasonality) provide valuable information about the typical conditions for each AWE: wind direction, mean wind speed, pressure, stability and wave characteristics. The other parameters in the dataset added little value to the analysis and are therefore not discussed. We use the pressure observed at 21 m because data availability is higher than for sea-level pressure that was measured with the wave buoy. Stability is determined by evaluating the bulk Richardson number between the tower top and the platform:

$$Ri_B = \frac{\frac{g}{\theta_v} \frac{\Delta\theta_v}{\Delta z_\theta}}{\left(\frac{\Delta u}{\Delta z_u}\right)^2 + \left(\frac{\Delta v}{\Delta z_v}\right)^2} \quad (2.5)$$

where  $g$  is the gravitational constant, and  $\theta_v$  is the virtual potential temperature. The sign of  $Ri_B$  depends on the density stratification: if density increases with height (virtual potential temperature decreases), the buoyancy term (numerator) is negative and the atmosphere is said to be unstably stratified, and turbulence will be enhanced. If the density decreases with height,  $Ri_B$  is positive and turbulence is suppressed by buoyancy. The shear term (denominator) is by definition positive as it will always lead to dissipative mixing. If  $Ri_B$  is close to zero, the atmosphere is said to be neutrally stratified. We refrain from defining a range of neutral stability, because this bulk layer is relatively thick and elevated, and typical surface-layer characteristics are hence not applicable. It would also be possible to evaluate  $Ri_B$  between the platform and the sea surface, which may be more appropriate to reflect surface layer properties. Unfortunately, missing buoy data prevented us from doing so. However, for the period where both layers are available, we find that they agree very well. Additionally, we estimated the Obukhov length  $L$  with an iterative procedure explained in Van Wijk *et al.* (1990). The results were qualitatively similar to the climatology of the bulk Richardson number, but we prefer the latter because

Monin-Obukhov similarity theory is only valid in the surface-layer (which can be very shallow offshore), and moreover, because we are especially interested in the conditions for which this theory is not applicable (e.g. low-level jets).

The wave characteristics will be summarized in the Charnock parameter as defined by Donkers *et al.* (2011):

$$\alpha = \frac{H}{T\sqrt{gd}} \quad (2.6)$$

where  $H$  is the wave height,  $T$  is the wave period,  $g$  is the gravitational constant and  $d$  is the water depth. The parameter thus encompasses information about wave height as well as wave steepness. The value of the Charnock parameter represents the strength of the coupling relation between wind profile shape and wave height. As such, it is used to correct the logarithmic wind profile for roughness over sea. A high value of the Charnock parameter implies a strong coupling, i.e. a direct interaction or equilibrium between wind and waves. Lower values indicate that the wave field is dominated by older waves or swell, and accordingly the surface is smoother.

To analyse the anomalous winds in relation to these conditions, we will consider a batch of relevant events and classify them in a boolean framework. For LLJs this is straightforward, as we used a fixed threshold which must be exceeded. Having no physical considerations at hand that lead to a natural choice of threshold values, we used a different criterion for all other AWEs. Instead of analysing how often a certain threshold is exceeded, we will report the threshold that is exceeded exactly 5 % of the time (i.e. the 95-percentile). This threshold yields a suitable balance between the ‘extremeness’ of the events, and the amount of observations needed for a climatological analysis.

For extreme wind speed, we will find exactly twelve quantile values, one for each altitude. However, for our boolean framework we need a criterion that allows us to classify the entire wind profile as one extreme wind event, rather than identifying extreme events for each measurement height. Therefore, we will use the maximum wind speed in each record. For wind shear and veer, we will use the three layers between 50–200 m that span most of the rotor plane of a current-day wind turbine. We exclude the lowest layer (0–50 m) as it would dominate the complete climatology. For turbulence we take the maximum of the profile, and for wind ramps we use the 115 m wind as it is close to the hub height of current-day wind turbines.

The relation between the occurrence of each AWE and each of the external conditions was studied by dividing the external condition into bins and counting the number of AWEs in each bin. In an absolute sense, these counts often resemble the distribution of the underlying external condition. In a relative sense, i.e. the counts of AWEs for each bin divided by the total number of observations in that bin, some distinctive patterns are revealed (see section 2.4).

## 2.4 Results

### 2.4.1 Background meteorological conditions

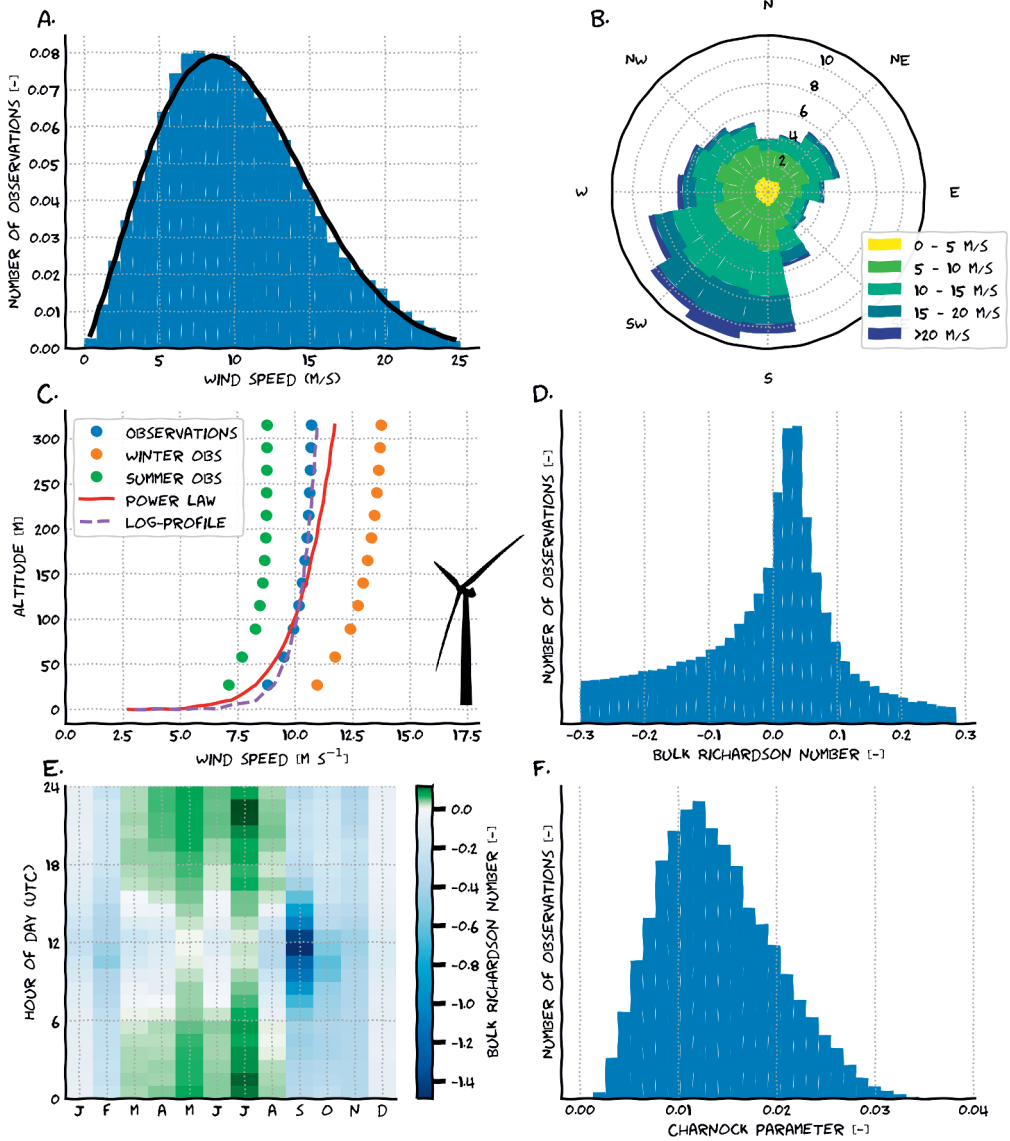
Before we proceed with the classification of anomalous wind events, it is important to understand the meteorological background on which the AWE climatology is projected. The mean

wind speed at hub height (115 m) over all seasons amounts to  $10.2 \text{ m s}^{-1}$ . The 2-parameter Weibull distribution, that was fitted to the data using a maximum likelihood function, provides a reasonable fit to the measurements (figure 2.4A). Its shape and scale parameters amount to 2.17 and  $11.49 \text{ m s}^{-1}$ , respectively. The wind climate is dominated by the frequent passage of mid-latitude cyclones, resulting in relatively strong winds from the southwest (figure 2.4B). The mean wind profile, obtained by time-averaging of the wind speed observations at all altitudes demonstrates that most wind shear is confined to the lowest layer, even below the lowest observation height (figure 2.4C). In winter, the mean wind profile is about  $5 \text{ m s}^{-1}$  stronger than in summer and more wind shear is present in the rotor layer of the wind turbines. figure 2.4 also illustrates Weibull-weighted mean wind profiles obtained by extrapolation of the 115m wind using the power law  $u/u_{ref} = (z/z_{ref})^x$  (with fixed exponent of  $x = 0.14$ ) and a neutral logarithmic wind profile  $u = u_{ref} \ln(z/z_0) / \ln(z_{ref}/z_0)$  with  $z_0 = 0.0002 \text{ m}$  (see Holtslag, 1984; Foken, 2006; Burton *et al.*, 2011). A few points were added in the lowest layer to compare the behaviour of the two surface-layer parametrizations. It appears that the power law overestimates the amount of wind shear throughout the wind profile. The logarithmic wind profile performs better.

The bulk Richardson number is positive for 41 % of the available observations, and negative for the remaining 59 %, i.e. a 2:3 ratio (figure 2.4D). Even though more observations are missing in spring relative to other seasons, this appears to be a robust figure. If we first calculate monthly means (as such ruling out seasonal variability in the average), the total time fraction of stable and unstable conditions remains nearly unchanged. The magnitude of  $Ri_B$  is inversely proportional to the bulk wind shear, which leads to high absolute values during calm conditions. For the unstable regime, the magnitude of  $Ri_B$  is somewhat higher because the boundary layer is typically well-mixed and hence there is less wind shear than under stable conditions.

Stability over the North Sea has previously been studied by Van Wijk *et al.* (1990) and Coelingh *et al.* (1992, 1996). They used seawater temperature observations from ships, and wind and air temperature observations from various stations, among which also the K13-a oil rig. They estimated the Obukhov length (see Van Wijk *et al.*, 1990), and found predominantly unstable conditions (48–64 % versus 23–36 % stable). Additionally, a range of neutral stability was identified. As already mentioned, we refrained from defining a range of neutral stability profiles for  $Ri_B$ , since the thresholds would be too arbitrary. However, the distribution of  $Ri_B$  (absolute values are higher on the unstable side of the figure) indicates that more stable than unstable profiles would be classified as neutral. Hence our findings correspond quite well to previous results. This provides confidence that  $Ri_B$  in the 21–90 m bulk layer is a suitable proxy for stability.

$Ri_B$  exhibits a pronounced annual cycle with more stable conditions in spring and summer and unstable conditions in autumn and winter (figure 2.4E). Spring and summer are characterised by an evident diurnal cycle, with more neutral conditions in the afternoon. In September afternoons the Richardson number is substantially lower than the rest of the day, while in winter the diurnal pattern is almost absent. The seasonal cycle can be related to sea surface temperature, which drops to  $7^\circ \text{C}$  in spring leading to cooling at the bottom boundary and increases to  $17^\circ \text{C}$  in autumn. In contrast, there is no straightforward explanation as to why we observe a diurnal cycle, since the thermal inertia of the sea is too large to follow the diurnal cycle of solar radiation. A plausible explanation seems to be the advection of onshore boundary-layer structures. Dörenkämper *et al.* (2015) have shown that coastal structures can still be observed over 100 km off the coast under stable conditions, because of limited vertical mixing.



**Figure 2.4:** A: 2-parameter Weibull fit to the 115 m wind speed frequency distribution. B: Wind rose of 115 m wind speed and direction in the boundary layer. C: Mean wind profile from observations, and by using two Weibull-weighted extrapolation methods (see text). D: Histogram of  $Ri_B$  evaluated between the top of the mast and the platform. The range was limited to 0.6. E: Median of bulk Richardson numbers grouped by hour of day and month of year. The colour map was constructed such, that positive values are green and negative values are blue, and their gradients are optimized for their respective ranges. F: Histogram representing the distribution of the Charnock parameter.

The observed Charnock parameter varies around a mean value of 0.017 (figure 2.4F), with higher values in winter than in summer, but no clear diurnal variability (not shown). This value is quite high compared to previous studies (e.g. Garratt, 1977; Fairall *et al.*, 2003)), but within the range reported by Smith (1988). Oost *et al.* (2001) studied the Charnock parameter for a near-coastal site in the North Sea, and estimating from their Figure 1, they find an average value of 0.018 ( $\ln(\alpha) = -4$ ). In our study, the Charnock parameter has a high linear correlation (0.87) with the 115 m wind speed. This indicates that for stronger winds, the wave field is mostly ‘wind-dominated’, whereas for smaller wind speeds, swell is relatively more important.

## 2.4.2 Characterization of anomalous wind events

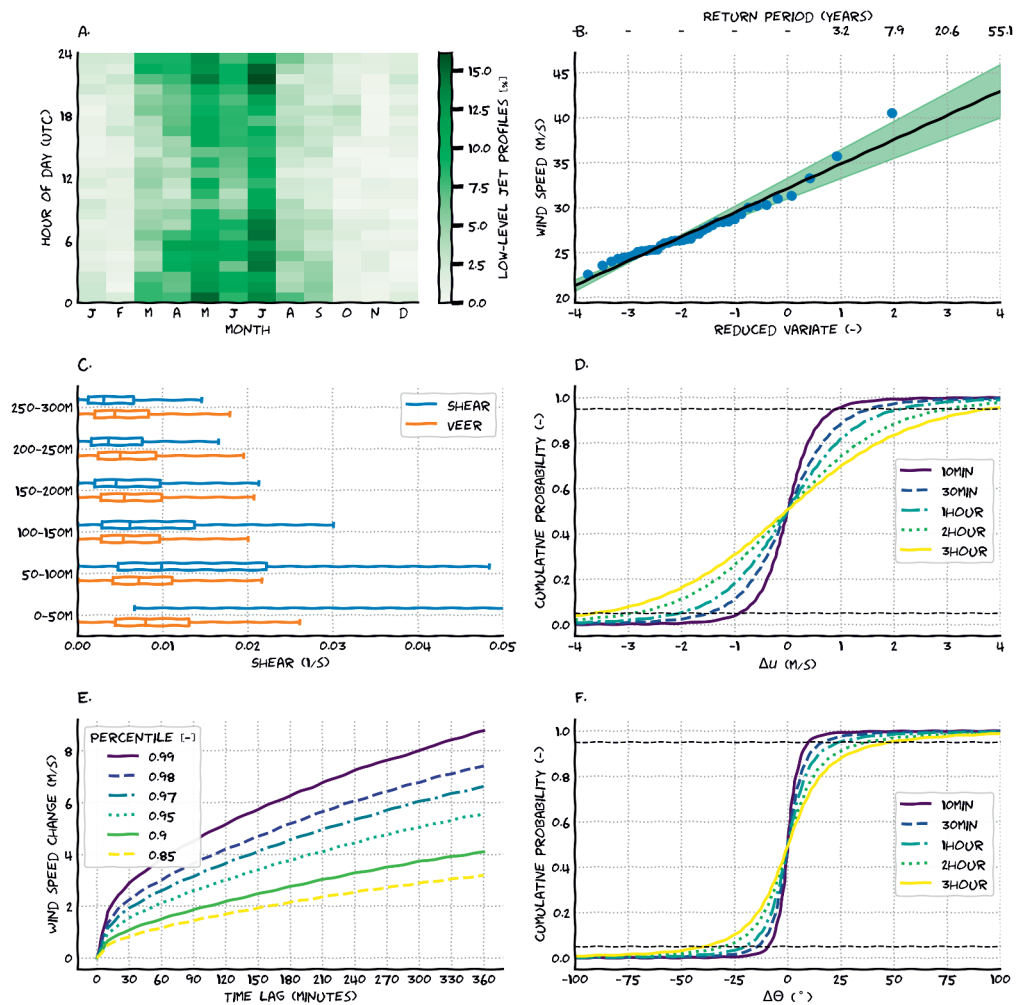
Low-level jets for which the fall-off exceeds  $2 \text{ m s}^{-1}$  are observed 4.6 % of the time, but this figure changes to 6.9 or 3.1 % when the threshold is increased or decreased by  $0.5 \text{ m s}^{-1}$ , respectively. It is noteworthy that most LLJs are observed at or around 100 m altitude, i.e. exactly in the rotor plane (not shown). This is lower than onshore, where they are typically observed around 250 m (Baas *et al.*, 2009), and underpins their importance for offshore energy purposes. Kettle (2014) already found similar characteristics for the FINO-1 tower in the German Bight, but because of the mast’s limited height, he was unable to identify the full spectrum of jets.

LLJs occur mostly in spring and summer (figure 2.5A). Their temporal distribution closely follows that of stability (compare figure 2.4D and 2.5A). In an absolute sense, there are approximately as many jets from the NNE as from the SW, but since the southwesterly winds are more abundant in the total dataset, the relative occurrence of LLJs is highest for north-northeasterly winds (figure 2.6B). Jets mostly occur in episodes of relatively high pressure and moderate winds, but some were observed during very low pressure (figure 2.6C). LLJs are usually associated with low values of the Charnock parameter, pointing to low wind-wave coupling or ‘decoupled’ conditions (figure 2.6F). In the light of these characteristics, it seems that multiple mechanisms can lead to wind profiles that are classified as LLJs. Most characteristics support the hypothesis that LLJs form under relatively calm, high pressure situations, when warm continental (NNE) air is advected over a cold ocean (spring/summer) leading to stable stratification ( $Ri_B > 0$ ) and decoupling of the surface layer (low Charnock parameter). Other jets, associated with somewhat stronger southwesterly winds and low-pressure situations may be related to the passage of fronts or squall lines. This is potentially much more damaging, but needs to be further investigated. There are few observations that support the baroclinity mechanism, since the jet direction would be parallel to the horizontal temperature gradient, which would be a meridional wind in the case of the Southern North Sea.

The hub-height 50-year wind extreme is estimated at  $42.7 \pm 2.4 \text{ m s}^{-1}$  (figure 2.5B). This is evidently lower than that specified in the IEC guidelines for a class I turbine (reference mean wind speed at hub height is  $50 \text{ m s}^{-1}$ ). The wind extreme decreases near the surface, but remains more or less constant with height above the tower. Assuming that for strong wind conditions the atmosphere can be described by a neutral logarithmic wind profile, and that there is a strong interaction between wind and waves (see section 2.4.1), we find that the following expression

$$u(z) = \frac{u_*}{\kappa} \ln \left( \frac{gz}{\alpha u_*^2} \right) \quad (2.7)$$

with  $u_*$  the friction velocity,  $\kappa = 0.4$  the von Karman constant,  $g = 9.81 \text{ m s}^{-2}$  the gravitational constant and  $\alpha$  the Charnock parameter (for which we take a high value of 0.035, consistent



**Figure 2.5:** A: Temporal distribution of low-level jet profiles. B: Gumbel plot for extreme value analysis. Return times are indicated on the upper x-axis. C: Box plots for wind shear and veer distribution over six layers of equal depth. D: Cumulative probability of wind ramps on various time scales. The 5- and 95-percentiles are indicated with dotted horizontal lines. E: Wind ramp magnitude versus time lag for different cumulative probabilities. F: Cumulative probability for wind direction ramps.

with high wind speed observations) provides a reasonable fit to our extreme wind profile with  $u_* = 1.7 \text{ m s}^{-1}$ .

The 95-percentile of wind speed increases with observation height and is around  $19 \text{ m s}^{-1}$  at hub height (115 m). Recall that we focus on this 95-percentile because we want to derive general characteristics of the episodes with the highest winds. Using higher quantiles, some generality might be lost because of insufficient data for a representative sample. Extreme wind speeds usually have a southeasterly direction and are predominantly observed under low-pressure situations in winter months. These characteristics support the hypothesis that the strongest winds are associated with low-pressure systems. The Richardson number is always near-neutral in these conditions because the lower atmosphere is typically well-mixed due to the strong wind-induced turbulence. Also, there is a strong interaction between wind and waves, characterized by a large value of the Charnock parameter.

The strongest shear is confined to the lowest layer (figure 2.5C). Above 150 m, the lateral shear component may become larger than the longitudinal shear. 70 % of the wind profiles veer with height (not shown) and on average, these profiles have a larger lateral component than the other 30 % of the data. This may be compensated by the fact that rotor blades are moving along with the lateral wind component in veering conditions, whereas they rotate opposite to the lateral wind in backing conditions.

Over the rotor plane, the 95-percentile of the longitudinal shear amounts to  $0.032 \text{ s}^{-1}$ . Over a 150-m layer, this constitutes a change of  $4.8 \text{ m s}^{-1}$ . Extreme shear is often observed under southwesterly and, to a lesser extent, northeasterly winds (figure 2.6). It peaks in April/May and December/January and is usually associated with a slightly stable stratification. The probability of extreme shear is higher under stronger winds and lower pressure. The Charnock parameter is typically quite high ( $> 0.03$ ), although some extreme shear was also observed for very low Charnock parameters. From this climatology we infer that extreme shear is mostly associated with strong SW winds and stable stratification. To understand this, we first note that the absolute value of the Richardson number is by definition small for strong shear. However, as extreme shear is observed more frequently for positive (stable) Richardson numbers than for negative values, it appears that the turbulence-suppressing buoyancy effects allow the stratification to be maintained even under high shear conditions. Conversely, for unstable conditions buoyancy acts to enhance the turbulent mixing and hence the wind speed gradient is smeared out. Additionally, the conditions with NE winds leading to LLJs seems to be favourable for extreme shear as well, although we cannot (yet) conclude whether this shear is indeed associated with low-level jets. This will be further explored in section 2.4.3.

The 95-percentile of veer over the rotor plane amounts to  $0.017 \text{ s}^{-1}$ , adding up to  $2.6 \text{ m s}^{-1}$  difference between the top and bottom of the rotor. Inspection of these extreme veer profiles demonstrated that, although veering with height is usually much stronger than backing with height, both cases contribute to the extreme conditions. The climatology of extreme wind veer is quite similar to extreme shear, although the annual cycle is even less distinctive (figure 2.6A). The fact that most veer is observed in stable conditions is in accordance with Van Ulden and Holtslag (1985).

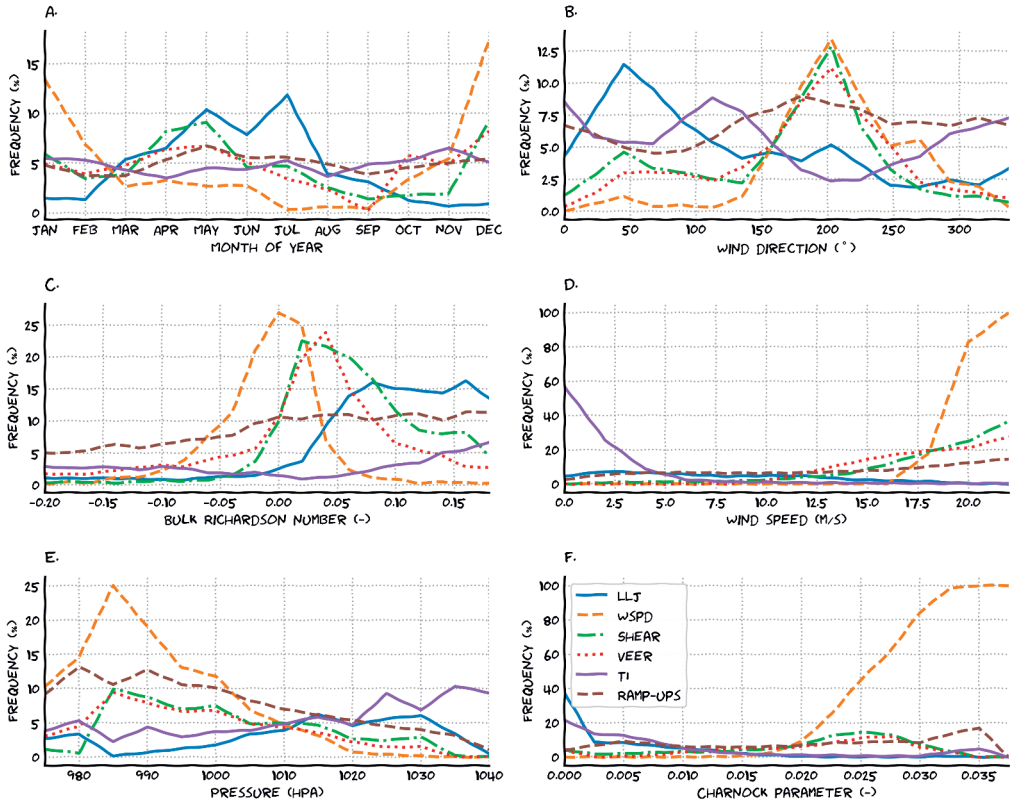
Figure 2.6 shows the temporal distribution of turbulence intensity, and its relation to external variables. Most notably, it appears that  $TI$  is highest for (very) low wind speeds, which is a direct result of the fact that  $TI$  is by definition inversely proportional to mean wind speed. Essentially, the other figures are just a reflection of this feature. The 95-percentile of  $TI$  amounts



to 0.13, and the extremes occur under northerly or southeasterly winds, and somewhat more in autumn and winter. Since wind speeds are low, the Richardson number is typically far from neutral. It is predominantly negative (which is consistent with the seasonality, i.e. unstable conditions in autumn and winter), although high values of  $TI$  are also observed under stable conditions, as long as the winds are weak. High  $TI$  is abundant for high pressure situations and the Charnock parameter is typically low. If, for a brief moment, we put aside our focus on extremes, we can compare the  $TI$  climatology with previous work. Hansen *et al.* (2012) evaluated turbulence intensity for a coastal site in Denmark and found that turbulence is highest for low wind speeds and unstable conditions. This is in line with the findings of Westerhellweg *et al.* (2010), who analysed the characteristics of  $TI$  as observed at the different FINO towers and with Kettle (2014), who focused on FINO-1. It is thus not surprising that we find the extremes for these low wind, unstable conditions. But although our findings make sense from a physical perspective, it is questionable whether these situations with extreme turbulence intensity are indeed the most penalizing conditions for wind turbines. Therefore, we also analysed the climatology of extreme standard deviations of the winds. The result is exactly opposite, i.e. the largest standard deviations are found under the strongest winds and the climatology now resembles that of extreme mean wind speeds. The 95-percentile of  $\sigma_u$  amounts to  $1.30 \text{ m s}^{-1}$  and high  $\sigma_u$  typically occurs in winter under strong southwesterly winds. With different proxies leading to different climatologies, we are left to conclude that based on the 10-minute means, we cannot identify typical conditions for extreme turbulence and wind gusts in a similar manner as the previously discussed AWEs.

The cumulative probability of wind speed ramp events on different timescales is given in figure 2.5D. A typical 10-minute ramp-up event (95-percentile) appears to be characterized by a wind speed change of  $0.9 \text{ m s}^{-1}$ . The ramp distribution appears to be very symmetric, such that typical ramps-downs (the 5-percentile) are of comparable magnitude with the ramp-ups. Organizing the percentile values against time intervals summarizes the expected wind ramps (figure 2.5E), which can be used as a reference for future wind climatologies and ramp studies. The ramp magnitude turns out to be well-described by a square function of the form  $y = a\sqrt{x}$ , with  $a = (1.5, 1.2, 1.1, 0.9, 0.7, \text{ and } 0.5)$  for percentiles (0.99, 0.98, 0.97, 0.95, 0.9, and 0.85). Although changes in the order of  $1\text{--}2 \text{ m s}^{-1}$  seem rather minimal, the actual power production is proportional to the wind speed cubed. For the LEANWIND reference wind turbine (Desmond *et al.*, 2016) with a rated power of 8 MW at a wind speed of  $12.5 \text{ m s}^{-1}$ , a typical 1-hour ramp at MMJ of  $-2 \text{ m s}^{-1}$  (from  $12.5\text{--}10.5 \text{ m s}^{-1}$ ) leads to a power reduction of 25 % of the rated power, which is a serious ramp event (Truwind, 2008; Gallego-Castillo *et al.*, 2015). However, if a similar ramp occurred at a wind speed well above the rated power, the power would not decrease at all. From this, it is apparent that the actual ramp impact is dependent on the wind speed at which it occurs. Therefore, figure 2.7 shows the typical ramps that can be expected, given a certain wind speed and time lag. Different rows give three different probabilities. Not surprisingly, strong ramp-downs are more probable for very high wind speeds and vice versa. The range between  $5\text{ and } 20 \text{ m s}^{-1}$  is most relevant for most contemporary wind turbines. Especially the 1-percentile plot for extreme ramp-downs is important, because for this ramp probability the wind speed change is very much dependent on wind speed, as can be seen from the sloping contour lines.

There is no clear seasonality for wind ramps (figure 2.6). Ramp-downs exhibit a slight preference for southwesterly winds, while ramp-ups are equally probable for any wind direction except for easterly winds, for which they occur somewhat less. The strongest ramps are typically observed



**Figure 2.6:** The relative occurrence of AWEs given A: month of year, B: wind direction, C: stability, D: wind speed, E: pressure, F: Charnock parameter. Except for LLJ, all anomalous events are defined by exceedance of the 95-percentile. wspd denotes wind speed.

under relatively strong winds and low pressure, although some are also observed at lower wind speeds. Ramps are slightly more frequent if  $Ri_B > 0$ . Unfortunately, this climatology does not reveal important clues as to the governing mechanisms for wind ramps.

For changes in wind direction (figure 2.6F), the 95-percentile amounts to 10, 17, 26, and 54° for time shifts of 10, 30, 60, and 180 min, respectively. However, this distribution is asymmetric, and the probability of direction ramps backing in time is a bit smaller; the 5-percentile amounts to 9, 15, 23, 36, and 48°, for the same times shifts, respectively. Large direction ramps are more probable for northerly and southeasterly winds (not shown), but more importantly, the largest direction changes occur under very weak wind conditions and are therefore probably not relevant. Figure 2.8 shows the expected direction changes with specified probability, given a wind speed and time lag. Under stronger wind speeds, wind veering is more probable than wind backing in time. A plausible mechanism associated with these ramps is the passage of cold fronts (Wallace and Hobbs, 2006). The fact that strong direction ramps are more probable for northerly and southeasterly winds, in combination with the notion that the strongest shifts occur under weak winds, suggests that these shifts are related to local circulations in the absence of strong synoptic forces, such as the sea breeze (Zack, 2007).

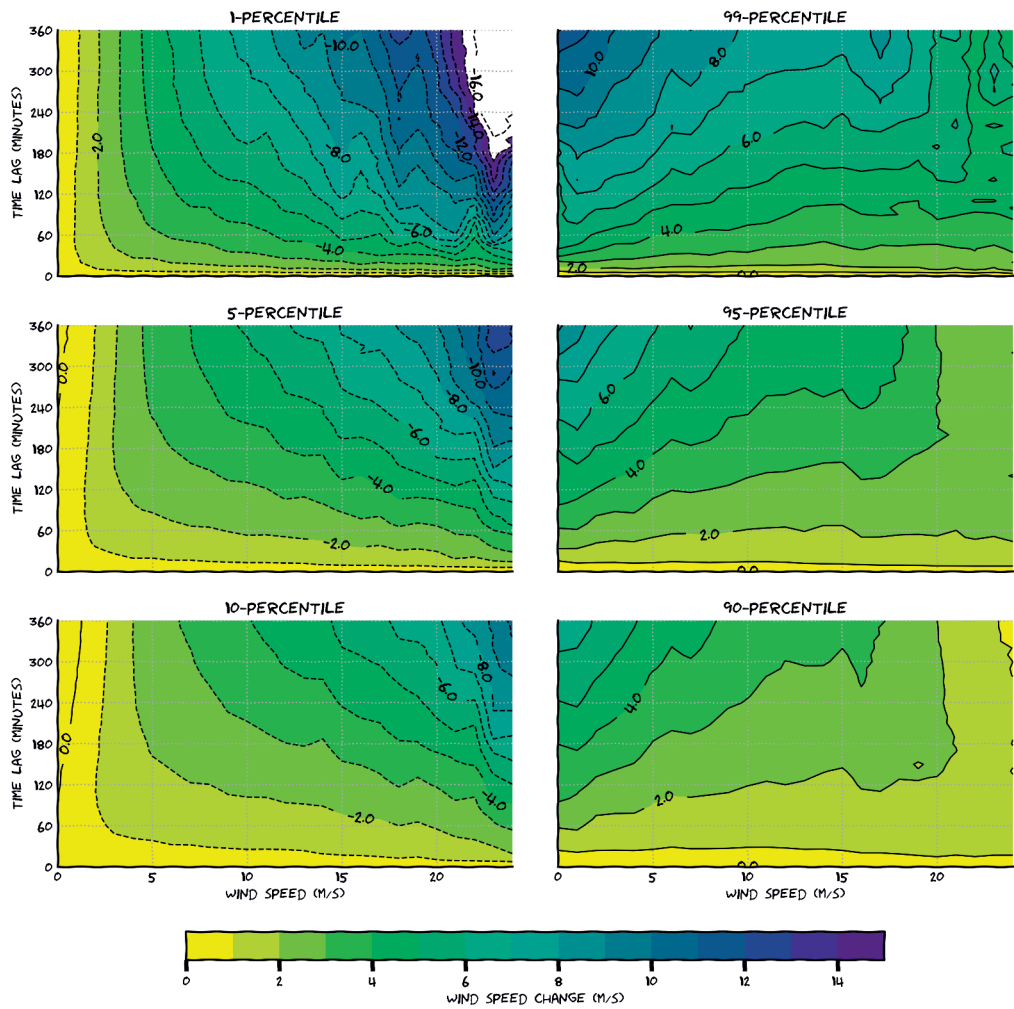


Figure 2.7: Expected wind speed ramp event with given probability (panel title), given a wind speed and time lag. Dashed contours indicate negative values; the colour bar shows ramp magnitude only.


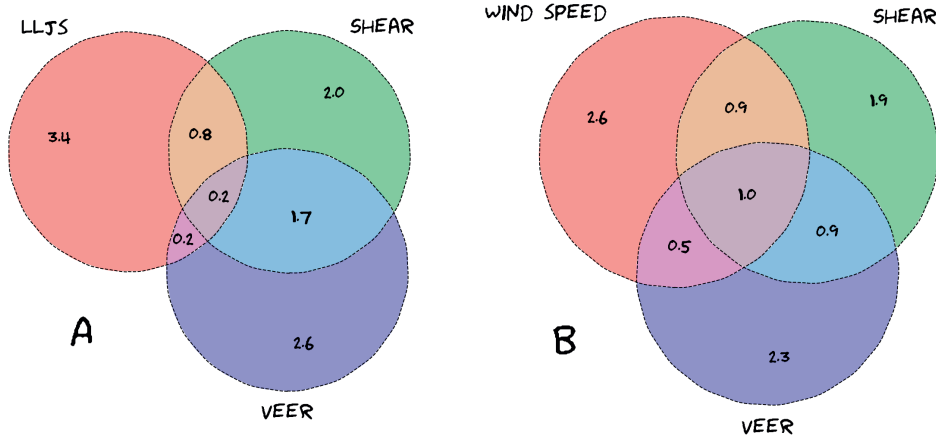


Figure 2.8: Expected wind direction ramp event with given probability (panel title), given a wind speed and time lag. Dashed contours indicate negative values; the colour bar shows ramp magnitude only.



**Figure 2.9:** Venn diagram showing relations between anomalous wind events. A: Low-level jets, extreme shear profiles and extreme veer profiles. B: Extreme wind speeds, extreme shear and extreme veer profiles. Numbers are in percent of total data.

### 2.4.3 (In)dependence of wind events

We have seen that the conditions for which various anomalous wind events occur are sometimes similar, and it is possible that they overlap in time, e.g. strong shear might be associated with low-level jets. Therefore, we analyse in this section the concurrence between the various wind events.

We study the overlap of events using Venn diagrams representing the 95-percentiles of each time series (figure 2.9). Diagram A illustrates that, although part of the LLJs is associated with extreme shear, still most extremes are independent. Likewise, there is some relation between extreme shear and extreme veer, but they do not always coincide. Even though LLJs are typically characterized by changes in wind direction (Baas *et al.*, 2009), the dependence between LLJs and wind veering is only small. In our analysis the veer is expressed as normal (lateral) wind component that scales with the magnitude of the wind. Therefore, most veer is related to strong winds, whereas LLJs typically occur under moderate winds. The amount of shear averaged over all low-level jet profiles is  $0.025 \text{ s}^{-1}$  over the layer 50–150 m, which is indeed lower than the extreme shear profiles.

Diagram B illustrates that extreme shear and veer regularly coincide with extreme wind speeds. Stability may explain why approximately half of the extreme wind profiles does not overlap with shear and veer: for strong winds and stable conditions, the buoyancy force acts to maintain the stratification, whereas under unstable conditions the shear-induced turbulence is only reinforced by buoyancy forces, leading to well-mixed boundary layers without much shear or veer in the mixed layer. The median of  $Ri_B$  is 0.017 for all extreme wind speeds, but 0.028 for extreme winds that coincide with extreme shear.

In a similar manner, we found that LLJs and extreme wind speeds never coincide, but wind ramps sometimes coincide with LLJs and sometimes with extreme wind speeds (not shown). Little overlap between  $\sigma_u$  and turbulence intensity appears, as noted before. While the former often coincides with extreme wind speeds, neither of these proxies overlap with extreme shear

profiles. Again, this can be explained by the fact that extreme shear is mostly observed under stable conditions, where buoyancy acts to suppress turbulent mixing. The fact that both turbulence proxies give the same result in this case, supports that under stably stratified conditions, even under relatively strong wind shear, turbulence is weak.

## 2.5 Discussion

We have found that low-level jets occur up to 12 % of the time at the North Sea, exactly in the rotor plane of contemporary wind turbines. Traditional wind profile parametrizations, such as the logarithmic wind profile and the power-law approach, are unsuitable to describe these wind profiles and therefore a reference climatology (figure 2.1) is incomplete and biased. Instead, one must separately account for LLJs, and these records should be removed from the dataset before deriving wind speed distribution parameters that can be used in conjunction with profile parametrizations.

Our findings suggest that the formation of LLJs is most frequent under stable stratification in spring and when the wind is directed offshore, similarly as found over the Baltic Sea (Dörenkämper *et al.*, 2015). It would be worthwhile to evaluate their representation in mesoscale models, to assess the spatial characteristics and physical mechanisms governing these jets and their relation to low-level jets that formed onshore.

Stability also plays a key role in distinguishing between extreme wind speed profiles and extreme shear profiles. The strongest wind speeds typically occur in winter, when the sea is relatively warm with respect to the overlying or advected air. This leads to unstable stratification, where buoyancy forces act to enhance turbulent mixing, ultimately resulting in well-mixed profiles. In contrast, one finds the highest wind shear under stable conditions, when buoyancy suppresses turbulent mixing.

Traditionally, shear is thought to be a turbulence-producing process; e.g. Moeng and Sullivan (1994) systematically studied the differences between buoyancy- and shear-driven boundary layers for unstable conditions, and found strong wind streaks in the shear-driven boundary layer, together with a high vertical momentum flux ( $\overline{w'w'}$ ). Our results suggest that offshore, stratification is the dominant mechanism for the manifestation of shear, rather than vice versa. We therefore suggest that more systematic modelling studies are conducted to further understand this subtle balance.

We found that different proxies for turbulence intensity lead to different climatologies, illustrating that more detailed, higher frequency data is essential for an accurate description and classification of turbulence. Our analysis was limited by the instrumentation of the mast and the 10-minute intervals at which the data were available. Observations closer to the surface and flux data from sonic anemometers would greatly enhance the possibilities to relate this dataset to other datasets as well as to evaluate it with classical boundary layer theory. For future studies it is imperative that measurements are taken closer to the surface, and include typical boundary-layer parameters such as shear stress and higher-order moments.

We have assessed the sensitivity of the climatology to the thresholds used. For low-level jets, the sensitivity was reported to be within 2–3 %. If, for the other AWEs, the 93- or 97-percentiles are used rather than the 95-percentile, the absolute quantile values may show non-negligible differ-

ences, but the patterns that were discovered in the external conditions remain nearly identical. This small sensitivity underlines the robustness of the current findings.

Analysing accumulated shear over different layers allowed us to include LLJs, which would not have been possible in traditional methods that would quantify shear as the ratio of wind speeds at two heights, or with a power law. By treating the occurrence of anomalous events in a boolean framework we were able to present a novel and consistent climatology of anomalous wind events.

Finally, it is worthwhile to consider the impact of anomalous wind events on turbine loads and power production. Low-level jets, lateral wind shear, and wind ramps are not considered in the contemporary load assessment guidelines. Extreme wind speed and turbulence are considered, but in a very artificial manner. The way in which we treated these anomalous wind events fits very well with the load assessment procedure: there are load cases for normal conditions and load cases for the extremes. It would be very informative to set up additional load cases for low-level jets, wind profiles with directional shear and load cases with transient mean wind profiles (wind ramps). Our observational study can be used as a basis for these load cases, but an approach in which meteorological models are coupled to the load assessment procedure is even better. A proof of concept on this methodology has already been published by Park *et al.* (2014).

## 2.6 Conclusions

We have analysed observations up to 300 m from met mast IJmuiden (MMIJ), situated in the North Sea 85 km off the Dutch coast. Six types of ‘anomalous wind events’ (AWEs) were defined to characterize certain extremes in the wind climate: low-level jets, extreme shear, both in wind speed and direction, extreme turbulence and gusts, and wind ramps (sudden transitions). This list covers the most important phenomena for wind industry, but it is certainly not exhaustive. Other AWEs can be defined and analysed following a similar methodology.

Our first objective was to objectively define the anomalous wind events, and to describe them in terms of typical values. We found that 4.6 % of the profiles should be classified as low-level jets, but with a distinct seasonal variability (up to 12 % in July) related to sea-surface temperature. The 50-year wind extreme at hub height was found to be  $42.7 \pm 2.4 \text{ m s}^{-1}$ . For the classification of wind ramps we introduced several new figures that provide a comprehensive overview of wind ramps on different time scales and for different wind regimes.

To assess the temporal variability of anomalous wind events as well as their relation to other weather conditions, we took the values that are exceeded exactly 5 % of the time. We found different weather regimes that are favourable for the formation of AWEs. This provided clues as to the underlying mechanisms, but in-depth case studies will be required to draw solid conclusions. Stability seems to play a very important role in the manifestation of several AWEs. Overlap between the various AWEs was assessed with the use of Venn diagrams. Although some extremes coincide, most AWEs are independent of each other.


This reference climatology will act as the point of departure for case studies with meteorological models to enhance mechanistic understanding of the presented AWEs.





# CHAPTER 3





## EVALUATION OF THE IFS, WRF, AND HARMONIE WEATHER MODELS WITH OBSERVATIONS OVER THE NORTH SEA

---

ALTHOUGH MODELLING IS A CENTRAL COMPONENT OF MODERN SCIENCE, SCIENTIFIC MODELS AT BEST ARE APPROXIMATIONS OF THE OBJECTS AND SYSTEMS THAT THEY REPRESENT - THEY ARE NOT EXACT REPLICAS. THUS, SCIENTISTS CONSTANTLY ARE WORKING TO IMPROVE AND REFINE MODELS.

- KARA ROGERS, ENCYCLOPAEDIA BRITANNICA

Much of our knowledge about the wind climate is obtained through weather models. They are not only used to produce forecasts, but also to investigate flow patterns, to reconstruct long-term climatologies, and to provide inflow fields for wind energy engineering models.

In this chapter, the observations from met mast IJmuiden are used again, but now to assess the performance of three commonly used weather models. All models appear to be very ‘case-sensitive’, with better performance for strong wind regimes and well-mixed wind and potential temperature profiles. Conditions characterized by moderate wind speeds combined with stable stratification lead to the largest misrepresentations. Thus, the conditions associated with anomalous events in chapter 2 are found worst represented by weather models.

A secondary objective of this chapter is to demonstrate a new case selection strategy, which is inspired by the idea that the large-scale circulation determines to a large extent the local manifestation of the flow. This strategy is shown to provide a good representation of the local climatology and may inspire future downscaling applications, such as discussed in chapter 7.

This chapter is published as:

Peter C. Kalverla, Gert-Jan Steeneveld, Reinder J. Ronda and Albert. A. M. Holtslag (2019). Evaluation of three mainstream numerical weather prediction models with observations from meteorological mast IJmuiden at the North Sea. *Wind Energy*, 22(1):34–48.

### 3.1 Introduction

Flow modelling is an integral activity in the development of wind energy systems. Applications range from resource assessment (Rodrigues *et al.*, 2015; Archer and Jacobson, 2005) to load calculations (Park *et al.*, 2014; O'Brien *et al.*, 2017; Schulz *et al.*, 2017) and from wake simulations (Sanderse *et al.*, 2011; Mehta *et al.*, 2014; Dörenkämper *et al.*, 2015) to power forecasting (Foley *et al.*, 2012; Jung and Broadwater, 2014). A comprehensive review of flow modelling in the wind energy sector was recently published by Sanz Rodrigo *et al.* (2017b). In that study, the authors distinguish between meso- and microscale models.

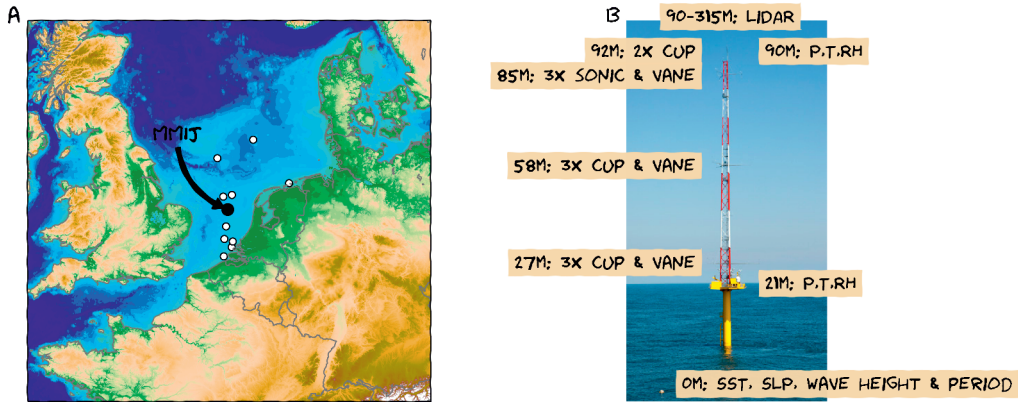
Mesoscale models typically use a grid spacing of a few kilometres and support domains that are large enough to incorporate large-scale weather patterns such as low-pressure systems, fine enough to resolve regional circulation patterns such as the sea breeze, but too coarse to resolve turbulence and its interaction with wind turbines (Talbot *et al.*, 2012). These models are typically used to refine the results of global modelling products, which are run on coarser grids, for a limited area. Important applications of mesoscale models in the wind energy sector are in forecasting and resource assessment (Rodrigues *et al.*, 2015; Foley *et al.*, 2012). Currently (01-2018), several institutes are working on the production of a New European Wind Atlas (Petersen *et al.*, 2013) based on mesoscale model results. At the same time, these models are further developed, e.g. in the context of the Wind Forecast Improvement Project 2 (WFIP2; McCaa *et al.*, 2017).

Microscale models use finer grids in order to resolve turbulence and interactions with wind turbines, but their domains are too small to incorporate regional or large-scale circulation patterns. These models typically use idealized, stationary forcing, which poses limitations when dealing with transient phenomena such as fronts (Muñoz-Esparza *et al.*, 2014). Coupling between meso- and microscale models is an active research field (Talbot *et al.*, 2012; Sanz Rodrigo *et al.*, 2017b; Muñoz-Esparza *et al.*, 2014; Mirocha *et al.*, 2014; Sanz Rodrigo *et al.*, 2017a) that is expected to substantially reduce uncertainties originating from the use of idealized inflow fields.

For a successful downscaling procedure and other applications, it is first and foremost important that all models are thoroughly validated. Since model performance varies with location (orography, land use), atmospheric conditions, and quantity of interest, and the models themselves are continuously improved, evaluation is an ongoing activity. For example, a downscaling product for resource assessment from 2012 was based on a 45 km grid (Petersen and Troen, 2012), whereas the New European Wind Atlas (Petersen *et al.*, 2013) will be based on a 3 km grid. The growing interest in offshore wind farm development also calls for model evaluation that focuses on offshore conditions, which have received considerably less attention than onshore conditions in the past due to a lack of observations.

In this study, we use a unique dataset that was obtained over a four-year period between 2012-2015 to evaluate the performance of three mainstream numerical weather prediction models. “Met mast IJmuiden” (MMIJ) was installed in the North Sea between the Netherlands and the UK (figure 3.1). The exceptional distance to shore (85 km), length (4 yr) and observation height (90 m mast, remote sensing up to 315 m) make the dataset well-suited for an evaluation tailored to wind energy purposes at the Southern North Sea. The models included in this study are:

1. The Integrated Forecasting System (IFS), the global operational weather model of the European Centre for Medium-range Weather Forecasts (ECMWF)



**Figure 3.1:** A. Map showing the location of the met mast and available automated weather stations (white dots) in the North Sea. B. A photograph of the mast with annotations of the instrumentation. P, T and RH are pressure, temperature and relative humidity, respectively. SLP and SST are sea-level pressure and sea-surface temperature.

2. HARMONIE-AROME, a short-range regional weather forecasting model developed and used operationally by a consortium of European weather institutes, and
3. The Weather Research and Forecasting (WRF) model, a community model maintained by the US National Center for Atmospheric Research (NCAR).

The results presented in this paper are inspired by the following research question: “What is the typical model performance concerning (the vertical structure of) wind, temperature, and humidity and to what extent does this performance depend on the model and on atmospheric conditions?” A wide variety of cases is analysed, each characterized by a unique combination of wind speed, direction and large-scale atmospheric stability. This characterization provides an interesting perspective on the results.

More details on the models and observations are given in sections 3.2.1 and 3.2.2. Subsequently, we explain the case selection procedure and the statistics used (sections 3.2.3 and 3.2.4). The results are presented and discussed in section 3.3. The paper concludes with a summary and conclusions section 3.4.

## 3.2 Methodology

### 3.2.1 Description of the models

We use three mainstream weather models (IFS, HARMONIE and WRF) that are all widely used for wind energy purposes, but have a different focus.

The IFS is a global model that provides forecasts of up to 15 d ahead of time on a grid of approximately 16 km horizontal spacing (at the time of the MMIJ observations). It is a hydrostatic model, which means that vertical motions are not prognostically resolved. Instead, a sub-grid parametrization is used to represent the effect of the unresolved convection. This approach is acceptable for relatively coarse grids, but becomes invalid for finer grids where convection

(strong up- and downdraughts such as in thunderstorms), local circulations, and turbulence start to be resolved. The latest operational analysis, which constitutes the best known state of the atmosphere at that time, is used as initial field. During the MMIJ observation period, the model received some updates, most notably a change from 91 to 137 vertical levels in June of 2013. Also, a change in the formulation of turbulent mixing in stable conditions was made in November of 2013 (Riddaway, 2014). Appendix 3.B shows which cases are before and after these changes. Further documentation of the IFS can be found on the ECMWF website<sup>1</sup>.

HARMONIE is a limited-area model that produces short-range forecasts of up to 48 h ahead of time for the region of Northwest Europe at a horizontal grid spacing of 2.5 km. It uses 65 vertical levels with the model top at 10 hPa (for comparison: the IFS model top is at 0.01 hPa). Initial and lateral boundary conditions are taken from the IFS; however, additional regional observations from satellites, stations, buoys, ships and aircrafts are assimilated into the initial condition before the forecast is started. In contrast to the IFS, HARMONIE is a non-hydrostatic model and it is assumed that deep convection reaching far above the boundary layer is completely resolved. However, the model still includes a sub-grid parametrization for shallow convection at boundary-layer scales. An extensive description of the HARMONIE model can be found in Bengtsson *et al.* (2017). For the studied period, the model version was 36.

The Weather Research and Forecasting (WRF) model, specifically the Advanced Research WRF (ARW) version 3.9 (Skamarock and Klemp, 2008), is a very flexible model which, in addition to forecasting and downscaling, is well-suited for case studies, sensitivity analysis, and for developing and testing new physical parametrizations. WRF is also popular among researchers and industrial users because of its flexibility and the fact that it is completely open source. WRF is a non-hydrostatic model which we force with initial and lateral boundary conditions from the IFS. However, where HARMONIE used the IFS forecasts as lateral boundaries (and assimilates additional observations), we use the subsequent operational analyses instead (and did not use data assimilation). For our reference simulations, we used 3 km horizontal grid spacing and 91 vertical levels up to 50 hPa (corresponding to the lowest levels in the 137 level formulation of IFS). In WRF, many sub-grid parametrizations are available. We used a scale-adaptive convection parametrization (Grell and Freitas, 2014) which gradually deactivates as grid spacing is refined. This scheme was found to perform well for resolutions up to 3 km (Fowler *et al.*, 2016). Further details of the WRF configuration can be found in appendix 3.A.

It is clear that the model set-up for each model is different. The WRF simulations were performed in-house, with a set-up that is very suitable for case studies. The IFS and HARMONIE data were taken from the operational archives of ECMWF and KNMI (Dutch national weather service), respectively, and as such represent the operational model set-up at the time of the MMIJ observations. Consequently, we evaluate the performance of each model in its typical set-up, as if it were to be used straight away. However, to bridge the gap between the horizontal grid spacing, we performed an additional set of WRF simulations at a resolution (and vertical grid spacing) comparable to the IFS. These simulations will be referred to as 'WRF-coarse'. Additional sensitivity runs with WRF-coarse were performed concerning the strength of the turbulent mixing in stable conditions, the impact of prognostics sea-surface temperature, and the impact of momentum entrainment at the boundary-layer top.

<sup>1</sup><https://www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model/ifs-documentation>

### 3.2.2 Observations and model data

The IJmuiden dataset is described in detail by Werkhoven and Verhoef (2012) and was previously analysed by Kalverla *et al.* (2017) and Holtslag *et al.* (2015). It is a four-year dataset (2012–2015) with wind observations from vanes and cup anemometers at 27, 58, and 92 m, from sonic anemometers at 85 m, and from a Zephir lidar between 90 and 315 m at 25 m intervals. All instruments have an uncertainty in the order of 1–2 % and they correlate very well ( $R^2 = 0.99$  Maureira Poveda and Wouters, 2015). Additionally, pressure, temperature and humidity were observed at 21 m (platform) and 90 m (tower top). All data is freely available from <http://www.meteomastijmuiden.nl/data/>. For more details we refer to Kalverla *et al.* (2017). In a preliminary study we found that the model performance in terms of wind speed and direction at this location is very similar to routine wind observations at other locations at the North Sea (also indicated in Fig 1). However, the observation height in combination with the distance to shore and the high quality of the MMIJ data is quite unique and therefore we focus on MMIJ in the presentation of the results.

The observations are available as 10-minute averages, whereas the model data from IFS and HARMONIE were available as hourly instantaneous model fields. WRF data for the MMIJ location was available at each time step, but for consistency in the comparison we used hourly instantaneous fields as well. These data were aligned with the corresponding observational data (i.e. the 10-minute average closest to the full hour). This approach was validated using the WRF dataset. By using the high-frequency output and trying various combinations of averaging and interpolation, it was found that this approach leads to the smallest ‘artificial’ model error statistics. Note that the instantaneous model values can be interpreted as spatial analogues of the corresponding time-averaged observations, although the exact analogy is dependent on the mean wind speed (e.g. at a mean wind speed of 10 m/s, a grid cell of 3 km would represent the average wind over an interval of 300 s).

Model data from the nearest neighbouring grid cells were linearly interpolated between the original model levels to match the MMIJ observation levels. All models were initialized at 12 UTC the day before the selected case date and forecast time series from 12–36 h ahead were used for validation. This is a relevant time scale for power forecasts (e.g. for power trading) and moreover, for this relatively short forecast horizon the effect of non-linearities in the model equations (the butterfly effect) is still small. As such, varying model performance will mainly be due to the (in)ability of the models to accurately capture the various atmospheric conditions.

### 3.2.3 Case selection: the $uvs \cdot t_2$ method

The WRF simulations are run in house on a rather large domain, which limits the number of cases. Since we expect that the model performance will be different for different weather situations, we use a clustering algorithm to objectively select a wide variety of weather types. From previous experience (Kalverla *et al.*, 2017), we know that wind speed, direction and atmospheric stability are useful parameters to distinguish between various wind regimes. However, these parameters can vary substantially in time or space within one case. In an intermittent turbulence regime, for example, there is no single characteristic Obukhov length to characterize the stability of the flow. This is because the Obukhov length is a very sensitive, *internal* parameter, based on local fluxes or gradients that are the result of (rather than the precondition for) the occurrence of boundary-layer processes.

To avoid this issue, the case selection can be based on larger-scale, *external* variables instead. The idea is that the observed conditions (e.g. intermittent turbulence) are the manifestation of the external forcing: an airmass driven over a water surface by an ambient pressure field. It follows that for similar characteristics of the airmass, water surface and pressure field, we expect similar flow regimes. This rationale is common in meteorological (downscaling) studies (Philipp *et al.*, 2010). However, existing classifications are not optimized for the current application. First of all, the correlation with local observations is not always clear. For example, the Großwetterlagen classification (Hess and Brezowsky, 1969) is not suitable because there can be many wind directions over the North Sea within one weather type. Moreover, existing classifications do not usually account for stability, while we know it can make a large difference for the manifestation of the wind field. Therefore, we introduce a new clustering algorithm that better suits our needs.

We find that a suitable balance between large-scale weather-type and the local flow is achieved if wind speed and direction are characterized in terms of the geostrophic wind components  $u$  and  $v$ . Following Jones *et al.* (2013), we use the south-north pressure difference over  $10^\circ$  centred over the area of interest to calculate the west-east ( $u$ ) wind component and vice-versa for  $v$ . To account for large-scale density stratification  $s$ , we use the difference between the potential temperature at the 850 hPa level and the sea-surface temperature (SST). Here, the 850 hPa potential temperature comprises information about the air mass while the SST encompasses the characteristics of the water surface. The difference between the two was found to be a more robust clustering parameter than each of them separately (based on the criteria presented in section 3.3.1).


Daily averages of large-scale  $u$ ,  $v$  and  $s$  are derived from ERA-interim (Dee *et al.*, 2011) mean sea-level pressure, SST and 850 hPa potential temperature fields in the 2012–2015 MMJ observation period. To improve the characterization of the airmass, the history of the flow is also relevant (Philipp *et al.*, 2010). Therefore, we additionally include the previous day to obtain a six-parameter dataset (hence the subscript in  $uvs \cdot t_2$ ). We use a hierarchical clustering algorithm that iteratively (1) splits the largest cluster (judged by average cluster radius) along a plane perpendicular to its principle component, (2) computes new cluster centres, and (3) reassigns all points to the closest cluster centre (Eggels *et al.*, 2018).

Once the algorithm has identified 30 clusters, we select the cases that are closest to the cluster centres. The number 30 is similar to other weather classification such as the Lamb Weather Types (Jones *et al.*, 2013) and Großwetterlagen (Hess and Brezowsky, 1969), and is aimed at selecting a wide range of characteristic weather conditions rather than exactly representing the climatology. Nevertheless, we perform several checks to verify the representativeness of our cases, which will be further discussed in section 3.3.1. For HARMONIE, model data for six of these cases were unavailable. Therefore, the second closest (to the cluster centre) dates from these clusters were used instead (see appendix B).

### 3.2.4 Statistical analysis: error diagrams

The difference between model predictions  $p_i$  and observations  $o_i$  is defined as  $e_i = p_i - o_i$  for each record  $i$ . A visualization in the form of a histogram or box plot (figure 3.2A) immediately provides insight into the nature of the errors. If the errors follow a normal distribution (which is a reasonable first approximation for our data), the complete set can be characterized by a mean error or bias,  $\mu$ , and a standard deviation of the error,  $\sigma$ . It can be shown (e.g. Murphy,





**Figure 3.2:** Various representations of a conceptual error distribution (left). In the box plot visualization, the middle vertical bar shows the median, the width of the box represents the interquartile-range (IQR), the whiskers, located 1.5 IQR outside Q1 and Q3, delimit the total range, and circles denote outliers. The histogram visualization closely resembles the Gaussian probability distribution function characterized by mean  $\mu$  and standard deviation  $\sigma$ . These can be used to visualize many cases in an *error diagram* (right).

1988) that these two fundamental properties are directly related to a third commonly used error statistic, the Root Mean Square Error (RMSE) as:

$$RMSE^2 = \mu^2 + \sigma^2 \quad (3.1)$$


We can use this property to concisely visualize model performance for many cases in a single diagram by plotting systematic errors  $\mu$  and random errors  $\sigma$  on the x- and y-axes, respectively (figure 3.2B). Interestingly, the distance to the origin in such diagrams represents the root mean square error, by virtue of Pythagoras' theorem. For convenience, we will refer to this type of figure as *error diagrams* in the remainder of this paper.

### 3.3 Results and discussion

In this section the results are presented and discussed, including cross-comparison with existing studies. The results will be presented in the following order. Section 3.3.1 deals with the case selection. The overall results of the model validation are presented in section 3.3.2. Special attention is paid to the role of resolution in section 3.3.3 and to the representation of vertical gradients in section 3.3.4. Finally, the results are related to large-scale characteristics in section 3.3.5.

#### 3.3.1 Case selection

We start with an inspection of the selected cases. A summary of each case is available in appendix 3.B. Figure 3.3 shows all 30 cases on a  $v$  versus  $u$  diagram. We can see that the strongest winds come from the southwest, which corresponds to the typical climatology of the North Sea. Observed static stability at MMIJ varies within each case, but the relative occurrence of stable stratification is consistently larger for cases with a stronger large-scale density stratification




**Figure 3.3:** Illustration of the clustering results. For each of the selected cases, the x- and y-axes show the  $u$  and  $v$  pressure gradients, respectively. The colour bar indicates the large-scale stability.

(the latter is indicated by the colour bar). The Weibull shape and scale parameters for both the full dataset (2.11 and  $11.80 \text{ m s}^{-1}$ ) and the subset of 30 cases (2.23 and  $11.87 \text{ m s}^{-1}$ ) turn out to be very similar. We have further checked that the cases adequately sample the annual cycle of sea-surface temperature: the cases are equally distributed over the year. Surface pressure charts (see Appendix B) indicate that we have sampled a broad range of synoptic circulation types and moreover, the surface pressure chart for each case does not deviate too much from the average surface pressure chart for its cluster. The correlation between the large-scale  $u$  and  $v$  and the IJmuiden observations is  $\sim 0.9$ , demonstrating that these clustering variables are indeed a very suitable ‘golden mean’ between large-scale circulation and local conditions.

To assess the uncertainty in the computed error metrics due to the case selection, we downloaded additional IFS forecasts for the full MMIJ period, and calculated error statistics for a varying number of cases. The most important results are plotted in figure 3.4. The full lines indicate the algorithm used in the  $uvs \cdot t_2$  selection (pca-partitioning), while the shading is based on a k-means algorithm (using the same input variables). The latter involves a random initialization step allowing us to estimate to some degree the uncertainty in the estimate. The figure demonstrates that this uncertainty is relatively small for most standard deviations, but can be relatively large with respect to the mean bias since the bias is often very small. For wind speed and potential temperature, the uncertainty has more or less converged at 30 cases, while for mixing ratio, the pca-partitioning deviates from the k-means uncertainty bands. With respect to the representation of vertical gradients, the uncertainty is larger and especially the standard deviation converges slowly.

Summarizing, we have sampled a variety of cases that allows us to investigate model performance under a wide range of forcing conditions. Care should be taken to generalize these results quantitatively, because the uncertainty due to the limited amount of cases cannot always be neglected. However, the quick convergence of the error statistics is promising and suggests that this method could be extended to downscaling climatologies as well. Perhaps our choice of large-scale variables could be combined with the subset selection strategy of Rife *et al.* (2013),



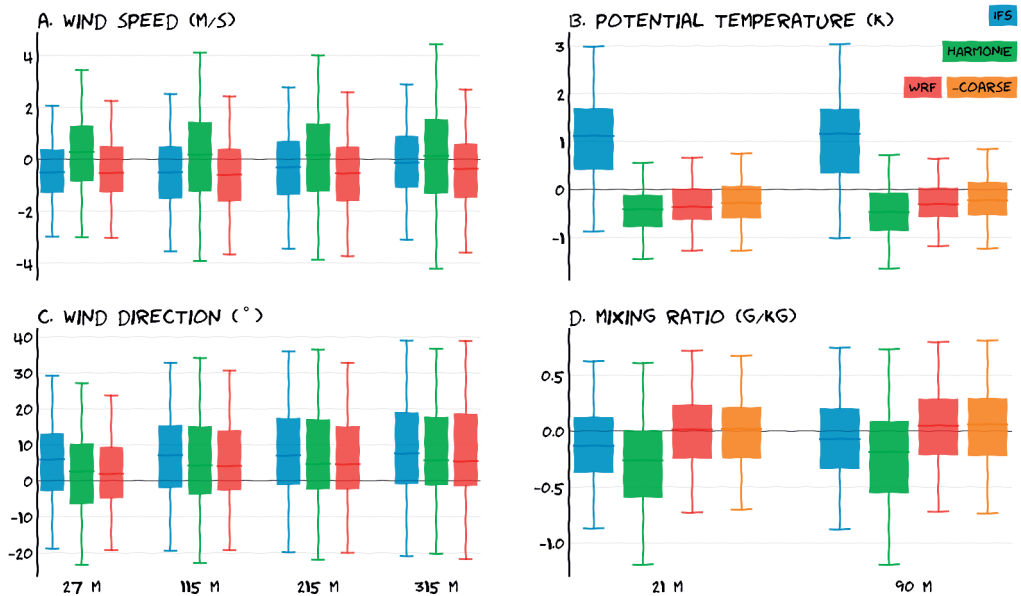
**Figure 3.4:** Uncertainty in selected error metrics as a function of the number of clusters. A: 115 m wind speed, B: 90 m potential temperature, C: 90 m mixing ratio, and D: Vertical difference in virtual potential temperature between 315 and 90 m.

who demonstrate that it is possible to obtain a fully representative subset of a climatology using 180 days.

### 3.3.2 Overall performance

Figure 3.5 shows the error distribution of wind speed, wind direction, potential temperature and mixing ratio at different altitudes for each model. Means and standard deviations of these distributions are listed in table 3.1. These results are weighed for the cluster size (see Appendix B). Note that the box plots are quite symmetric, which supports the assumption that the errors approximately follow a normal distribution.

Typical wind speed biases are within  $\pm 0.5 \text{ m s}^{-1}$  and  $\sigma$  below  $2 \text{ m s}^{-1}$ . On average, all models slightly underestimate the wind speed. The IFS is virtually unbiased at 315 m but underestimates winds closer to the surface. HARMONIE has the smallest wind speed bias, but shows a larger spread than the other models. These findings are in line with previous results. The ECMWF publishes standard validation metrics on their website, demonstrating that the mean bias and standard deviation of the errors in 10 m wind speed over Europe at forecast hour 60 are within  $0.5$  and  $2.0 \text{ m s}^{-1}$  (see Haiden *et al.*, 2014, figure 27). For the IFS ensemble system, the RMSE for short-range wind speed forecasts over land is reported to be around  $2 \text{ m s}^{-1}$  (Pinson and Hagedorn, 2012). For (earlier versions of) WRF, Hahmann *et al.* (2015) find small systematic



**Figure 3.5:** Box plots of the error distribution for wind speed, direction, temperature and humidity for all models at different heights. Lead time 12–36 h. See figure 3.2 for the definition of the box plot. Outliers not shown.

errors ( $0.09 \text{ m s}^{-1}$ ) in 100 m wind speed predictions at the FINO-1 mast, and Krogsæter and Reuder (2015) find Mean Absolute Errors of about  $2 \text{ m s}^{-1}$ . Similar errors are reported in Draxl *et al.* (2014), who find a wind speed bias ranging from  $-0.33$  to  $0.68 \text{ m s}^{-1}$  and RMSE ranging from  $1.95$  to  $2.69 \text{ m s}^{-1}$  and attribute the diverging results to differences in boundary-layer parametrizations and atmospheric stability. Carvalho *et al.* (2014) report wind speed biases larger than  $0.5 \text{ m s}^{-1}$  and  $\sigma$  roughly between  $1.5$  and  $2.5 \text{ m s}^{-1}$ . While Carvalho *et al.* attribute the positive bias they find to an excessively smooth sea surface, Hahmann *et al.* (2015) and Draxl *et al.* (2014) find their results to be sensitive to the boundary-layer parametrization used. Given this sensitivity, recent changes in the WRF model and the MYNN boundary-layer scheme might well explain the minor differences we see with our results. The smaller wind speed bias in the HARMONIE simulations might be the result of the additional data assimilation, although we were not able to run additional experiments to verify this hypothesis. With respect to the modifications to the IFS of November 2013 (Riddaway, 2014), it is worth mentioning that in our results, there is no marked improvement in the representation of surface winds. Observed differences are generally small, not consistent with height, and may be related to the frequency of stable conditions or other factors rather than the modifications to the IFS.

In all models, the wind direction is veered (turned clockwise) with respect to the observations by up to  $10^\circ$  on average (table 3.1). Performance is better near the surface than at 315 m. At 27 m, HARMONIE and WRF represent the wind direction better than the IFS, but at the other altitudes performance is similar. The misrepresentation of surface winds has previously been attributed to excessive mixing under stable conditions and the recent modifications to the IFS are expected to improve this deficiency (Riddaway, 2014). The ECMWF website reports a decrease in the systematic wind direction error near the surface over land since then ( $< 5^\circ$  after 60 hours). Also in our results, there seems to be an improvement, but we cannot conclude that this is due

**Table 3.1:** Mean error ( $\mu$ ) and centred root mean square error ( $\sigma$ ) for each model for different variables: wind speed  $v$  ( $\text{m s}^{-1}$ ) and direction  $d$  ( $^\circ$ ), (virtual) potential temperature  $\theta_{(v)}$  (K), and mixing ratio  $r$  ( $\text{g kg}^{-1}$ ); subscripts are altitudes.  $\Delta$  indicates difference between 315 and 27 m for wind and between 90 and 21 m for  $\theta_v$ . Bold numbers emphasize best scores.

Model	IFS		HARMONIE		WRF		WRF-coarse	
Variable	$\mu$	$\sigma$	$\mu$	$\sigma$	$\mu$	$\sigma$	$\mu$	$\sigma$
$v_{27}$	-0.48	<b>1.26</b>	<b>0.19</b>	1.87	-0.25	1.77	-0.24	1.95
$v_{115}$	-0.48	<b>1.59</b>	<b>0.02</b>	2.17	-0.49	1.93	-0.51	2.13
$v_{215}$	-0.33	<b>1.61</b>	<b>-0.02</b>	2.23	-0.47	2.02	-0.60	2.17
$v_{315}$	-0.12	<b>1.59</b>	<b>0.05</b>	2.27	-0.33	2.05	-0.50	2.19
$d_{27}$	5.97	<b>20.22</b>	<b>1.49</b>	25.77	2.07	23.71	3.12	26.66
$d_{115}$	8.60	<b>22.24</b>	<b>5.74</b>	25.48	6.52	24.42	6.83	26.41
$d_{215}$	9.74	<b>23.57</b>	<b>8.41</b>	27.68	8.67	26.71	9.01	27.14
$d_{315}$	1.02	<b>24.40</b>	<b>9.00</b>	26.77	10.27	28.19	9.02	29.93
$\theta_{21}$	1.00	0.94	-0.43	<b>0.58</b>	-0.19	0.75	<b>-0.10</b>	0.87
$\theta_{90}$	1.03	1.06	-0.54	0.95	-0.17	<b>0.83</b>	<b>-0.03</b>	1.03
$r_{21}$	-0.12	<b>0.40</b>	-0.27	0.51	<b>-0.01</b>	0.48	<b>-0.01</b>	0.45
$r_{90}$	-0.14	0.61	-0.24	<b>0.60</b>	-0.03	0.64	<b>-0.02</b>	0.62
$\theta_{v,21}$	0.99	0.94	-0.48	<b>0.57</b>	-0.20	0.75	<b>-0.10</b>	0.89
$\theta_{v,90}$	1.01	1.04	-0.58	0.97	-0.18	<b>0.80</b>	<b>-0.04</b>	0.99
$\Delta v$	0.37	<b>1.35</b>	<b>-0.06</b>	1.72	0.07	1.76	-0.13	1.82
$\Delta d$	<b>5.69</b>	14.12	7.87	20.87	7.30	17.40	6.30	<b>13.56</b>
$\Delta \theta_v$	0.03	0.58	-0.10	0.70	<b>0.02</b>	0.62	0.07	<b>0.57</b>

to the modifications, because the frequency and intensity of stable stratification is generally lower for the cases after November 2013.

The IFS shows a warm bias of roughly 1 K and also a large spread ( $\sigma \approx 1$  K). Oppositely, WRF and HARMONIE underestimate the temperature (at 21 m by 0.19 and 0.43 K, respectively) and show a smaller spread. Draxl *et al.* (2014) also evaluated temperature at 2 and 100 m from WRF simulations and found a positive bias at 2 m and either a positive or negative bias at 100 m depending on the boundary-layer parametrization. They report typical RMSE scores of about 1.3 K. Thus, our findings show improved performance. The most important difference is that Draxl *et al.* (2014) only used data from October 2009. The high-resolution verification chart of 2 m temperature at the ECMWF website shows that the bias is consistently within  $-0.5$  and  $0.5$  K, which is substantially smaller than our results. This is an average over a large area and based on synop observations on land, so it might not be directly comparable. Nevertheless, we expected that offshore performance would be relatively well-represented, because the surface is much more homogeneous. Mixing ratio is best represented by WRF. Both IFS and HARMONIE results show a dry bias, approximately  $0.1$  and  $0.2 \text{ g kg}^{-1}$ , respectively. Standard deviation amounts to about  $0.5 \text{ g kg}^{-1}$  for all models.

### 3.3.3 Role of model resolution

Since we include model results from both the relatively coarse grid of the IFS and the refined grids of HARMONIE and WRF, the question arises whether differences in results can be attributed to differences in model resolution, and whether a higher resolution contributes to an improved forecast (as is often assumed). Theoretically, we would expect that the higher-resolution models are able to incorporate local features that are not resolved by the IFS. Also, a coarser resolution would effectively smooth the time-series at a given location as it can only represent the grid-cell average. Indeed, table 3.1 shows that the IFS has the lowest standard deviation of the error for most variables, which is consistent with a smoother signal.

However, there are many other differences between the three models, so in order to isolate the effect of increased resolution we performed additional WRF simulations on a resolution of 16 km, comparable to the IFS. Some of these additional results are included in the right subplots of figure 3.5. It turns out that the change in resolution has only a small impact on these results. For most variables, the spread of the errors is larger for the coarser resolution runs, instead of smaller. Thus, the difference in model resolution does not seem to be responsible for differences between model results in this study and increased resolution does not necessarily lead to better results.

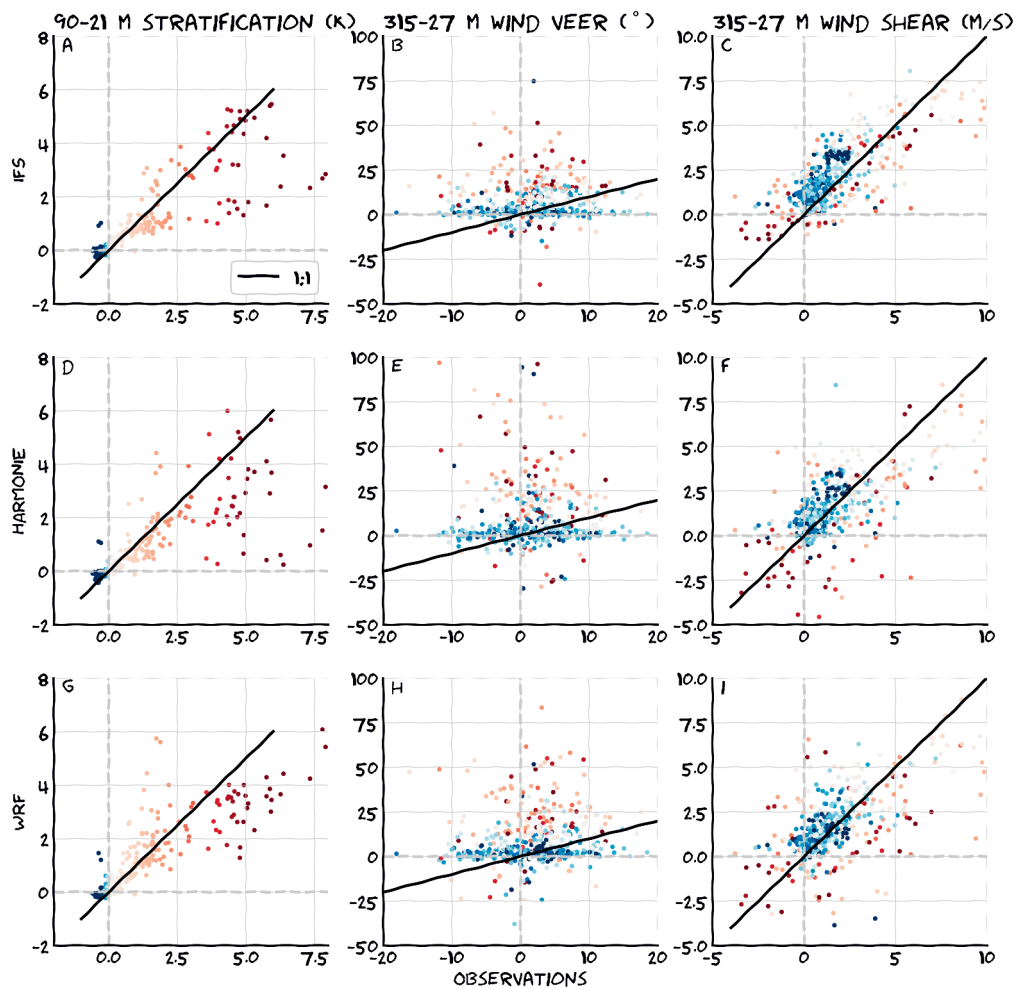
### 3.3.4 Boundary-layer structure

In the previous sections we evaluated the model performance at individual altitudes. For wind energy purposes it is especially important that the models adequately reproduce the vertical boundary-layer structure as well. Misrepresentation of wind shear or density stratification can have adverse effect on turbine loads (Park *et al.*, 2014) and power yield (Dörenkämper *et al.*, 2015). Besides, the gradients of wind and temperature are intimately linked through the process of turbulent mixing. Table 3.1 also lists statistics of vertical gradients, but in this case scatter plots are helpful to further understand the model performance (figure 3.6).

We first discuss the temperature stratification, which was calculated between the top of the mast (90 m) and the platform (21 m). While neutral conditions represent most of the data, most scatter occurs for stable conditions. In the IFS results, we observe two different types of behaviour: in one branch, the strong stable stratification is well represented, whereas in the other branch, the stratification is underestimated. The HARMONIE results also show a small branch in which the stratification is not captured by the model, while WRF shows a more continuous behaviour where intermediate to strong stratification is generally underestimated.

The difference in wind direction with height (calculated between 315 and 27 m) is not well represented by either of the models. Mostly, the models do not capture any wind veer at all, irrespective of the observed wind veer. If they do represent wind veer, it is often far outside the range of observed veer. As indicated by the colour scale, this is usually associated with stable conditions.

Modelled wind shear over the same layer correlates reasonably with the observations, with most scatter during stable conditions. The bottom-left quadrant in these figures indicate negative shear that is captured by the model. This is a characteristic feature of low-level jets, which are always associated with stable conditions.



**Figure 3.6:** Modelled versus observed: virtual potential temperature difference between 90 and 21 m (left), wind direction difference between 315 and 27 m (middle) and wind speed difference between 315 and 315 m (right) for the IFS (top), WRF (middle) and HARMONIE (bottom). Red and blue are used respectively for stable and unstable observed stratification (the colour scale is not linear).

The misrepresentation of stable boundary layers has received considerable attention in the past. For example, Brown *et al.* (2005) evaluated the representation of wind veer over sea in the IFS model and found a significant underestimation of wind turning across the boundary layer. The largest misrepresentation occurred for stable conditions or cases with strong baroclinity (horizontal temperature contrasts). Svensson and Holtslag (2009) related this issue to excessive mixing under stable conditions. This excessive mixing, which was also reported by Holtslag *et al.* (2013), inspired Sandu *et al.* (2013) to make the aforementioned modifications to the IFS in November 2013, which were expected to improve the representation of wind under stable conditions (Riddaway, 2014). In our results, we do not see a clear difference in the representation of vertical gradients for the cases before and after November 2013. However, the frequency and strength of stable stratification was considerably lower in the cases after this update, so we cannot conclude whether the representation of (very) stable conditions improved or not based on our results. With respect to the other modification, the addition of vertical levels in June 2013, there is one case with very stable stratification after this date (but before November 2013), which is reasonably well represented. Obviously, this is not enough to be conclusive, but it supports the hypothesis that the diverging model behaviour that we noticed in the scatter plot of temperature stratification, is related to modifications to the model.

### 3.3.5 A case against the average

So far we have focused on overall error distributions and obtained an impression of average model performance. However, if we focus on individual cases, we find substantial variability in model performance. To gain insight in the variability between individual cases, we will inspect error diagrams of 115 m ( $\sim$  hub height) wind and 90 m potential temperature and humidity (figure 3.7). Recall that the absolute distance to the origin represents the overall RMSE.

For hub-height wind speed (figure 3.7A), most cases cluster in an area with a mean bias between  $-1.5$  and  $1 \text{ m s}^{-1}$  and  $\sigma$  of about  $1.5 \text{ m s}^{-1}$ . Again, we find a tendency to underestimate wind speeds on average. Exceptionally poor performance is found for cases 5 (WRF); 6 (IFS); 2, 9, and 21 (HARMONIE); 24 (WRF and HARMONIE); and 18 (all models). These are almost exclusively spring and summer cases, in which the sea water is relatively cold and stable conditions occur frequently. For wind direction (figure 3.7C), two-thirds of the cases are clustered within  $15^\circ$  on both axes, without obvious differences between models. Cases 5 and 18, which are among the worst represented cases, are both characterized by weak-wind regimes, for which wind direction is usually less well defined. Nevertheless, the positive direction bias (modelled wind direction is veered w.r.t. the observations) is clearly present for almost all individual cases.

For temperature (figure 3.7B) we find a remarkable difference between the models. The cold bias of about  $-0.5 \text{ K}$  in WRF and HARMONIE is clearly evident for all cases, and very well-defined. In contrast, the IFS bias is very “case-sensitive”: it varies between  $-1$  and  $3 \text{ K}$ . Cases 5 and 18, which we already highlighted for poorly representing wind speed are also outliers in the temperature plot. These cases share a characteristic large-scale stable stratification ( $\theta_{850} - SST > 10 \text{ K}$ ). Errors in humidity are quite concentrated (figure 3.7D), but remarkably, all cases that have a larger variability also have a negative bias. This points to an underlying process that creates variability in moisture and the misrepresentation of which is associated with too dry boundary layers. Examples of such processes are frontal passages, precipitation, or dry-air entrainment at the boundary-layer top.



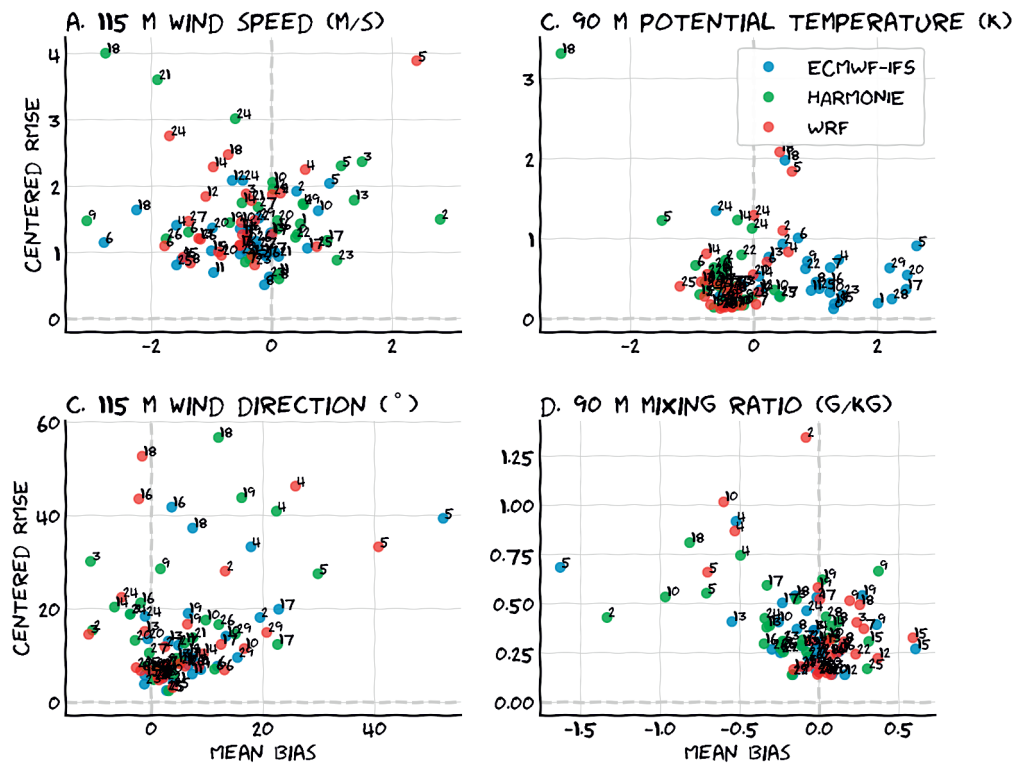


Figure 3.7: Error diagrams of 115 m wind speed and direction (A, C), and 90 m potential temperature and mixing ratio (B, D) at met mast IJmuiden for the IFS, HARMONIE and WRF models for lead times up to 36 h. See section 3.2.4 for an explanation of the error diagrams.





Figure 3.8: Like figure 3.7, but for 90 – 21 m virtual potential temperature (A), and 315 – 27 m wind speed (B) and wind direction (C) differences.



**Figure 3.9:** Error distributions of 90 – 21 m virtual potential temperature (A), and 315 – 27 m wind speed (B) and direction (C) differences for WRF, sorted by increasing large-scale stability ( $\theta_{850} - SST$ ).

A similar figure for the vertical difference in virtual potential temperature (figure 3.8A) indicates that all models typically perform very well (see also table 3.1), but a few cases are actually responsible for almost all of the overall error. Cases 2, 4, 5, 6, and 18 perform particularly poorly. Exactly these cases also underestimate the wind speed gradient (figure 3.8B) and have a large positive bias and spread in the wind direction gradient (i.e. too little backing or too much veering with height; figure 3.8C). These cases are characterized by (strong) stable stratification, or weak winds, or both.

The role of stability becomes especially clear from figure 3.9, where we sorted the box plots by the characteristic large-scale stability for that case. The representation of vertical boundary-layer structure is especially poor if the large-scale stability is strong, i.e. when relatively warm air mass is advected over a cold sea surface. This is indisputable for the vertical differences shown, but also apparent for 90 m potential temperature and 115 m wind direction (not shown). For 115 m wind speed and 90 m mixing ratio, the difference in model performance between stable and unstable conditions is less pronounced. Similar results are found for the other two models. It is worth noting that within each case, we find a both stable and unstable conditions locally (based on the temperature gradient near the surface). However, stable conditions are much more frequent in the cases characterized by a strong large-scale stratification.

### 3.4 Summary and conclusions

We have presented an evaluation of three mainstream meteorological models against unique observations up to 315 m at the North Sea. The integrated forecasting system (IFS) is the operational global weather model of ECMWF, HARMONIE is the operational limited-area model used by a number of European weather agencies, and WRF-ARW is a community research model. These models are used in wind energy systems for forecasting, resource assessment, and increasingly to provide input for microscale simulations of wind-turbine interactions. The observations from met mast IJmuiden provide a unique opportunity to validate these models at a far (85 km) offshore location up to 315 m.

Our results show that there is substantial variability in model performance. Both WRF and the IFS consistently underestimate the wind speed by about  $0.5 \text{ m s}^{-1}$ . The spread of the errors, expressed as standard deviation of the error distribution, is below  $2 \text{ m s}^{-1}$ . HARMONIE has a smaller wind speed bias than the other models, which may be due to the additional data-assimilation used in this model. The spread in the errors is smallest for the IFS, and we have shown that this is not due to the coarser grid.

Wind direction is found to be biased by up to  $10^\circ$ , i.e. the model winds are slightly veered (turned clockwise) with respect to the observations. One might argue that a wind direction mismatch of  $< 10^\circ$  is not particularly harmful. However, as Gaumond *et al.* (2014) point out, the wake losses are very sensitive to wind direction due to the typical offshore wind farm layouts with close turbine distances for certain directions. Furthermore, both WRF and HARMONIE show a consistent temperature bias of about  $-0.5 \text{ K}$ , while the IFS shows a positive bias of about  $1 \text{ K}$ .

Additionally, we find that the model results are very “case-sensitive”. All models struggle with the representation of stable conditions, leading to unrealistic boundary-layer structures. This is especially important because both power production (through enhanced wake losses) and turbine loads (through enhanced wind shear) are affected. Recently, Baas *et al.* (2016) and Kalverla *et al.* (2017) reported that over sea, strong wind shear can be maintained even under moderate to strong wind conditions if the density stratification is strong enough.

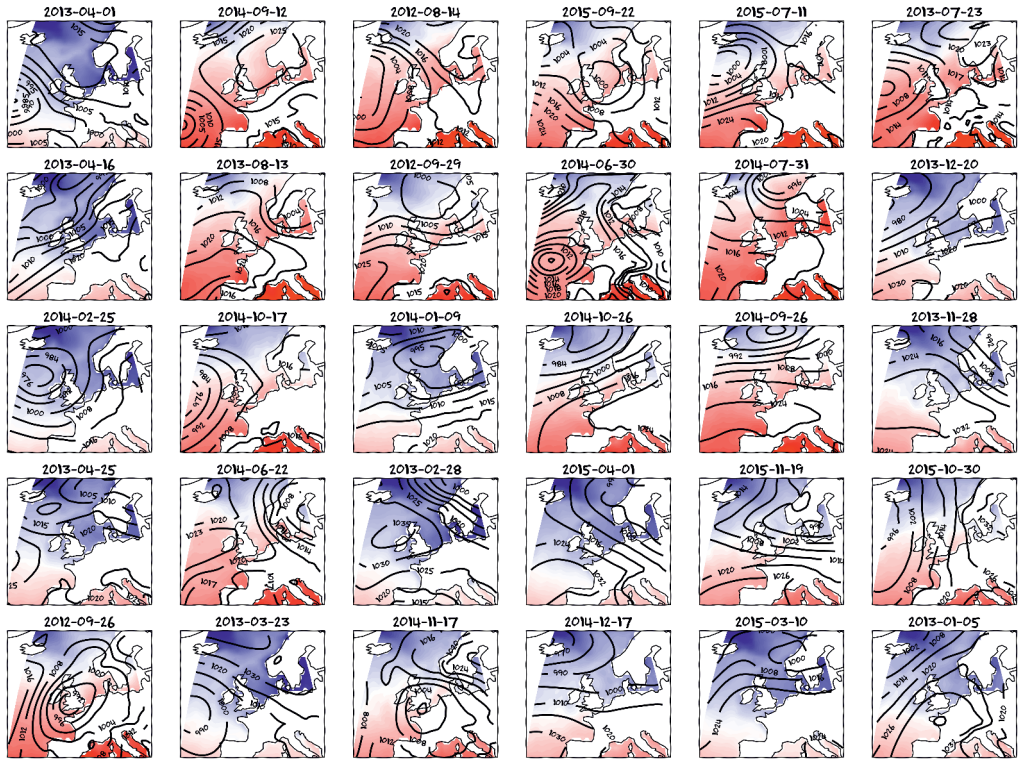
Our case selection strategy, based on the clustering of relevant large-scale variables, was not only able to select a wide range of circulation patterns, but also provided a reasonable representation of the underlying climatology. However, with the limited number of 30 cases included in this study, one should be careful to generalize the overall model scores presented in this study. This is especially true because the results in section 3.3.5 illustrate that the overall model performance as listed in table 3.1 cannot always be projected on individual cases.

Finally, our results suggest that the large-scale stability, expressed as the difference between the potential temperature at 850 hPa and the sea-surface temperature, can be used to make an a-priori estimate of the accuracy of model results. Since the large-scale stratification is easily determined, and also more robust and generic than the local stability, it could serve as guidance during the interpretation of model results.

### 3.A WRF set-up

We used WRFV3.9 (Skamarock and Klemp, 2008). The WRF settings were chosen such that they represent the latest trends and developments in mesoscale modelling. Most of the parametrization schemes have undergone active development in recent years. Settings are roughly similar to the NOAA high-resolution rapid-refresh model and the NCAR 3 km ensemble (Benjamin *et al.*, 2016). The following physics options were used:

- MYNN2.5 with mass-flux (stem) and updated length scale; tke-advection on
- MYNN surface layer
- RRTMG for both short- and long-wave radiation
- Thompson microphysics
- Grell-Freitas scale-adaptive convection



**Figure 3.B.1:** Synoptic charts of sea-level pressure (contours) and sea-surface temperature (shading) for each case. Numbering 0–29 is from left to right and top to bottom.

- Noah land surface

The single domain consisted of 600 by 600 grid cells of each  $3\text{ km} \times 3\text{ km}$  centred around the IJmuiden mast using a Lambert projection with true latitudes at  $45^\circ$  and  $55^\circ$ . In the vertical we used the ECMWF L137 formulation up to 50 hPa (91 levels in total). Additional sensitivity tests on a coarser resolution used the ECMWF L91 formulation for cases before 2013/06/25 to maximize the similarity to the IFS simulations. Simulations were run with a time step of 10 s.

### 3.B Selected cases

Table 3.B.1 lists details of the selected cases. Synoptic charts for these dates are shown in figure 3.B.1.

**Table 3.B.1:** All cases, sorted by month, with their average large-scale wind speed, direction and stability and number of days in the corresponding cluster. \*Before update IFS of June 2013; \*\*Before update IFS of November 2013. Bold numbers are second choice cases, because HARMONIE data was missing for the first choice (i.e. closest to cluster centers).


#	Date	Wind speed (hPa/10°)	Wind direction (°)	Stability (K)	N
29*	2013-01-05	11.5	289.5	16.5	12
<b>14</b>	2014-01-09	18.0	264.4	5.2	37
20*	2013-02-28	7.9	20.0	13.5	53
12	2014-02-25	25.5	205.4	6.7	33
<b>25*</b>	2013-03-23	27.5	121.9	5.1	3
28	2015-03-10	12.7	284.8	9.0	24
<b>0*</b>	2013-04-01	11.7	86.3	1.2	58
6*	2013-04-16	20.3	230.6	11.1	54
18*	2013-04-25	5.5	236.4	15.0	35
<b>21</b>	2015-04-01	29.4	303.5	1.4	16
19	2014-06-22	6.6	345.0	4.4	60
9	2014-06-30	7.6	331.4	1.7	75
5**	2013-07-23	4.0	125.4	12.7	68
10	2014-07-31	7.7	263.1	4.6	94
4	2015-07-11	6.2	238.0	10.0	106
2*	2012-08-14	10.8	157.4	7.8	65
7**	2013-08-13	15.0	315.6	1.3	55
<b>24*</b>	2012-09-26	10.0	179.3	1.3	25
8*	2012-09-29	16.7	278.9	0.2	54
1	2014-09-12	7.4	62.6	5.9	94
<b>16</b>	2014-09-26	15.6	275.0	7.2	33
3	2015-09-22	7.7	261.6	1.4	106
13	2014-10-17	9.1	220.2	5.2	95
15	2014-10-26	22.9	242.1	6.9	51
23	2015-10-30	19.5	181.4	10.9	28
17	2013-11-28	13.9	318.9	8.8	14
26	2014-11-17	12.8	127.0	1.0	29
22	2015-11-19	29.0	280.1	1.9	17
11	2013-12-20	24.6	249.5	2.7	41
27	2014-12-17	20.6	262.9	5.6	25



# CHAPTER 4







## LOW-LEVEL JETS OVER THE NORTH SEA IN ERAS AND OBSERVATIONS: TOGETHER THEY DO BETTER

---

YOU CAN OWN THE EARTH, AND STILL; ALL YOU'LL OWN IS EARTH UNTIL... YOU CAN PAINT  
WITH ALL THE COLOURS OF THE WIND

- STEPHEN SCHWARTZ (FROM POCAHONTAS)



In chapter 2, low-level jets stood out as an important anomalous event. However, only the measurements from met mast IJmuiden were used. In chapter 3, the performance of three weather models was evaluated, among which the ECMWF integrated forecast system (IFS). However, low-level jets were not explicitly considered there.

This chapter addresses both limitations. ERA5, a state-of-the art reanalysis dataset based on the IFS is combined with additional lidar observations from ten offshore platforms to investigate low-level jet characteristics over the Dutch North Sea. The objective of this study is to combine the best of two worlds: (1) ERA5 data with a large spatio-temporal extent but inherent accuracy limitations due to a relatively coarse grid and an incomplete representation of physical processes and (2) observations that provide more reliable estimates of the measured quantity but are limited in both space and time. We demonstrate the effect of time and range limitations on the reconstructed wind climate, with special attention paid to the impact on low-level jets. Various methods to overcome these limitations are discussed as well.

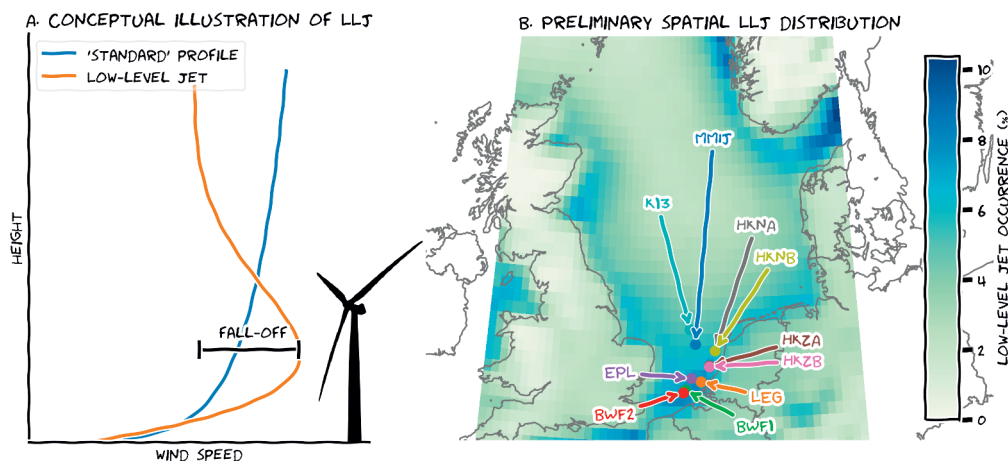
This chapter is published as:

Peter C. Kalverla, James B. Duncan Jr., Gert-Jan Steeneveld, and Albert. A. M. Holtslag (2019). Low-level jets over the North Sea based on ERA5 and observations: together they do better. *Wind Energy Science*, 4(2):193–209.

## 4.1 Introduction

On average, wind speed increases with height above the surface and the rate of increase can be described using simple formulas (e.g. power-law or logarithmic profile, see Sedefian, 1980). Due to their simplicity and ease of use, these *wind profile parametrizations* have been widely adopted in the wind energy community. However, in some situations these formulas cannot adequately capture the observed wind profile. During these situations, application of a simplified wind profile parametrization can introduce error or ‘uncertainty’ into the reconstructed wind climatology. This is clearly the case for low-level jets, for which wind speed reaches a maximum not far (i.e. roughly less than 500 m) from the surface (figure 4.1A). Wind shear and turbulence intensity associated with low-level jets also differ substantially from that assumed under ‘standard’ conditions.

Low-level jets modify wind power performance and loading by impacting wake recovery rates and vertical profiles of wind speed, direction and turbulence (Wharton and Lundquist, 2012; Bhaganagar and Debnath, 2014; Park *et al.*, 2014; Gutierrez *et al.*, 2017). Thus, for a complete assessment of loads and power, it is important to have a broad understanding of the site-specific low-level jet characteristics: how often do they occur, under which circumstances, at what height and with what strength, and what mechanisms are responsible for their formation? A large body of literature exists on low-level jets, the majority focusing on the onshore phenomenon. We refer to Rife *et al.* (2010) for a global climatology and to Shapiro *et al.* (2016) for a synthesis of the underlying mechanisms. In coastal areas, the occurrence of low-level jets has been attributed to the thermal contrast and differences in surface roughness between land and sea (e.g. Mahrt *et al.*, 2014; Nunalee and Basu, 2014). Dörenkämper *et al.* (2015) linked the occurrence of coastal jets to their onshore counterpart. In certain areas, other mechanisms like orographic forcing may play an important role (e.g. Moore and Renfrew, 2005). Concerning



**Figure 4.1:** A. Example low-level jet profile as compared to the ‘standard’ logarithmic wind profile. B. Preliminary spatial distribution of annual low-level jet occurrence based on ten years of ERA5 data up to 500 m. Overlaid are the location of the ten measurement platforms used in this analysis: Met Mast IJmuiden (MMJ), Hollandse Kust Noord A (HKNA) and B (HKNB), Hollandse Kust Zuid A (HKZA) and B (HKZB), Lichteiland Goeree (LEG), Borssele Wind Farm Lots 1 (BWF1) and 2 (BWF2), Europlatform (EPL) and K13. Colour coding is consistent across all figures in this chapter.

the spatial and temporal variability of the coastal jets we refer to Ranjha *et al.* (2013) and Lima *et al.* (2018), who presented global maps based on reanalysis data. Their analyses highlight a number of large-scale global ‘hotspots’ that, in effect, overshadow more regional phenomena. Consequently, a systematic long-term characterization of coastal jets is lacking for the North Sea.


In a previous publication (Kalverla *et al.*, 2017), we reported on low-level jet characteristics at a prospective wind power site 85 km off the Dutch coast (MMIJ, aka “IJmuiden ver”), using four years of mast and lidar observations. The climatology consisted of: the diurnal and seasonal variability in low-level jet occurrence, jet speed, jet height, jet direction, etcetera. Inherently, this low-level jet climatology is only valid for the single observation site examined. In order to generalize the results from this study, and to improve our overall understanding of low-level jets across the North Sea, we now present a spatial climatology of low-level jets based on ERA5 (section 4.2; C3S, 2017) reanalysis data and an extended set of observations.

Preliminary results based on ten years of data in the lower 500 m of the atmosphere (figure 4.1B) shows that ERA5 provides interesting information about the spatial distribution of low-level jets. However, without observational support this information is of little value. Therefore, we incorporate additional lidar observations to provide this support, but knowledge gained of the Dutch offshore wind climate from these measurements is inhibited by the relatively short duration of measurement collection (i.e. typically  $\sim 1$  yr) and the limited vertical measurement range (i.e. typically less than 300 m; see appendix A for details on measurement time and range). Consequently, the aim of this study appears twofold: (1) observations will be used to validate the ERA5 climatology of wind and low-level jets and (2) ERA5 will be leveraged to infer long-term low-level jet characteristics based on a limited set of observations. Absolute agreement in low-level jet characteristics between the two data sources would enable perfect execution of these objectives; however, that is unlikely. Therefore, we formulated the following research question to serve/blend both perspectives:

*How can observations and reanalysis data be combined to obtain a spatial climatology of low-level jets that is both rich (in its spatial and temporal extent) and reliable (in terms of its correspondence with available in-situ observations)?*

The paper is structured as follows. A brief description of the data and an elementary evaluation of wind speed itself is provided to illustrate how both datasets are biased. Thereafter, low-level jet representation within both datasets is discussed, starting with jet detection and morphology (e.g. jet height). A common thread throughout the paper is how these characteristics are impacted by time and (vertical measurement) range limitations. Using the seasonal cycle of low-level jets as an illustrative example, we experiment with various methods to post-process the ERA5 data and extend the observations based on identified correspondence and/or differences. This exercise is repeated for the diurnal cycle, atmospheric stability and various circulation patterns. Finally, all of these characteristics are combined to demonstrate that the ‘true’ low-level jet rate can be reconstructed with reasonable accuracy if sufficient observations are available. The paper ends with a comprehensive discussion of the implications and future research directions.

The focus of this paper is to obtain a reliable spatial representation of the low-level jets. This provides clues as to the physical mechanisms that govern them, but a detailed treatment of these processes is outside the scope of the current work.



**Figure 4.2:** A. Time-height plots of wind speed for each platform, illustrating the data collection periods, temporal overlap between platforms and episodes of missing data. B. Site-specific measurement heights. Reference elevation for the ERA5 data have been included for comparison. The colour coding in A highlights episodes of high (yellow/green) and low (blue) wind speed.

To facilitate transparency and reproducibility, a series of Jupyter notebooks is available as supplementary material to Kalverla *et al.* (2019a). Consequently, some technical details are left out of the main text, that is intended as a pleasant and coherent treatise of the most important results.

## 4.2 A brief description of both datasets and their shortcomings

Observations are available from seven sites (figure 4.1B). Three of these sites had two lidars operating simultaneously and one site (MMIJ) also featured a 90 m met mast. The temporal span of measurements ranges from six months to over four years (figure 4.2A). Some of the lidars were placed in the vicinity of existing wind farms, and are appropriately filtered to remove any potential wind farm wake effects. More information on quality control and post-processing of the lidar data can be found in appendix A. The observations are available as 10-minute averages, but to facilitate comparison with ERA5, the data were converted to hourly averages.


ERA5 (C3S, 2017) is the latest reanalysis dataset from the European Centre for Medium-range Weather Forecasts (ECMWF). Re(trospective) analysis is the procedure of fitting a state-of-the-art weather model to historical measurements (e.g. satellites, weather stations, etc.) to obtain a long-term dataset that is both spatially and physically consistent and depicts the state of the atmosphere as it evolved through time. ERA5 is the successor of ERA-interim, and similarly ERA5 is expected to be widely used for wind resource assessment studies (Olauson, 2018). Compared to its predecessor, ERA5 has a finer horizontal grid of about 30 km and also enhanced vertical resolution (for this study, data was retrieved on a  $0.3^\circ$  by  $0.3^\circ$  latitude/longitude grid). ERA5 is based on a newer model version and moreover, provides output at hourly intervals, enabling a comprehensive analysis of sporadic features such as low-level jets. ERA5 data from the North

Sea domain between 2008 and (end of) 2017 in the lowest 500 m demonstrates the ability of the model to resolve low-level jets (figure 4.1B).

Before analysing the morphology of these jets, we illustrate the limitations of both datasets concerning the representation of wind speed. Figure 4.3A shows averaged wind profiles for the grid points closest to each of the measurement locations (we verified that this approach is comparable to spatial interpolation between multiple neighbouring grid points). The full lines represent all ten years of ERA5 data<sup>1</sup>, whereas the dashed lines indicate averaged wind profiles derived from data subsets, which only incorporate ERA5 data when observations are available. The full lines are all quite close together, while the data subsets exhibit a much larger spread. Variability between the full lines can be related to physical differences between sites (e.g. distance to coast). Dissimilarity between the ERA5 ten-year datasets and the ERA5 data subsets indicates that, due to the limited time extent of the observations, the data subsets are not representative of the site climatology. For some sites, this *representativity bias* reaches up to almost  $2 \text{ m s}^{-1}$ , and even for MMIJ, wherein measurements occurred for the longest period, it still amounts to  $\sim 0.5 \text{ m s}^{-1}$ . The primary reason for this bias at MMIJ is that the data contain more winter than summer months, and the wind is generally stronger in winter. Because the MMIJ dataset spans more than four years, some of it can be discarded in order to ensure an equal representation of the seasons within the data. However, at the other stations, the temporal period of observation is limited, and using a similar seasonality filter would result in almost half of the data being removed, which is not desirable. Worse still, HKN observations do not encompass a complete year, and even if they did, inter-annual variability can be substantial. Available observations therefore cannot be used to derive the long-term wind climatology directly. However, by correlating a short-term dataset with long-term observations at a nearby site, the long-term wind characteristics at the target site can be inferred with reasonable accuracy. This procedure is known as measure-correlate-predict (MCP, see Carta *et al.*, 2013). While not discussed here, application of similar techniques to the low-level jet phenomena will be examined later in this document.

ERA5 also demonstrates bias in its representation of site winds. An *error diagram* of the wind speed in ERA5 (subsets) versus observations is provided in figure 4.3C. In this diagram (co-opted from Kalverla *et al.*, 2019b), the mean error (BIAS) is plotted on the x-axis, the standard deviation of the error distribution (STDE) is plotted on the y-axis and, by virtue of the relation  $\text{BIAS}^2 + \text{STDE}^2 = \text{RMSE}^2$ , the distance to the origin represents the root mean square error (RMSE). Wind speed data from all observation levels were aggregated in this figure to evaluate the overall performance of ERA5 at each measurement site. For example, the HKZ lidars show a strong bias (i.e. systematic error), but have a relatively small standard deviation (i.e. random error). ERA5 site-specific RMSE values, ranging from  $1.25$  to  $1.5 \text{ m s}^{-1}$ , can be caused by multiple model aspects such as the limited grid resolution and the incomplete representation of physical processes. Uncertainties in the observations can also contribute to overall error statistics. Based on the manufacturer information and previous validation (Maureira Poveda and Wouters, 2015), the uncertainty in the observations can only account for about 2 % of the errors. Finally, displacement in space or time, as well as discrepancies between point-based measurements and modelled control-volumes can contribute to errors, although we did our best to minimize these effects, e.g. by using appropriate time-averaging of the observations (see supplement of Kalverla *et al.*, 2019a).

<sup>1</sup>Some lines are exactly on top of each other because they are in the same grid point. Both are plotted, though, to preserve colour coding



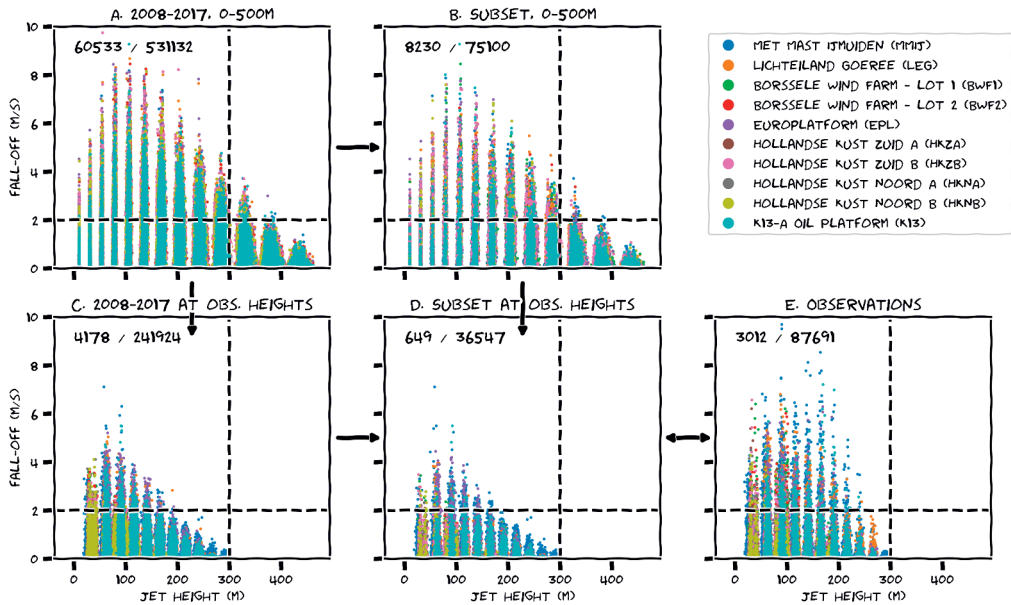
**Figure 4.3:** A. Averaged wind speed profiles for each measurement location, based on ten years of ERA5 data (full lines) and data subsets (dashed lines). B. Mean (full lines) and standard deviation (dashed lines) of the error between ERA5 (subsets) and the observations, for each measurement site, as function of the time of the day. C. Error diagram of wind speed in ERA5 (subsets) versus observations for all lidar datasets. Colour coding is the same in all subplots, so C can serve as legend.

The observed biases exhibit a strong diurnal variation. During the night (figure 4.3B), the bias is roughly between 0 and  $-0.5 \text{ m s}^{-1}$ , depending on the location. However, at 10 UTC, there is a sharp decrease in the bias of  $\sim -0.5 \text{ m s}^{-1}$  for most stations. The reason for this discontinuity can be found in the IFS (Integrated Forecasting System) documentation (ECMWF, 2016). ERA5 is produced with a 4D-VAR data-assimilation algorithm that uses two twelve-hourly windows running between 9–21 and 21–9 UTC. This means that all hourly fields up to the 9 UTC analysis are based on the nighttime observations, while data from 10 UTC onwards are based on the daytime observations. We hypothesize that the impact of the data-assimilation is magnified during the nighttime, because nighttime boundary layers are generally shallower; the difficulty of appropriately assimilating observational data within the (stable) boundary layer is discussed in Reen and Stauffer (2010) and Tran *et al.* (2018). Discontinuity in the diurnal cycle is present at each model level up to 300 m, irrespective of the season and platform; however, it seems to be slightly stronger for those stations closer to the coast.

### 4.3 Jet detection: a precarious procedure

Low-level jets are identified by seeking local maxima in the wind profiles. Having identified a local maximum, the jet strength, height and fall-off are analysed. Fall-off, as indicated in figure 4.1A, is defined as the difference between the maximum and the subsequent (moving upwards) local minimum or, if no local minimum is present, the top of the wind profile. Most results in this study are based on an absolute fall-off threshold of  $2 \text{ m s}^{-1}$ . Figure 4.4 demonstrates how this threshold influences the low-level jet detection rate, and further how the detection of low-level jets is influenced by both time and (vertical measurement) range limitations. The figure consists of five scatter plots, each depicting the fall-off versus the jet height for each wind profile that was detected with a local maximum. The differences between the panels are the underlying data analysed - i.e. observations and varying subsets of ERA5 data.

The first panel (figure 4.4A) is based on ten years of ERA5 data and the model levels contained within the lower 500 m of the atmosphere. The two dashed lines represent limiting factors: (1) the fall-off threshold of  $2 \text{ m s}^{-1}$  (horizontal dashed line) and (2) limitations due to observation height (vertical dashed line). The model data extend up to 500 m, but the observations reach only

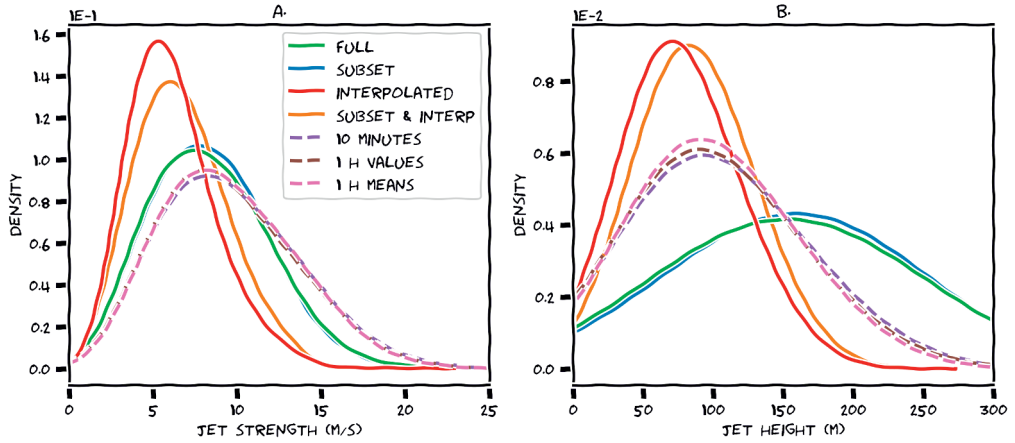


**Figure 4.4:** Scatter plots of fall-off versus jet height for various representations of model data, and observations. In A and B, the jet height is represented by discrete model levels. Since these are specified in terms of pressure rather than height, they can exhibit small height variations in time. In C, D and E, jet height is represented by fixed measurement heights, and to improve readability of the graph we added small random perturbations to these heights. See text for further explanation of the figure.

up to about 300 m (depending on the platform). All platforms are overlaid (shorter datasets on top). Only points above the horizontal dashed line are included in the low-level jet climatology that is presented in the next sections. The numbers in the top left corner of each panel give the number of jets above the fall-off threshold and the total number of jets plotted.

Figure 4.4B-D are based on subsets of the ERA5 dataset. In panel B, ERA5 data is incorporated only if observations are available; as expected, this substantially limits the total number of low-level jets (85 % reduction). In panel C, we have retained all ten years of data, but only at observation heights (i.e. data above 300 m were discarded and the remaining data were vertically interpolated – using a cubic spline – between the remaining model levels to obtain the ERA5 wind speeds at the exact observation height). The effect of this step is that 93 % of the meaningful jet events (i.e. those exceeding the fall-off threshold) vanish, and not just those above 300 m. In order to classify a wind profile as a jet, fall-off above must be properly resolved. This explains why a jet at 100 m can also vanish from the climatology if data from above 300 m are removed. The pronounced impact of this vertical range limitation on the ERA5 data raises the question whether the *observed* low-level jet climatology would be much different if we could observe higher-altitude winds. Increased measurement range might reveal not only low-level jets above hub-height, but also new low-level jets at hub-height that are currently not identified as such.

Height and time limitations are combined in panel D in order to develop an ERA5 dataset that is fair to compare with observations (panel E). Judging from the figure, it seems that ERA5 does not perform well. Much fewer jets are found above the fall-off threshold in the ERA5 data as compared to the observations. Indeed, a more quantitative comparison in the form of a con-



**Figure 4.5:** Kernel density estimates of the probability distribution of jet strength (A) and jet height (B) for various representations of the ERA5 data (full lines) and the observations (dashed lines), aggregated over all stations.

tingency table, based on one-to-one (1:1) jet correspondence between the two datasets, shows a very low critical success index ( $\sim 0.2$ ) and probability of detection ( $\sim 0.2$ , see supplement of Kalverla *et al.*, 2019a). In other words, only 20 % of low-level jets are correctly represented by ERA5. Does that imply that ERA5 is useless? No! Figure 4.4A indicates that potentially relevant information was filtered out. Even though the fall-off is typically much smaller (to the extent that it falls below the fall-off threshold), the height distribution of the ERA5 jets seems similar to the observations (also see section 4.4). Perhaps the ERA5 jets appear vertically displaced or just not strong enough? This would not come as a surprise: weather models have long been known to generate excessive vertical mixing under stable conditions, effectively ‘smearing out’ low-level jets (Holtslag *et al.*, 2013). If the height thresholds for the ERA5 data is modified to 500 m, the 1:1 correspondence is still quite poor (critical success index  $\sim 0.2$ ; probability of detection  $\sim 0.5$ ), but despite an inability to accurately denote total number of low-level jets, other characteristics appear to be captured quite well – e.g. the average monthly low-level jet rate. Therefore, the remainder of this paper is devoted to the analysis of such low-level jet characteristics and methods to consolidate ERA5 and measurement data.

## 4.4 Vertical range affects perceived jet morphology

Jet height and jet strength are of paramount importance for wind energy applications. Small variations in height can result in either symmetric or asymmetric loads on the turbine, and typical strengths in the rated part of the power curve are probably less critical than typical strengths in the cubic part. It turns out, though, that the concepts of ‘typical’ height and strength are not self-evident.

Figure 4.5 displays probability distributions of jet strengths (panel A) and jet heights (panel B) for various representations of the ERA5 data and observations<sup>2</sup>. It shows that the jet height and

<sup>2</sup>Obviously, it is physically impossible to have a jet strength or height below zero. This is an artefact of the visualization - it has a smoothing effect. We experimented with other visualizations (smaller bandwidth, or histograms), but found that this visualization best represented the underlying data.



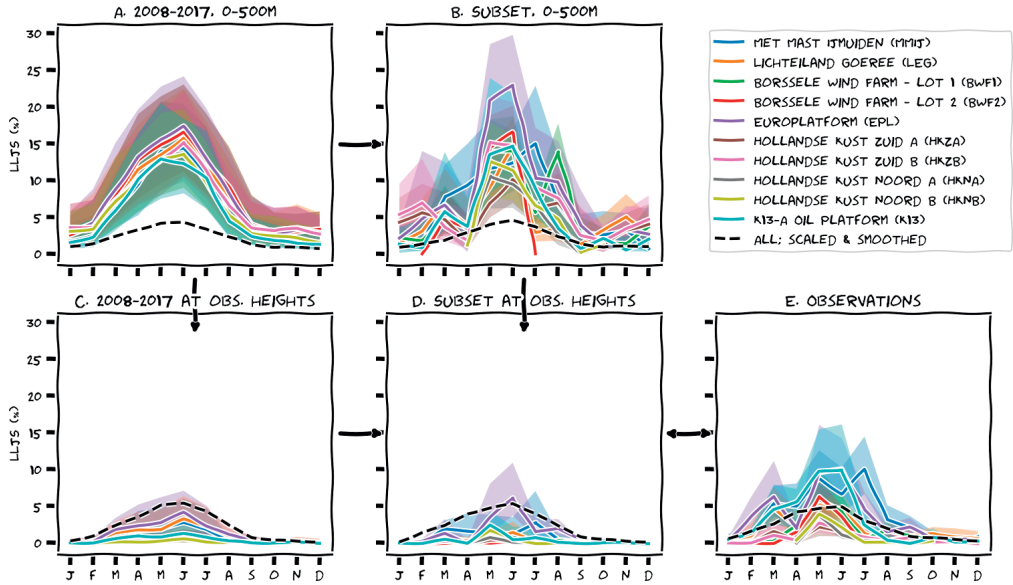
strength distributions are sensitive to the range limitation. The median observed jet strength is about  $8 \text{ m s}^{-1}$ . This is quite well reflected in the ERA5 data if we consider all levels up to 500 m, but after imposing the range limitation, the jet strength is underestimated by about  $3 \text{ m s}^{-1}$ . The observed median jet height is around 80 m. The ERA5 jet height distribution is broader with greater jet heights for the data up to 500 m, while it is narrower with lower jet heights for the range-limited data. To obtain a robust result, this figure is based on the aggregated data from all platforms. Separate figures for each individual platform show similar characteristics, although the jets near the coast seem to be somewhat closer to the surface than jets further offshore (not shown).

Three different representations of the observations are included in figure 4.5. The first one is based on the ten-minute data. The second is based solely on the data of each full hour; in other words, we discarded five-sixth of the data. With this strategy, (small) discrepancies in low-level jet timing can have a disproportionate impact on the results. A more permissive evaluation (the third representation) is based on hourly averages obtained with a sliding window, where each full hour is an average including the ten-minute data from the preceding two and the following three time stamps. This last version of the observations is used throughout the remainder of the paper. This figure demonstrates that the differences between various resampling methods in terms of jet height and jet strength are small.

## 4.5 Datasets agree: most jets in spring and summer

Figure 4.6 displays the seasonal cycle of low-level jets and, in a similar fashion as figure 4.4, how this cycle is subject to time and range limitations. Over ten years' time and 500 m (panel A), the seasonal cycle is smooth and differences between the individual platforms are small. Ideally, we would compare this to ten years of observations up to 500 m, but since those data are not available we take spatial and temporal subsets of the ERA5 data instead. By investigating how this affects the seasonal cycle, we identify methods to extend upon the limited observations. Over the shorter measurement periods (panel B) the seasonal cycle appears much more erratic than the 10-year climatology. Some years are not very representative, and some datasets do not even cover a complete cycle. As we will see later on, a favourable weather pattern for low-level jets is a weak large-scale forcing typically associated with high-pressure systems. Such 'blocked' weather patterns can last for several weeks, and their occurrence can thus cause large differences in monthly low-level jet rates. In other words, the seasonal cycle based on only one or a few years is very sensitive to inter-annual variability. Upon vertical subsetting/interpolation to measurement heights (panel C) the seasonal cycle is still visible, albeit with a much smaller amplitude. The combined effect (panel D) leads to a very uninformative climatology, because the monthly low-level jet rates are all (close to) zero except for some unrepresentative spikes. Based on panel B, we expect that the observations are similarly affected by the limited time window of the observations. Indeed, panel E shows an erratic seasonal cycle with an amplitude somewhere between panels B and D.

Thus, both datasets agree on the presence of an annual cycle, but the amplitude differs between (various representations of) ERA5 and the observations. Moreover, the observation periods are too short to obtain a reliable climatology. To distil a more robust signal from the observations, we combined the data from all sites before computing the monthly means, and smoothed the resulting signal with a moving average of three months. The result is the dashed black line



**Figure 4.6:** Seasonal cycle for various representations of model data, and observations. Shading is the sensitivity to  $\pm 0.5 \text{ m s}^{-1}$  for the LLJ fall-off threshold. The dashed lines represents an aggregated seasonal cycle of all platforms, smoothed with a rolling average of 3 months (2 at the edges) and scaled with the ratio of the mean jet frequency in the respective representations of ERA5 and the mean jet frequency in the observations.

in panel E. We then repeated these steps for the ERA5 data (panels A-D), but before plotting these lines, we scaled them with the observations, using a fixed scaling factor that is simply the ratio between the mean low-level jet rate in the respective representation of ERA5 (panel A-D) and the mean of the observations (panel E). The result is promising: the seasonal cycle is similar for all datasets, peaking at about 5% in June. The crude manipulation of the data leads to a large error margin, though, and we wonder whether we can find a more sophisticated approach to achieve a similar result. Furthermore, because valuable information is lost if we discard the ERA5 data above observation heights, we will continue to work with the ERA5 data up to 500 m in the remainder of this paper.

## 4.6 Simple scalings for the seasonal cycle

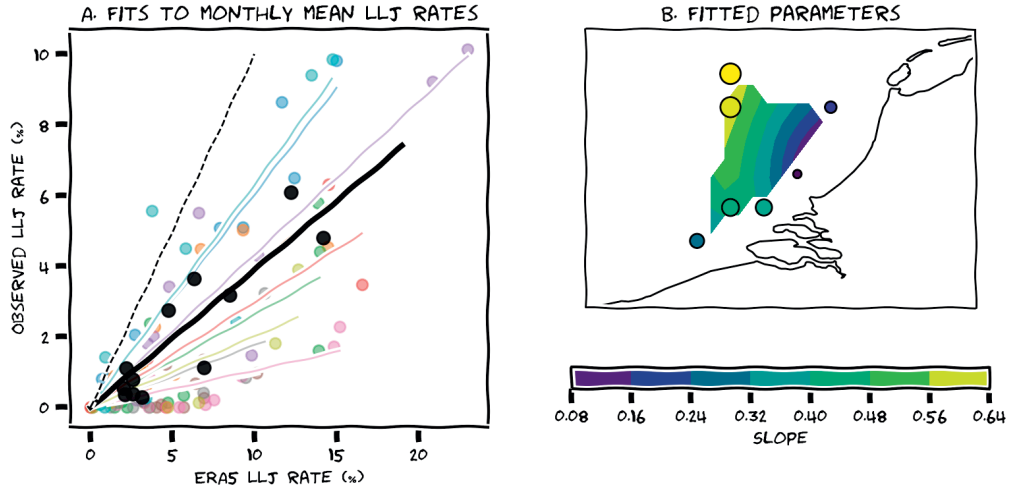
In the previous section we learned that ten years of ERA5 data lead to a smooth seasonal cycle, but shorter observation periods lead to an erratic seasonal cycle because the months in the subset are not representative of the long-term monthly means. We also saw that upon aggregation and smoothing, both ERA5 and observations show similar seasonal cycles that differ mostly in their amplitudes. In this section we seek to combine the information from both data sources to reconstruct the ‘true’ seasonal cycle of low-level jets over the North Sea. We considered two different approaches.

The first method applies a correction to the observations, based on information about their representativity. For each month and each platform, we calculated the ratio between the low-level jet occurrence in the full- and subsets of the ERA5 data. Months for which this factor is much

smaller (or larger) than 1 are characterized by above(below)-average low-level jet occurrence. We then applied these ratios as correction factors to the observed monthly means to adjust the outliers and obtain a more representative seasonal cycle. However, this method did not lead to satisfactory results, because the correction factors were not robust: if only one year of data was available, and a month was very unrepresentative, the correction factor would become very high/low and the adjustment would overcompensate. Consequently, the reconstructed long-term seasonal cycles still appeared erratic and were deemed unreliable (this result is therefore not shown here, but is available as supplement 4/6 in Kalverla *et al.* 2019a). For MMIJ the measurement period spanned more than four years and consequently, the monthly low-level jet occurrence already started converging to the climatological seasonal cycle. For this platform, the correction factors were closer to 1 and we obtained a reasonably smooth seasonal cycle. This emphasizes that for this correction method, at least several years of measurement data are required obtain a reliable estimate of the long-term low-level jet climatology.

Whereas the first method was aimed at correcting the observations (using ERA5 as a ‘vehicle’ to assess their representativity), with the second method we aim to correct the long-term ERA5 data based on prior evaluation of its performance during the short-term period for which we have observations. This can be readily understood from figure 4.6. We compare panels B and E, and seek a fixed scaling factor that minimizes the difference between each pair of monthly observed and simulated LLJ frequencies. Denoting the *monthly mean* low-level jet frequency in ERA5 and collocated observations with  $\mathbf{x}$  and  $\mathbf{y}$ , respectively, an optimized scaling factor can be found by solving for  $a$  in  $\mathbf{y} = a\mathbf{x}$  (using linear least squares regression). We do this for each platform individually, and also for their combined signal.

The results are illustrated in figure 4.7A. The lighter colours represent the individual platforms, while the black line and scatter points represent the combined monthly means. The overall fit, based on all available data, has slope 0.44, but there are substantial differences between the individual platforms, with slopes between 0.15 and 0.73 and relatively large scatter. The difference between platforms could be random, due to the limited availability of measurement data, or systematic, in which case different sites need different scaling parameters. If the difference is random, the global optimum indicated by the black line in figure 4.7A could do justice to all individual platforms, because it incorporates a much larger body of measurement data than any single-site regression. Applying this factor of 0.44 to the full ERA5 data provides us with a smooth seasonal cycle with reduced amplitude (similar to the black dashed line in (figure 4.6A), but this time based on an optimized scaling factor). In other words, the seasonal cycle of low-level jets based on ERA5 data up to 500 m overestimates the observed cycle (based on measurement up to 300 m) by a factor of  $\sim 2$ . However, as shown in figure 4.7B, there seems to be a spatial dependence in the scaling factors with larger slopes away from the coast, implying that the different sites need different scaling parameters. In order to cross-validate the single-platform regressions, we need to split the measurement data in train and test datasets, but this poses a challenge. Like before, the data record at MMIJ is long enough to obtain a reasonable prediction of the test data, but some of the other data records are very short and splitting them would leave e.g. only three months of training data, which obviously leads to very poor statistics, especially since there are hardly any low-level jets in winter. Without cross-validation more data is available for regression, but this introduces the risk of over-fitting and therefore quantitative evaluation will be biased. Qualitatively, the resulting seasonal cycles still appear erratic (see supplement 4/6 of Kalverla *et al.*, 2019a)).



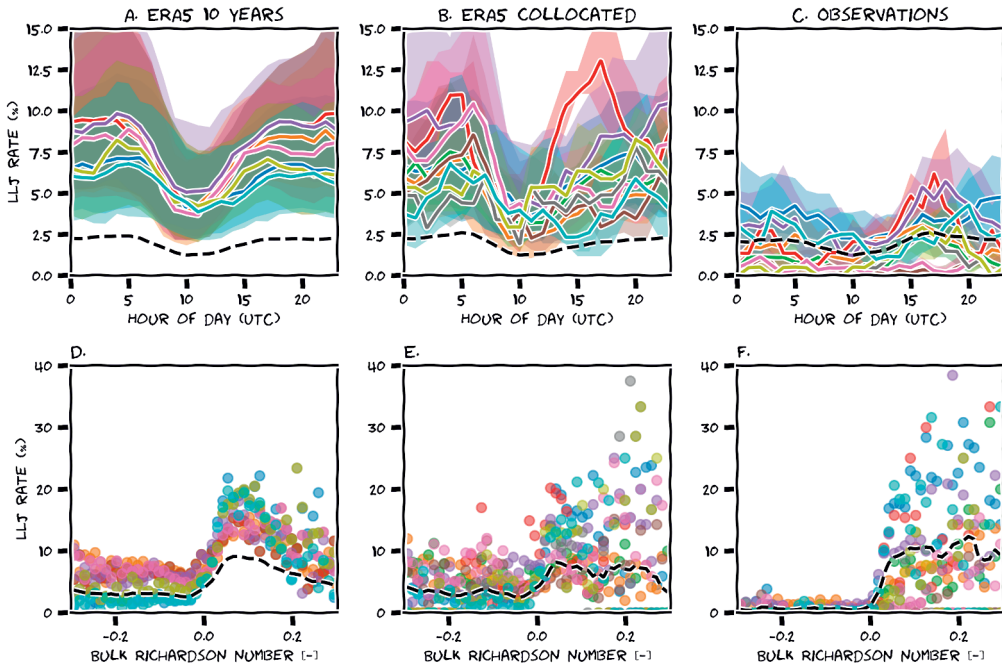
**Figure 4.7:** A. Illustration of linear regression between monthly low-level jet rates in the ERA5 data (subset, up to 500 m) and the observations. Black line and scatter points represent aggregated data of all platforms, while the other colours correspond to fits for individual platforms. Dashed black line indicates a 1:1 correspondence. B. Spatial distribution of the obtained fit parameters for each individual platform. Like the colour coding, marker size is scaled with the slope of the regression.

Thus, despite similarities between the datasets, it is not straightforward to either correct the observations using ERA5 representativity factors, or to correct the ERA5 data using a scaling factor derived from collocated observations. In this section, we used the seasonal cycle to obtain aggregated low-level jet characteristics (i.e. monthly means), but perhaps we can identify other characteristics that lead to better results.

## 4.7 Other jet characteristics and their scaling potential

### 4.7.1 Diurnal cycle and stability

After analysing the seasonal cycle of low-level jets in-depth, we now briefly consider some other variables that describe relevant characteristics of the low-level jet climatology, starting with the diurnal cycle. Figures figure 4.8A-C are again similar to figure 4.6, now only including the ERA5 data up to 500 m. From the observations, it appears that the low-level jets occur throughout the day, but with a small dip around 11 UTC. Panels B and C, based on short temporal subsets, are so erratic that it is difficult to distinguish this diurnal cycle by eye. After aggregating all platforms and smoothing the data (black dashed lines), we find that the observations and ERA5 agree on the general shape, but again we needed to scale the ERA5 signals because they differed in magnitude: the diurnal cycle in ERA5 is much more pronounced. At this point, we think it is good to stress that several mechanisms can lead to low-level jets in coastal areas (see sections 4.1 and 4.9), and the resulting diurnal signature should not be confused with that of the typical onshore nocturnal jet that is often found over land. As in the previous section, we performed linear regression to identify optimal scaling parameters for the dashed black lines in panels A and B. The difference with the previous section is that the regression is now based on pairs of *hourly* instead of monthly observed and simulated low-level jet frequencies. The scatter in this



**Figure 4.8:** Average low-level jet rate for each hour of the day (A, B, C) and as function of the Bulk Richardson number (D, E, F), for the full (ten years) of ERA5 data up to 500m (A, D), a subset of this data collocated with the observations (B, E), and the observed data (C, F). As in figure 4.6, the black dashed lines represented a scaled and smoothed aggregated signal based on all platforms.

data is larger than for the seasonal cycle, but the spatial distribution of the fitting parameters is similar (not shown).

The second row in figure 4.8 shows the relation between low-level jet occurrence and atmospheric stability (expressed by the bulk Richardson number based on the ERA5 surface data: two-metre temperature, skin temperature, and ten-metre wind). Scatter points represent mean aggregated low-level jet frequencies over 50 stability bins. Both ERA5 and the observations agree that low-level jets are typically associated with stable stratification, although for some platforms in panel D and E, there seems to be a substantial number of jets for unstable conditions as well. In the subsets (panel E) this distinctive behaviour is not as clear, and in the observations it seems mostly absent. Without going into detail, we note that low-level jets can be formed by different mechanisms, and it is possible that ERA5 represents one mechanism better than another, or perhaps one mechanism is actually over-represented. Also note that in panels E and F there are (positive) values of the Richardson number for which no low-level jets are observed. In panel D, this is not the case, which indicates that the measurement periods are too short to adequately sample the full range of stability conditions. Finally, we note that in panel D, the low-level jet rate seems to decrease again for very stable situations. This could be an artefact of the bulk Richardson number, or a physical limit: a stable atmosphere leads to a low-level jet, but the low-level jet produces wind shear and consequently, the bulk Richardson number decreases. The fact that this behaviour is not reflected in the observations suggests that the true stability (that would have been observed) was actually smaller than what ERA5 predicted. Again, we tried to scale the amplitude of the stability signature by performing linear

regression between pairs of low-level jet frequencies in ERA5 and observations (now based on stability bins instead of monthly or hourly groupings). The slopes are larger than those based on the seasonal and diurnal cycle ( $\sim 1.0$ ), but qualitatively they seem to be less robust (not shown).

#### 4.7.2 Weather types and the spatial distribution of low-level jets

We also investigated the relation between low-level jet frequency and typical circulation patterns. We used Lamb Weather Types (LWT; Jones *et al.*, 2013) to perform this analysis. To derive these weather types we used the ERA5 mean sea level pressure on a  $5^\circ$  latitude/longitude grid of 16 points as laid out in the appendix of Jones *et al.* (2013), but centred over the area of interest. The method distinguishes three main groups: those with a dominant cyclonic (anticlockwise, low-pressure area) circulation, those with a dominant anticyclonic (clockwise, high-pressure area) circulation, and those with a ‘pure directional’ flow. These three groups are further subdivided based on the main direction of the flow over the North Sea (north, north-east, east, etc.). If there is no dominant direction, the LWT is ‘pure (anti)cyclonic’. Pressure fields characterized by the absence of a dominant forcing are ‘undefined’. In total this yields 27 different circulation patterns. We computed average low-level jet rates for each group.

To illustrate the association between the circulation type and the low-level jet occurrence, figure 4.9 shows the average low-level jet rate per weather type in the North Sea domain, based on ten years of ERA5 data up to 500 m. The streamlines show the dominant flow pattern for each weather type: the columns represent different wind directions over the North Sea, while the full rows represent different rotation types. In the first full row, the rotation is predominantly clockwise, in the bottom full row, the rotation is mostly anticlockwise, and the middle full row is characterized by the absence of rotation. Notice how the same wind direction can be associated with different large-scale flows – and how this can impact the low-level jet rate. Like before, we will not go in-depth on each individual feature in this figure, but we will focus on overall characteristics. In general, we see that low-level jets concentrate along the coastlines. This extends and refines the global findings of Ranjha *et al.* (2013) and Lima *et al.* (2018) for the North Sea domain. Low-level jets are much more dominant for certain Lamb Weather Types. Most notably, the ‘undefined’ weather type often gives rise to the formation of jets. This makes sense, as low-level jets are subtle phenomena, and the absence of a strong large-scale flow eases their development. Furthermore, we observe that low-level jets occur frequently during large-scale flows with a pronounced easterly component. Note that easterly flows bring in continental air, while westerly flows originate from the Atlantic. Low-level jets are uncommon for westerly flows. Closer inspection reveals that the differences in spatial distribution of the low-level jets (e.g. comparing the Dutch and Norwegian coastlines) seems to be related to whether the large-scale flow is directed offshore. The British isles are different in this respect, since for westerly flows we do not observe an increased low-level jet rate off the eastern coast of the UK.

Like with the previous characteristics, we performed linear regression between ERA5 and observed low-level jet frequency, this time aggregated over the various Lamb weather types. We found similar patterns in ERA5 and the observations (not shown), but the spatial distribution of the scaling parameters is different. Most slopes are around 0.4, but LEG stands out with a slope of 0.65. This is not a huge difference, but it implies that our earlier hypothesis – that the slope increases with distance to coast – does not hold for all predictors. Indeed, one could argue that with Lamb weather types as predictor, the scaling parameters are spatially more robust. Thus,

while we believe that the spatial distribution in figure 4.9 is actually meaningful, the absolute low-level jet rate (as indicated by the colour bar) is still off by a factor of  $\sim 2$ .

## 4.8 Combining multiple predictors to extend observations

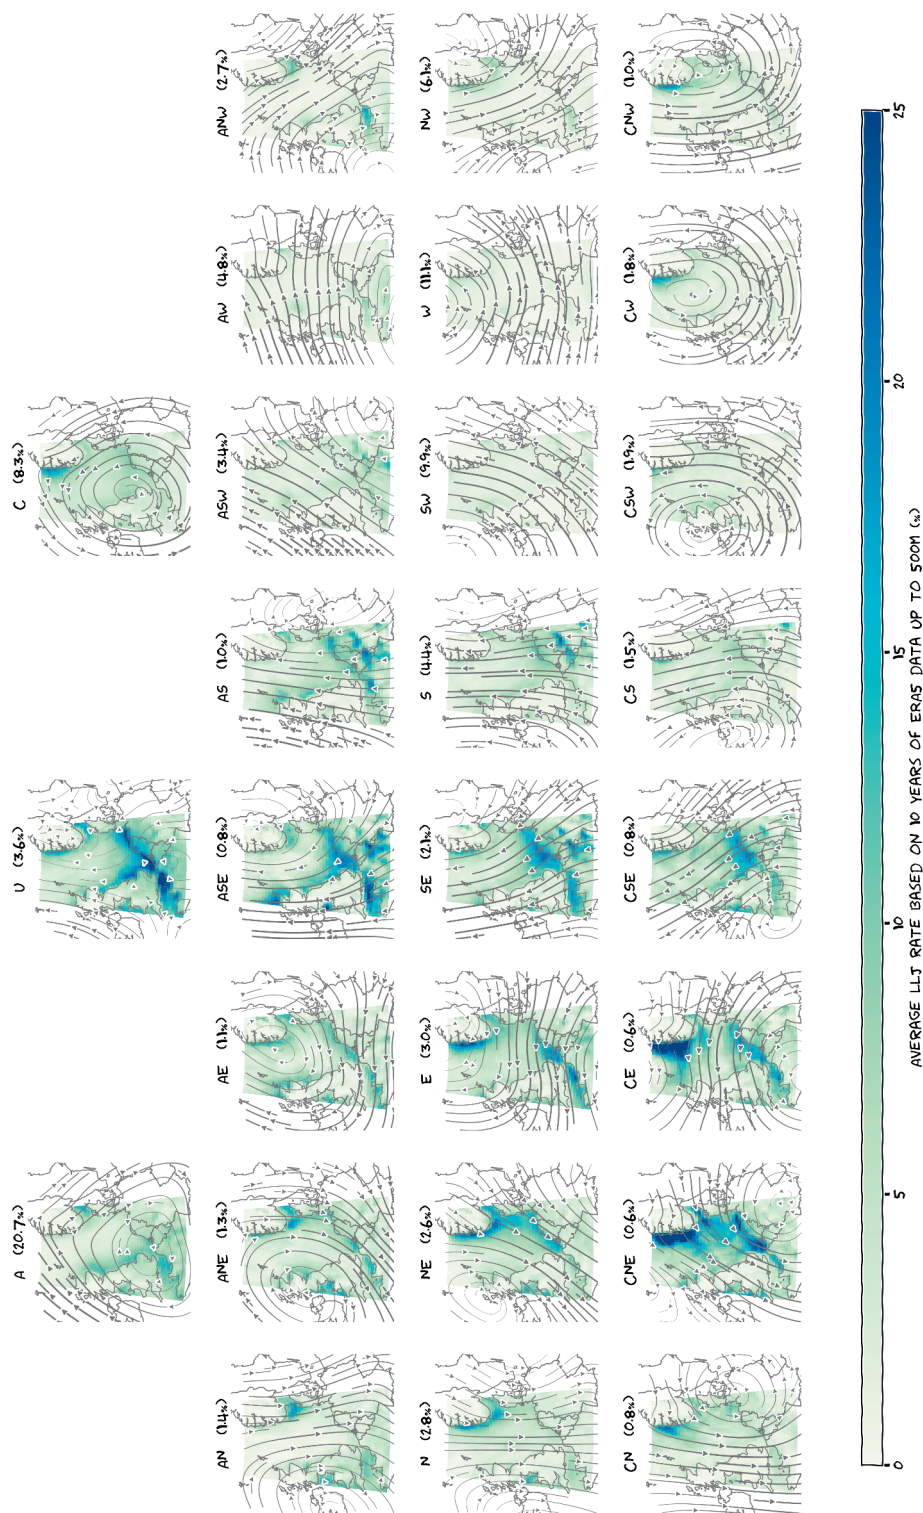
So far, we have tried to scale the low-level jet climatology with simple linear factors applied to individual characteristics (e.g. seasonal cycle). Perhaps, we can find a more sophisticated transformation function by combining multiple predictors? In this section we use the MMIJ data to illustrate how this could be applied in practice. In contrast to the previous sections, which focused on aggregated low-level jet *frequencies*, here we consider *individual* wind profiles. The procedure resembles the Model Output Statistics (MOS) forecasts that are widely used for weather forecasts (e.g. Glahn and Lowry, 1972; Carter *et al.*, 1989; Wilks, 2006, chapter 6.5.2) and is similar to the Measure-Correlate-Predict methods mentioned in section 4.2 (Carta *et al.*, 2013). We use a machine learning package to perform this task and for readability, we will not highlight all the technical details here. However, Jupyter notebooks are available as supplement to Kalverla *et al.* (2019a) to facilitate reproducibility.

The general idea is illustrated in figure 4.10A: we have a short time series with observations and a long reanalysis dataset. Based on the overlapping part of the data, we determine the optimal parameters of a statistical model (depicted by the red box). We then use this model to predict the value of the observations, given the available long-term reanalysis data. In the illustration, it seems as though one reanalysis variable is used for this purpose, but in fact, we can use as many variables as we want. In our case, the variable we want to predict is the probability that a low-level jet will be observed, given various predictor variables from the ERA5 data. Because this is a binary outcome (a jet either occurs, or not), our model of choice is a logistic regression model, which predicts the probability of a positive outcome as function of one or several predictor variables. The general form of this model is

$$p = \frac{1}{1 + e^{-(\beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots)}}$$

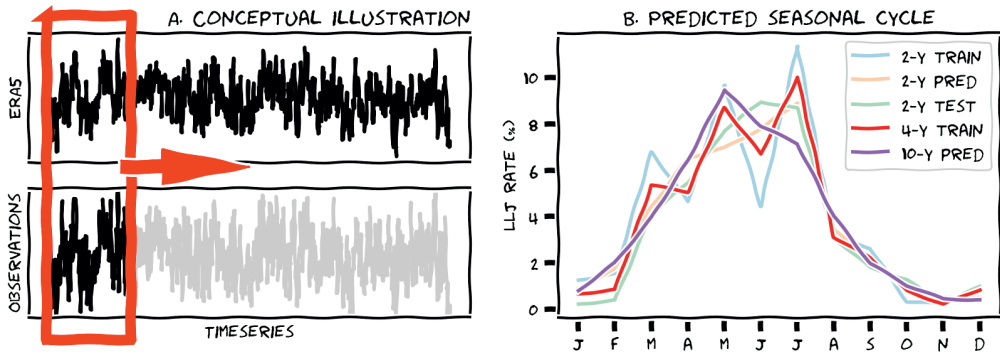
where  $\beta_i$  are the coefficients of the corresponding predictor variables  $x_i$ . In a short exploratory phase, we experimented with various combinations of predictor variables. We found the best performance for a small set of predictor variables consisting of: time of the year, atmospheric stability, and Lamb weather type. This makes sense, as together these variables encompass information about wind speed, direction and history of the flow, as well as the probability of stable stratification and baroclinic conditions. Indeed, each of these variables alone already provided valuable information in the previous sections. For optimal performance, these variables were preprocessed as follows: to truthfully represent its cyclic nature, date was encoded by splitting the day of year in a sine and cosine contribution. The Lamb weather type is a categorical variable, and to make it suitable for regression it was encoded by converting it to the binary representation of the numbers up to 27 (the total number of weather types) and treating each digit as an individual binary variable. Stability was represented by the difference between the two-metre temperature and sea-surface temperature, which provided better results than the bulk Richardson number. We also experimented with various training algorithms to determine the coefficients  $\beta_i$  of the logistic model (intermediate results can be found in the supplement of Kalverla *et al.* 2019a). In the end, we settled on a stochastic gradient descend algorithm.





**Figure 4.9:** Spatial distribution of low-level jet rate in ERA5 data. As explained in the text, the values shown here overestimate available observations and should be interpreted with caution. A=Anticyclonic, C=Cyclonic U=Undefined; N, NE, E, etcetera are eight wind direction sectors; combinations of a direction and a rotation type are 'hybrid' weather types, while weather types without a dominant rotational component are 'pure directional'. Streamlines illustrate the dominant large-scale flow pattern (averaged over each LWT). The relative occurrence of each weather type is indicated as well. A larger version of this figure will be available online.





**Figure 4.10:** A. Illustration of the MCP/MOS/ML procedure in which a (logistic) model is trained with observation data and then used to predict long-term characteristics. B. Illustration of the MMIJ seasonal cycle of low-level jets based on: two years of observed data (train), two years of predicted (pred) and observed data (test), four years of observed data (train) and ten years of predicted data (pred).

First, we took only half of the MMIJ dataset (a bit more than two years) to train the model (in other words: we fitted the parameters of our logistic regression model to the first half of the data). The light blue line in figure 4.10B shows the seasonal cycle of low-level jets in those first two years of observations. Note that this seasonal cycle is very erratic. This can be expected for such a short period, but the question is whether the additional information contained in the predictor variables enables us to predict the other two years, despite the unrepresentative training data. Thus, in the next step, we used our trained model to predict the other half of the dataset. In fact, the model predicts the probability that a low-level jet occurs. An individual jet is predicted only if the probability is higher than 50 %, but this happens only occasionally. Therefore, rather than predicting individual jet events, we used the predicted probabilities directly and computed the monthly mean predicted probability (figure 4.10B, orange line). To evaluate the performance, we compared the predicted seasonal cycle with that based on the true observations during the second part of the dataset (figure 4.10B, light green line). The true seasonal cycle was indeed smoother than in the first two years, but it peaked a bit higher and earlier than predicted. To quantify this result, we computed the root mean square error between the monthly means of the predicted and test-data, and found it to be about 1 p.p.. This result confirms that the model generalizes well to new input data.

We then used the full MMIJ dataset to train the same model. With twice as much training data as before, we are confident that the model will achieve at least similar performance and thus predict the seasonal cycle to within 1 p.p. RMSE (but probably better). The observed seasonal cycle averaged over these four years of training data (figure 4.10B, red line) is still clearly affected by the unrepresentative months in the first half of the dataset. Apparently, four years of data is still not enough for the climatology to converge. Therefore, in the final step, we used the trained model to predict the 10-year seasonal cycle. The result (figure 4.10B, purple line) is a smooth seasonal cycle which peaks in May at about 9 %. This is our best estimate of the low-level jet seasonal cycle, based on the coalescence of reliable measurements and extensive reanalysis data. Compared to the results presented in section 4.6, we can conclude that we have adjusted the erratic nature of the short-term observations (figure 4.6E), resulting in a seasonal cycle similar to that shown in figure 4.6A, but with reduced amplitude. Compared to this final result, the crude amplitude adjustment with which we started in section 4.6 now appears far too strong.

The results presented in this section are intended as proof of principle, and for the purpose of illustration we tried to keep it conceptually simple. With respect to the selection of predictor variables, choice of model, and method of cross-validation, we realize that the possibilities are endless. The availability of sufficient measurement data is key to an exhaustive follow-up study.

## 4.9 Discussion

This paper has demonstrated our efforts to infer reliable low-level jet characteristics by combining observations and reanalysis data. We have deliberately chosen to illustrate how the results are impacted by limitations of the data and choices in the analysis. In this section we summarize our work, discuss the implications and offer an outlook to future research directions.

We started with a general validation of the ERA5 data for the observed wind speed at measurement locations at the North Sea. We found that the overall root mean square error is between 1.25 and 1.5 m s<sup>-1</sup>. The bias shows a clear discontinuity at 10 UTC, which is related to the data assimilation strategy that was used to produce ERA5. Users of the ERA5 data should consider a suitable bias correction (e.g. Staffell and Pfenninger, 2016), but we strongly suggest that future reanalysis products use sliding or at least partly overlapping observation windows. We also demonstrated that the observations alone can neither be relied upon, because the limited temporal extent of the measurement data leads to biased climatologies. Thus, in the remainder of the paper we focused on finding a suitable way to combine the two datasets. A procedure similar to measure-correlate-predict methods but tailored to low-level jets instead.

Low-level jet detection is very sensitive to the vertical extent of the data, and this has important implications for the interpretation of all results. Typical jet characteristics like jet height and jet strength cannot be reliably inferred from range-limited observations. With this restriction in mind, we can say that many of the observed jets occurred at heights fully or partly in the range spanned by contemporary wind turbine blades. Moreover, typical observed jet strength is about 8 m s<sup>-1</sup>, which is in the cubic part of the power curves of these turbines. We therefore expect that the low-level jet impact on loads and power can be substantial. ERA5 is not able to reliably reproduce these characteristics. There are some indications that the jets are ‘smeared out’: they appear higher and weaker than observed. Given this vertical displacement, a fair comparison between ERA5 and the observations is difficult. Considering the lower 300 m only, ERA5 drastically underestimates the amount of jets, but including heights up to 500 m, ERA5 shows more low-level jets than observed. We decided to include the data up to 500 m because it gives a stronger climatological signature.

Even though 1:1 correspondence between ERA5 and the observations is poor, both datasets agree on the following climatological characteristics: most jets occur in spring and summer; the diurnal cycle is weak, only around noon the chances for low-level jets are slightly lower; low-level jets are typically associated with stably stratified conditions; the absence of a strong large-scale forcing, or flow regimes with a pronounced easterly or offshore component are favourable for their formation. From the ERA5 data, we learned that low-level jets concentrate along the coasts. We then compared the frequency of low-level jets between ERA5 and the observations. In the most general terms, we can state that the mean low-level jet rates based on ERA5 up to 500 m typically overestimate the amount of low-level jets that would have been observed with lidars up to 300 m by a factor  $\sim 2$ . To improve upon this result we then illustrated how a logistic regression model was able to predict the seasonal cycle of low-level jets at MMIJ to within

1 p.p. RMSE. This is a promising result, and we expect that our results can still be improved upon. Longer measurement datasets would form a major contribution to further advancement as well.

The characteristics identified in this paper provide some clues as to the processes that govern these jets. The academic literature recognizes two dominant formation mechanisms, both of which are supported by our results. The first is frictional decoupling (Blackadar, 1957; Van de Wiel *et al.*, 2010). This theory describes a perturbed system attempting to re-establish equilibrium. As the accelerating wind field in the lower atmosphere is deflected by the Coriolis effect, it moves around its new equilibrium in a circular fashion. Over land, frictional decoupling has been linked to the decay of turbulent mixing around sunset and it has been suggested that a similar situation applies in coastal areas upon the abrupt surface (temperature and roughness) transition (Smedman *et al.*, 1993). This mechanism is supported by our results, which show that low-level jets are frequent for winds directed offshore and in stable conditions. The second mechanism relates low-level jets to horizontal temperature gradients (baroclinity, see Holton, 1967). According to this theory, the tilt of isobaric surfaces leads to a thermal wind component that under certain conditions can manifest as a low-level jet. This mechanism has been coupled to low-level jets over gently sloping terrain, but equally applies to coastal areas where large horizontal temperature differences can occur due to differential heating between the land and sea surface (Mahrt *et al.*, 2014). The fact that most low-level jets occur in spring and summer supports a baroclinic contribution, and possibly an interplay with the evolution of sea breezes, which show a similar seasonal cycle (e.g. Steele *et al.*, 2015). In the end, we expect that both processes are likely to contribute to the low-level jet climatology. Finally, we note that we also spotted a low-level jet with a clear frontal structure in the ERA5 data. It is unlikely that such events contribute significantly to the low-level jet climatology, but the characteristics of such jets may be very different and potentially much more harmful for (offshore) wind turbines. Other causes have been described in literature, such as orographic blocking. We don't expect this plays a major role along the Dutch coast, but for some of the low-level jets that are present in ERA5 along the British and especially the Norwegian coast it may play an important role (Christakos *et al.*, 2014). A more detailed investigation of the ERA5 data may allow us to separate these mechanisms. This is an interesting direction for further research.

With respect to future work, it would also be interesting to look at other datasets. In this paper, we have used ERA5 data to analyse the spatial characteristics of low-level jets directly. However, ERA5 is currently being used to develop higher resolution, down-scaled reanalysis datasets (e.g. the New European Wind Atlas (Petersen *et al.*, 2013) and the Dutch Offshore Wind Atlas), and it would be worthwhile to see if they improve upon ERA5. Another interesting alternative is COSMO-REA6 (Bollmeyer *et al.*, 2015), which is down-scaled from ERA-interim, but with its resolution of 6 km it might outperform ERA5. The current paper can serve as a guideline for the investigation of other reanalysis datasets.

Finally, a note on dealing with low-level jets in practice. It would be worthwhile to include a low-level jet case as standard inflow field for wake and load simulations. Recent papers have developed affordable methods to provide realistic inflow fields (Gebraad *et al.*, 2014; Englberger and Dörnbrack, 2018). Expensive CFD simulations have been used to derive parametrizations to generate realistic inflow fields for wind farm simulations. The second cited paper also includes low-level jet profiles in the early morning. These profiles can be compared with the morphology and frequency distributions detailed in the current manuscript to optimize yield and lifetime. Since the presence of the coastline turns out to have an important effect on the formation of

low-level jets, it would be interesting to perform an additional precursor LES simulation for such a heterogeneous terrain. This could also shed light on the mechanisms involved in jet formation.

## 4.A Lidar observations

Vertically pointing lidar provides efficient and non-intrusive measurement of boundary-layer winds. Compared to traditional meteorological masts, lidar typically expands the height and vertical sampling frequency of offshore wind measurements. Lidar data from seven measurement sites were used in this study to analyse North Sea LLJ spatio-temporal behaviour. Lidar type used included the WINDCUBE v2 pulsed lidar (only at LEG) and the ZephIR 300s continuous-wave (CW) lidar (all other platforms). The lidar were typically platform mounted, except within the Borssele wind farm and Hollandse Kust wind zones (Noord and Zuid) where the lidar was instrumented atop a floating metocean buoy. At these locations, two lidar-equipped metocean buoys were positioned simultaneously.

CW and pulsed wind lidar are coherent systems, meaning they both analyse Doppler shift frequencies to determine an estimate of the radial wind speed (Peña *et al.*, 2015). However, radial velocity and vertical wind profile extraction techniques differ between the two lidar types. Whereas pulsed wind lidar use range gates to near-simultaneously extract radial velocity estimates at multiple points in space, CW wind lidar can only extract a radial velocity estimate at the beam focus length. This beam focus length must be modified in time in order to measure the wind field at varying elevation levels. The radial wind speed is defined as the motion of the wind towards or away from the remote sensing system, and therefore unless the wind is moving along one of these radials, then the wind speed will not be fully resolved. Consequently, CW and pulsed wind lidar use varying adaptations of conical scanning techniques (Banakh *et al.*, 1995) to resolve the horizontal wind field at varying elevation levels. For brevity, these differences are not detailed here. However, because of these differences, the vertical wind profile was resolved at 17 s intervals for the CW wind lidar and at 4 s intervals for the pulsed wind lidar. These wind profiles are then analysed by the lidar software and output as a ten-minute average vertical wind profile. A summary of the lidar measurement heights and data collection periods for all sites is provided in figure 4.2.

Data quality control is imperative to ensure an accurate depiction of the offshore LLJ. Implementation of data quality control varied depending upon the lidar type (i.e. ZephIR 300s versus WINDCUBE v2), albeit considerations were made to ensure that data quality control was employed relatively uniformly between measurement sites. Wind lidar data from both the Borssele wind farm and Hollandse Kust (Noord and Zuid) wind zones have additionally had quality control measures implemented by Fugro Oceanor. An overview of these quality control procedures can be found online (<https://offshorewind.rvo.nl>). The data quality control procedures implemented are as follows. First, plausible value checks were implemented on the wind data. Any ten-minute observation that met the following criteria was removed from the data record:

1. The mean wind speed was either greater than the period maximum wind speed or less than the period minimum wind speed.
2. The mean wind speed was less than  $0.05 \text{ m s}^{-1}$ .
3. Turbulence intensity ( $TI$ ) for the period fell below 0.10 % (i.e. 0.001).

4. At the measurement height, the value of  $TI$  was 10 standard deviations ( $\sigma_{TI}$ ) greater than the mean ( $\mu_{TI}$ )  $TI$  value (i.e.  $TI \geq \mu_{TI} + 10\sigma_{TI}$ );  $\mu_{TI}$  and  $\sigma_{TI}$  were defined as the height-respective value for the entire data collection period. Because  $TI$  typically decreases with mean wind speed, this threshold was only imposed if the ten-minute mean wind speed exceeded  $4 \text{ m s}^{-1}$ .

Specific quality control measures were also applied to the lidar wind data. Any ten-minute observation that satisfied the following criteria were removed from the data record:


1. A lidar error code (e.g. 9998 or 9999) was reported.
2. The carrier-to-noise ratio (CNR) was less than  $-22$  (the value of CNR provides a measure of signal strength [i.e. quality]). CNR was only outputted by the WINDCUBE v2 wind lidar.
3. Backscatter magnitude was less than  $1 \times 10^{-5}$  or greater than 100 – backscatter served as a proxy for CNR for data reported by the ZephIR 300s lidar.

Prior analyses (e.g. Maureira Poveda and Wouters, 2015) demonstrate that the ZephIR 300s lidar can incorrectly measure wind direction by  $180^\circ$ . Analyses of wind data at MMIJ from 1 January 2012 through 1 January 2014 indicated that approximately 3.6 % of the measured wind data exhibited this flow reversal. Although mitigation (i.e. removal) of this data is possible, it requires independent wind direction measurements from a collocated meteorological mast. Because mast data was not available at each site, these wind direction errors were not removed. However, ZephIR 300s lidar wind direction errors did not appear to impact the measured wind speed, which is the main focus of this paper. In order to account for the wake effect of neighbouring wind farms on wind speed measurements, wind direction sectors were filtered and corresponding data (wind speed and direction) were removed. A generous estimate of 20 km was used to denote the maximum wind farm wake length.



# CHAPTER 5





## UNDERSTANDING THE PHYSICS OF COASTAL WINDS AND LOW-LEVEL JETS

---

WHAT IS HERE SAID, IS TO BE UNDERSTOOD OF THE SEA WINDS, AT SOME DISTANCE FROM THE LAND; FOR UPON AND NEAR THE SHORES, THE LAND AND SEA BRIZES ARE ALMOST EVERY WHERE SENSIBLE; AND THE GREAT VARIETY WHICH HAPPENS IN THEIR PERIODS, FORCE AND DIRECTION, FROM THE SITUATION OF THE MOUNTAINS, VALLIES AND WOODS, AND FROM THE VARIOUS TEXTURE OF THE SOIL, MORE OR LESS CAPABLE OF RETAINING AND REFLECTING HEAT, AND OF EXHALING OR CONDENSING VAPOURS IS SUCH, THAT IT WERE AN ENDLESS TASK, TO ENDEAVOUR TO ACCOUNT FOR THEM.

- EDMOND HALLEY FRS, 1686



Low-level jets have become the protagonists of this dissertation. They are perfect examples of anomalous wind events, their representation in weather models remains challenging, and they occur frequently enough to leave their mark on offshore wind projects. Moreover, as I'll show in this chapter, low-level jets are also intriguing from a physical point of view. An attempt to explain them with simple physical reasoning has grown out to be more of a tribute to mesoscale meteorology. The specific question I originally set out to address in this chapter was *which mechanisms are responsible for the formation of low-level jets in the coastal areas along the North Sea?* But in pursuing this objective it seems I've opened Pandora's box, or Aeolus' bag of winds. Consequently, this chapter now deals with a wide range of meteorological theories.

The chapter consists of two parts. The first part consists of a literature survey of meteorological processes in coastal areas and low-level jets. By combining theory that was until now quite scattered across the academic literature, a consistent theory on coastal jet mechanisms emerges. In the second part of the chapter, this theory is projected onto the ERA5 reanalysis data.

## 5.1 Introduction

Over three centuries ago, Edmond Halley was working on his general circulation theory that should explain the notorious equatorial trade winds and monsoons. In defining the scope of his paper, he decided to add the little disclaimer that is quoted at the beginning of this chapter. The casual clause that could equivalently be translated as “near-shore processes are way too complicated and are therefore not considered” bears witness to some great physical insights, such as the importance of coastal morphology, surface heat capacity, and surface conductivity for the manifestation of the sea breeze. Considering the narrow view that standard texts on offshore wind energy take on the coastal environment today, Halley’s statement “that it were an endless task, to endeavour to account for them” still appears to be valid. Indeed, while substantial progress has been made since Halley’s days, the complexity of coastal processes can still be daunting at times.


This chapter is dedicated to understanding the coastal low-level jets over the North Sea that have been identified and characterized in chapters 2 and 4. In order to account for these events in (wind-energy) applications, and to better appreciate and eventually improve the model performance, it is imperative to understand the mechanisms that lead to these jets. It turns out that a comprehensive explanation involves a breadth of meteorological processes that stretches far beyond the scope of low-level jets alone. In fact, it starts where Halley left us: with the sea breeze.

The first part of the chapter is dedicated to a survey of relevant literature. Section 5.2 will present an overview of the sea breeze and its relation to coastal jets. Against this background, I will highlight various mechanisms that may (also) be associated with low-level jets in section 5.3, and I will discuss how these theories can be projected upon, or reconciled with coastal flows and the theory of the sea breeze. In the second part of the chapter (section 5.4), ERA5 data will be used to investigate the formation mechanisms of low-level jets over the North Sea.

## 5.2 The kinship between sea breezes coastal jets

The sea breeze is extensively described in the academic literature, and a comprehensive overview of many important contributions may be found in Miller *et al.* (2003). The classical textbook illustration is sketched in Figure 5.1A and the corresponding idealized description is roughly as follows:

*On a brilliant day in spring or early summer, when the sea is still relatively cold but the sun is fiercely shining down on land and water alike, a temperature contrast develops between the adjacent surfaces due to their different capacity to absorb all that heat. While little still is happening in the offshore environment, the land surface passes its excess energy on to the overlying air, causing it to heat and expand, pushing up against the air still further aloft and to the sides. The situation so provoked is not particularly stable: cold marine air is heavy compared to the adjacent air over land, and the expansion of that heated air has caused a drop in surface pressure over land. Consequently, the cold marine air near the surface slides under the warmer and lighter air over land and the air aloft starts flowing out to sea. As this process continues for a while, the rotation of the*



**Figure 5.1:** A. Idealized illustration of the sea breeze circulation with warm air over land and cold air over sea. Wind vectors are plotted at the coastline, along with a vertical wind profile of the cross-shore wind component. B. Hypothetical equilibrium state corresponding to the traditional sea breeze situation as sketched in A.

*earth apparently<sup>1</sup> causes the wind to turn towards the right until the surface wind blows parallel to the coast with the land to its left, and vice versa higher up.*

A very relevant climatology of sea breezes along the coasts of the North Sea is available in Steele *et al.* (2015). Their seasonal cycle bears a striking resemblance to the seasonal cycle of low-level jets identified in chapters 2 and 4. Most sea breezes form in spring and (early) summer, because the temperature contrast between land and sea reaches the largest values in this period. Additionally, notice the jet-shaped character of the cross-shore wind component in figure 5.1A, which must vanish near the surface due to the effect of friction. The question thus arises how to distinguish between the two phenomena, and the simple answer is that I will not make a distinction at all. Both the sea breeze and the low-level jet are often treated in isolation, and for some reason the concurrence and shared origin of the two phenomena is often not explicitly discussed, leaving room for ambiguity.

The ultimate driver of the sea breeze is the thermal contrast between land and sea. The low-pressure area that develops over land is (sometimes) referred to as a thermal low, or heat low (Hoinka and Castro, 2003). Thermal lows have been associated with many low-level jet ‘hot-spots’ along coastal areas across the globe (Ranjha *et al.*, 2013). In arid coastal areas, they may be a semi-permanent feature during the summer season. Their diurnal evolution and association with the sea breeze and (coastal) low-level jets is illustrated very clearly in Rácz and Smith (1999). The sea breeze may be interpreted as the initial response to the temperature gradient, the shore-parallel low-level jet as its equilibrium state after adjustment. This adjustment will be illustrated more clearly in the remainder of this section. An important nuance is that the thermal contrast is not the *only* mechanism that may lead to coastal jets. Other mechanisms are discussed in section 5.3.

### 5.2.1 On the evolution of the sea breeze

While figure 5.1A captures the dominant forcing mechanism of the sea-breeze, it is far from complete and arguably incorrect. For example, the sea breeze is hardly a stationary circulation. It is well-known that after the onset of the sea breeze, the Coriolis effect will cause it to turn in a clockwise direction in the Northern hemisphere (historically, the wind was sometimes said to ‘turn with the sun’). Theoretically, rotation will continue until the sea breeze blows parallel to

<sup>1</sup>from an earthly observer’s perspective

the coast. At that point, the Coriolis force is balanced by the pressure gradient. This theoretical equilibrium state is called *geostrophic balance*.

A hypothetical balanced state consistent with figure 5.1A might look something like figure 5.1B: the cold marine air has moved towards the land, where it has slid under the continental air. The warmer air aloft has replaced some of the sinking marine air over sea. The cold air, which had initially started moving towards the land has been deflected 90° to the right and blows parallel to the coast, coming *towards* the reader in figure 5.1B. The warm air has also been deflected, but now blows parallel to the coast, *away* from the reader, since it initially started moving in the opposite direction. The wind now blows parallel to the isobars with the high pressure on its right. This figure is already more consistent with the theory on coastal low-level jets, and it is worthwhile to point out some features of (the adjustment towards) this hypothetical equilibrium state.

The rate with which the sea breeze turns is dictated by the Coriolis parameter, a function of latitude proportional to the earth's rotation. Together, the rate of the diurnal forcing and the rate of the Coriolis deflection determine the atmospheric response. Van Delden (1992) explains that a relatively abrupt forcing is followed by an oscillatory response, whereas a gradual forcing, such as that associated with the sea breeze, induces a relatively steady adjustment towards the new equilibrium. Rotunno (1983) adds that if the Coriolis frequency exceeds the frequency of the diurnal forcing, the resulting circulation will be confined to its neighbourhood, which is the case in midlatitudes. In the tropics, by contrast, the sea breeze may penetrate far inland (and offshore).

Rotunno derived an idealized expression for the horizontal extent/radius ( $L_h$ ) of the sea breeze in mid-latitudes, which reads:

$$L_h = \frac{NH}{\sqrt{f^2 - \omega^2}} \quad (5.1)$$


where  $N$  is the Brunt-Väisälä frequency (a measure for atmospheric density stratification),  $H$  is the depth over which the air is heated, and  $f$  and  $\omega$  are the Coriolis parameter and the frequency of the diurnal forcing, respectively. For a temperature difference of 6.5 K over a layer of 1000 m (i.e. typical tropospheric lapse rate) at 52° latitude,  $L_h = 164$  km (using a reference temperature of 300 K to compute  $N$ ).

Interestingly, as  $\omega$  vanishes, Rotunno's length scale reduces to the Rossby radius of deformation ( $Ro$ ), defined as:

$$Ro = \frac{NH}{f} = \frac{\sqrt{gH\Delta\theta/\theta_0}}{f} \quad (5.2)$$

where  $g$  is the acceleration due to gravity,  $\Delta\theta$  is the temperature difference between the two fluids and  $\theta_0$  is the background temperature. The *internal* and *external* Rossby radii, respectively, express the same concept for a continuously stratified fluid as opposed to a fluid with an abrupt density interface. The Rossby radius for the same fluid layer as described above is 127 km. Thus, the periodic nature of the diurnal forcing enhances the horizontal extent of the sea breeze.

The Rossby radius was first derived by Rossby (1937a,b), who described how the atmosphere (or ocean, or any other fluid under the influence of rotation) responds to a perturbation, in this case a horizontal temperature (or: density) gradient. This process is now generally known as geostrophic adjustment. While it seems to apply very well to the formation of sea breezes and



**Figure 5.2:** A. Zonal cross section at 52°N, illustrating the positioning of the North Sea with the Atlantic ocean to the west and the Eurasian continent to the east. The colours indicate a hypothetical large-scale temperature gradient with relatively warmer air over the continent than over the ocean. B. Enlarged view of the North Sea positioned between the Netherlands and the UK. Here, the colours suggest that the land on both sides has been heated by convection during the daytime, leading to an enhanced local baroclinity. Compare to figure 5.1B for an illustration of the corresponding local wind field.

their evolution as outlined above, the situation is not completely analogous<sup>2</sup>. Rossby assumed, for example, that the perturbation occurred instantaneously, whereas the forcing of the sea breeze is rather gradual. This difference is reflected in the equations above.

### 5.2.2 Baroclinity and thermal wind

Given that the forcing is relatively slow compared to the atmospheric response, a more realistic sketch of the situation is given in figure 5.2, which depicts a gradual horizontal temperature gradient across the coastlines. After several subsequent days of heating, a large-scale temperature gradient may also build up between the main continent of Europe and the Atlantic Ocean to the west (Miralles *et al.*, 2014). The situation in which the sea breeze is embedded in a large-scale temperature gradient is described in (Drobinski *et al.*, 2011). Notice how the sea breezes on the Dutch and UK coasts might amplify each other by sharing the localized high-pressure area over the North Sea, and how the sea breeze on the east coast of the UK opposes the large-scale temperature contrast, while the Dutch sea breeze reinforces it.

Following the AMS glossary, an atmosphere “in which the geostrophic wind varies with height and is related to the horizontal temperature gradient via the thermal wind equation” is called a

<sup>2</sup>As with the definition of the sea breeze, there is some controversy about stretching the term geostrophic adjustment beyond its original meaning; see e.g. Plougonven and Zhang (2014).

baroclinic atmosphere<sup>3</sup>. This thermal wind relation expresses the *difference* in the geostrophic wind over a layer ( $v_1$  at the top,  $v_0$  at the bottom):

$$v_T = v_1 - v_0 = \frac{R}{f} \ln \left( \frac{p_0}{p_1} \right) \frac{\Delta \theta}{\Delta x} \quad (5.3)$$

where  $R$  is the gas constant for dry air,  $p_1$  and  $p_0$  are the pressure at the top and the bottom of the layer of air,  $\theta$  is the potential temperature and  $x$  is the distance in the direction of the temperature gradient (I used  $v$  and  $x$  for analogy with previous and subsequent illustrations). Thus, the thermal wind ‘blows’ parallel to the isotherms with the warmer air to its right. If this situation applies to a low-level jet, some authors refer to it as a baroclinic mechanism (e.g. Mahrt *et al.*, 2014; Shapiro *et al.*, 2016).

It is possible to estimate the theoretical sea breeze and corresponding low-level jet characteristics that can be explained by thermal wind arguments by plugging in some numbers to equation (5.3). It is not immediately obvious which value to use for the temperature gradient, but a good approximation may be obtained by setting  $\Delta x$  equal to the Rossby radius given by equation (5.2). Interestingly, the thermal wind thus becomes independent of latitude. Taking the layer between 1013 and 900 hPa, which is approximately 1000 m deep, the same temperature difference as before (6.5 K against a background temperature of 300 K) yields a thermal wind of<sup>4</sup>  $\sim 15 \text{ m s}^{-1}$ . If this temperature difference is equally split between the near-surface flow and the flow aloft, the wind near the surface would be  $7.5 \text{ m s}^{-1}$  towards the south, and the wind speed would gradually decrease, reverse and reach  $7.5 \text{ m s}^{-1}$  towards the north aloft. In practice, this is not achieved because friction plays a role near the surface, the temperature gradient develops gradually, and the response might not be symmetric. For the sea breeze, the so-called ‘return current’ (Tijm *et al.*, 1999) is typically much weaker and deeper than the surface gravity current. However, this back-of-the-envelope estimate shows that adjustment of the sea breeze towards thermal wind balance is a good candidate to explain coastal low-level jets.


### 5.2.3 Interaction with background wind

The idealized picture of the sea breeze and its transition into an along-shore coastal jet is still highly idealized. Especially the presence of an ambient or ‘background’ wind may substantially alter the manifestation of the sea breeze. A moderate background wind can amplify or weaken the process depending on its predominant direction. The interaction between the sea breeze and the prevailing wind has been very neatly described in Adams (1997). He distinguishes between four sea breeze ‘types’, but even within these types he adds some nuances. To guide the interpretation, figure 5.3 shows the superposition of the sea breeze and 4 extremes of the background wind. These wind vectors are the *equilibrium* winds, and the sea breeze will first develop in an onshore direction before it turns towards the south by virtue of the Coriolis effect.

If the prevailing wind is directed offshore as in figure 5.3A, the onshore flow that marks the establishment of the sea breeze must first overcome the background wind. This favours the formation of a strong temperature contrast. The combined wind vector (as will be observed) thus first diminishes, then reverses, and finally turns towards the south. An onshore background wind as in figure 5.3B has the opposite effect. The synoptic onshore wind impedes

<sup>3</sup>Wallace and Hobbs (2006) refer to the specific configuration discussed here as an ‘equivalent barotropic flow’, because in this case, the pressure and temperature gradients are exactly aligned and therefore surfaces of constant density don’t intersect with surfaces of constant pressure, which is the formal definition of baroclinity.

<sup>4</sup>Alternatively, taking Rotunno’s length scale yields a wind shear of  $11\text{--}12 \text{ m s}^{-1}$ .



**Figure 5.3:** Illustration of the superposition of background geostrophic wind (green) with diurnally varying thermal forcing (red) for four different background wind configurations: easterly (A), westerly (B), northerly (C) and southerly (D). Land is on the right, sea on the left. Panel titles refer to the classification of Adams (1997), although he describes that the corkscrew and backdoor types usually also have an offshore component. In C and D, exaggerated surface wind vectors are drawn to illustrate the divergence and convergence associated with the synoptic surface pressure field.

the development of a strong thermal contrast, and if a sea breeze *does* develop, it is also much harder to recognize. Following the schematic picture in panel B, it might initially amplify the background wind, and the combined wind vector will eventually veer towards the south. Some authors (e.g. Miller *et al.*, 2003; Steele *et al.*, 2015) disregard this possibility altogether, stating that the sea breeze generally forms against a modest background wind directed offshore.

An along-shore background wind does not seem to interfere with the onset of the sea breeze at first sight. However, Adams (1997) notes that a synoptic high-pressure area over sea favours the formation of the sea breeze circulation, for air from aloft can easily sink into the surface divergence zone offshore. Therefore, the situation sketched in figure 5.3C leads to the strongest sea breezes, especially if the prevailing wind also has an offshore component. The development of this sea breeze is characterized by a backing of the observed wind in time, followed by wind veer as the new wind vector adjusts towards the equilibrium. Notice that aloft, the superposition of the prevailing wind and the sea breeze's return current is opposite, which leads to a very characteristic low-level jet-shaped wind profile with a pronounced fall-off, but not necessarily a reversal of the wind direction. This is also recognized by Steele *et al.* (2015), who note that coastal jets are typically associated with corkscrew sea breezes, although it does not become clear from the text how they distinguish between sea breezes with or without a coastal jet. Finally, in the situation sketched in figure 5.3D, the surface convergence offshore impedes the development of a circulation, and therefore this situation requires a very strong temperature contrast and a weak synoptic flow. If it forms, the evolution of the sea breeze is much more unsteady (Adams, 1997).

The inclusion of a background wind already provides a much more nuanced perspective on the sea breeze. Additionally, the characteristics of the coast, whether it is straight, or convex, or concave, the presence of mountains, cities, estuaries, or whatsoever – all these features leave their mark on the eventual manifestation of the flow (Miller *et al.*, 2003, section 3.1.3). Such effects, however, are beyond the scope of this chapter. The subsequent sections first and foremost require an appreciation of the role of thermal effects in coastal flows, which has now been covered sufficiently in-depth.

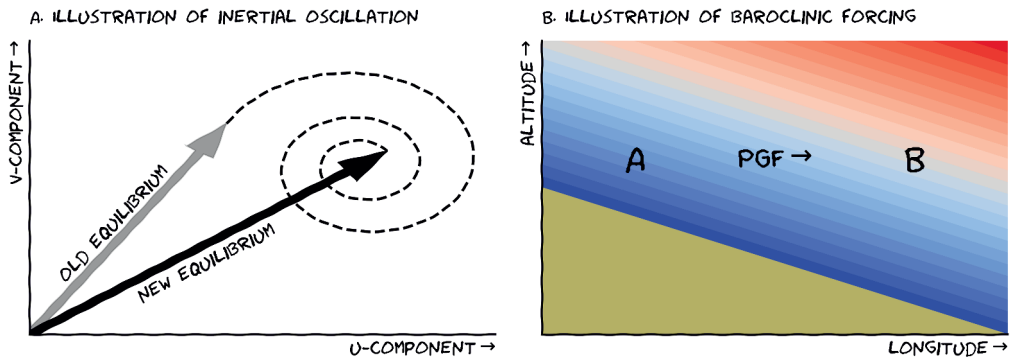
### 5.3 A survey of explicit low-level jet studies

While the thermal argumentation presented in the previous section is often used to explain coastal jets, it is not the only mechanism that can create jet-shaped profiles. Since the morphological definition employed in this thesis does not distinguish different low-level jet mechanisms, these processes cannot be readily discarded. Furthermore, several studies have found an interplay between various jet formation mechanisms, which calls for a brief survey of previous low-level jet studies.

#### 5.3.1 Blackadar's inertial oscillation

Much previous research on (onshore) low-level jets can be traced back to two main theories. The first is the inertial oscillation, proposed by Blackadar in 1957. This theory describes how the wind field responds to a change in environmental conditions. In principle, any change in environmental conditions might trigger inertial oscillations, but the characteristic (nocturnal) low-level jet over land is usually attributed to the decay of turbulence and the formation of a stable boundary layer after sunset. This relatively abrupt change in turbulent mixing upsets the three-way force balance between the pressure gradient, Coriolis effect, and friction. In the subsequent adjustment process the accelerating flow is modulated by the Coriolis effect to establish a rotation around the new equilibrium (figure 5.4A). Detailed explanations of this process may be found in Van de Wiel *et al.* (2010); Schröter *et al.* (2013).

Because the (change in) frictional drag affects mostly the surface layer, the inertial oscillation is confined to the lower layers of the atmosphere, and the corresponding wind speed maximum that occurs when the ageostrophic wind overshoots the equilibrium value is recognized as a low-level jet. The development of a stable boundary layer near the surface is important for an inertial oscillation to persist, because the negative buoyancy inhibits mechanical turbulence which could otherwise destroy the low-level jet. For this reason, the onshore nocturnal jet is often observed during cold nights with clear skies.



**Figure 5.4:** Conceptual illustration of the adaptation of a wind vector to a new equilibrium with less frictional drag. The adaptation is in the form of a damped inertial oscillation. Simplified momentum equations with linear damping were used to construct this figure, which is intended for illustrative purposes only. The exact trajectory is very sensitive to the damping coefficients. B. Conceptual illustration of baroclinity introduced by sloping terrain (exaggerated). The horizontal temperature difference introduces a pressure gradient force between point A and B. This pattern reverses when the air is heated from below.



At 52° latitude, the period of an inertial oscillation is  $\sim 15$  h. An interesting aspect of this mechanism is that the low-level jet only appears during a certain phase in the oscillation. If the low-level jets identified in previous chapters are in fact governed by inertial oscillations, this would imply that only half of the process was captured and characterized. Although this limitation might affect some of the low-level jet statistics presented in this and previous chapters, it should not impede our understanding, which is the main focus of the current disquisition.

Baas *et al.* (2009) provide a climatology of inertial low-level jets over Cabauw, a flat grassland site in the Netherlands approximately 50 km inland. They report very similar characteristics to those identified in chapter 2 for offshore site IJmuiden: typically occurs in spring/summer, mostly for easterly and southwesterly wind directions. The modal jet height is 140 m above ground-level and the typical strength is  $6\text{--}10\text{ m s}^{-1}$ . The theory has also been used to explain a low-level jet that is frequently found over the Great Plains and may reach wind speeds of over  $20\text{ m s}^{-1}$ . Baas suggest that this difference in low-level jet strength might be explained by the intensity of turbulent mixing during the day. Since the Great Plains are generally much dryer than the Netherlands, the sensible heat flux is larger, leading to more momentum entrainment and a larger ageostrophic wind – which translates into a larger oscillation after sunset (also see Fast and McCorcle, 1990).

A similar frictional decoupling may occur when an airmass is advected from land to sea (e.g. Garratt and Ryan, 1989). Especially in the situation where the land is heated during the daytime, but the sea is still cold in spring and early summer, the airmass will experience an abrupt change in surface roughness and stability. Cooling from below creates a stably stratified internal boundary-layer (Garratt, 1990). Högström and Smedman-Högström (1984); Smedman *et al.* (1993, e.g.) argue in favour of a frictional decoupling mechanisms to explain the occurrence of a low-level jet over the Baltic Sea. Angevine *et al.* (2006) find a strong low-level jet directed offshore at the northeast coast of the US, and speculate that it is formed by an inertial oscillation triggered by the collapse of the sea breeze circulation. Mahrt *et al.* (2014) considers both thermal and inertial arguments to explain this jet, but his study is inconclusive. Dörenkämper *et al.* (2015) illustrates a jet event in the Baltic Sea that forms in concert with an onshore inertial oscillation.

### 5.3.2 Holton's baroclinic mode

Just like differential heating between the land and sea, the diurnal heating cycle over a gentle slope introduces a temperature (or: density) gradient on a horizontal plane (figure 5.4B). Holton (1967) was the first to capture this effect in a conceptual model, with which he also attempted to explain the low-level jet over the Great Plains, for this area is characterized by a gentle, but widespread slope towards the east. Notice the analogy to the coastal situation: during the daytime, the baroclinic pressure gradient forces the air to flow uphill. However, as this forcing acts relatively slowly, the Coriolis effect deflects the uphill flow towards the north (in the illustration). By the time the new equilibrium is reached, however, the sun will have set, causing the wind to turn further, downhill. The oscillation is completed as the Coriolis effect finally deflects the flow further towards the south.

Interestingly, Shapiro *et al.* (2016, a.o.) argue that neither frictional decoupling nor baroclinity alone can explain all characteristics of the Great Plains low-level jet. They note: “the phase of the oscillations was not captured correctly, and the wind profiles were not as jet-like as in observed low-level jets.” They derive a set of equations that couples both mechanisms. Their

solutions reproduce most features of the observed low-level jets at least qualitatively. Moreover, a detailed sensitivity analysis revealed that inertial oscillations produce the strongest jets, but are reinforced by baroclinic contributions. A similar situation may very well apply to coastal low-level jets, for the conditions that favour baroclinity (cold sea surface, heated overland air) also favour the formation of internal boundary layers and inertial oscillations.

Shapiro *et al.* (2016) considered only cases with a large-scale southerly wind. Although this is the most frequent phenomenon, northerly low-level jets have also been observed. Gebauer *et al.* (2017) repeated the analysis with a reversed free-atmosphere wind. As expected, each individual mechanism behaved similarly, although the jet timing for the baroclinic mechanism was completely opposite. Consequently, the combined effect of the two mechanisms resulted in much weaker low-level jets than for the southerly case, and the jet occurred at much higher altitudes. This result emphasizes that the interplay between the local baroclinity and the large-scale wind is crucial for the eventual manifestation of the wind field, similar as shown for the coastal situation in figure 5.3.

As mentioned before, low-level jets are sometimes mentioned as a collateral feature in papers on sea breezes or thermal lows. A number of studies, however, focus explicitly on low-level jets and explain them in terms of Holton's baroclinic formation mechanism. Zemba and Friehe (1987) report on a thermally forced low-level jet along the California coast. Burk and Thompson (1996) extensively describe the characteristics of this jet, and Parish (2000) similarly concludes that this jet "is consistent with the geostrophic adjustment of thermally direct circulation forced by the horizontal temperature contrast between land and ocean". Colle and Novak (2010) describe a coastal low-level jet at the east coast of the United States, and find qualitative agreement with Burk and Thompson's findings. Mahrt *et al.* (2014) seem to suggest that both inertial effects and a baroclinic mode may contribute to the low-level jet along the northeast US coast.

### 5.3.3 Orographic effects

The presence of complex terrain has many effects on the flow that may or may not be recognized as low-level jets. Stull (2000), for example, describes a number of important mountain winds such as the Bora and Fohn, which may be interpreted as low-level jets, at least by the morphological definition adopted in this thesis. These winds form when air is forced over a mountain (range). Alternatively, the mountains may block the flow, causing accumulation of air and a pressure gradient force opposing the large-scale flow. The resulting wind that blows parallel to the mountains is called a barrier wind. Finally, gap winds can produce jet-like structures in places where there is an opening in the mountain barrier. The Venturi effect plays an important role in the formation of several of these mountain winds. Alternatively, thermal effects similar to those described before can play a role, such as in the formation of anabatic winds and drainage flows (Stull, 2000; Whiteman, 2000; Grisogono and Belušić, 2009; Sun and Sun, 2015; Duine *et al.*, 2019).

To appreciate the (potential) role of orographic effects in the present work, some findings of earlier studies are projected onto the low-level jets identified in the North Sea study area. As an entry into the relevant literature, I'll use the study of Moore and Renfrew (2005), who investigate the occurrence of four types of 'high-speed surface-wind events' along the coast of Greenland.

The first event is the barrier wind, which blows parallel to the east coast of Greenland with the coast to its right. The Froude number, or an alternative definition  $U/(Nh)$ , which is sometimes

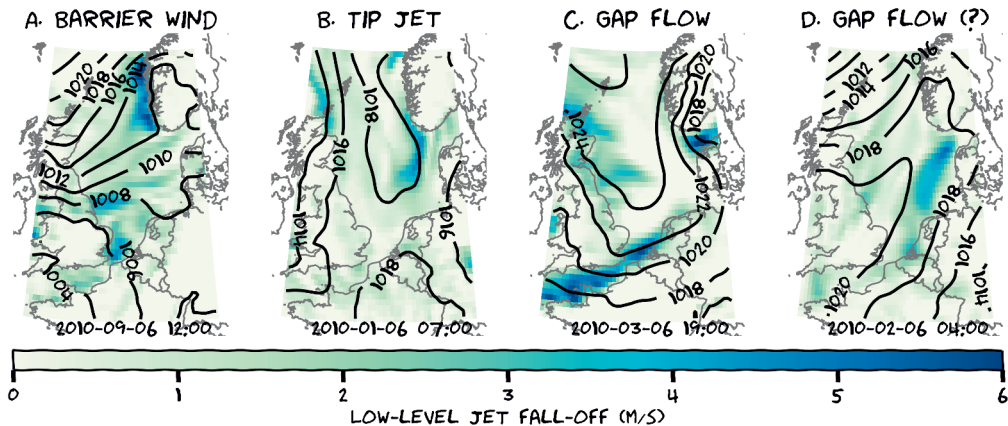


Figure 5.5: Snapshots of low-level jet events in the ERA5 dataset. Contours represent sea-level pressure in hPa, colours represent the fall-off where a low-level jet is present.

called the (inverse) non-dimensional mountain height, expresses the likeliness that wind will flow over the obstacle. In this ratio,  $U$  is the wind speed directed towards the obstacle,  $h$  is the obstacle height and  $N$  is the Brunt-Vaisala frequency, a measure of density stratification. If this ratio is substantially smaller than 1, the obstacle will partly block the flow (see Lott and Miller, 1997). According to Moore and Renfrew, a typical value for the southeast coast of Greenland is 0.3, thus, a blocking effect seems likely.

Within the North Sea study area, the most likely area for the occurrence of barrier jets is the west coast of Norway. For  $N = 0.02 \text{ s}^{-1}$  and taking  $H = 1000 \text{ m}$  (a rough estimate, in reality the coast is quite erratic), wind speeds of up to  $20 \text{ m s}^{-1}$  would give a Froude number below unity. An example of a low-level jet event in the ERA5 data that would match this description is shown in figure 5.5A. Notice how the isobars are deflected near the Norwegian coast. This localized area of higher pressure is found where the flow would be blocked by the mountains. The geostrophic wind, following the isobars, would thus blow towards the north here, whereas the predominant geostrophic wind direction in the domain is towards the east. This situation is described in more detail in Barstad and Grønås (2005) and a case study concerning a similar event may be found in Christakos *et al.* (2014).

The second event considered by Moore and Renfrew is the ‘tip jet’, after Doyle and Shapiro (1999). The following arguments are put forward to explain this type of jet<sup>5</sup>: (1) According to Bernoulli’s equation the total energy of the flow along a streamline must be conserved, which must lead to an acceleration of the flow when it descends while maintaining its temperature. (2) However, over the higher parts of Greenland (or any other complex terrain), the orography may cause breaking waves, which lead to a turbulent and dissipative region on the lee side of the mountains. These flows thus don’t conserve Bernoulli’s equation. (3) If part of the flow is forced over the high regions of Greenland, while an adjacent part of the flow is deflected slightly more to the south, where it will encounter gentler slopes at a less direct angle, the latter will conserve the Bernoulli equation whilst the former loses energy in the wave breaking region. (4)

<sup>5</sup>While Moore and Renfrew present these arguments as if they concern different mechanisms, my impression is that Doyle and Shapiro see them as part of the same mechanism.

This would lead to a sharp horizontal wind speed gradient downstream. (5) Hence the relatively high localized wind speed near the surface in the downstream region of the tip of Greenland. Moore and Renfrew also consider what they call *reverse tip jets* (their third event), which occur for different wind directions, for which Doyle and Shapiro's theory does not hold. To explain this feature, they propose that the wind downstream of the tip gains anticyclonic curvature. In this situation, the Coriolis effect not only balances the pressure gradient force, but also the centrifugal acceleration, enabling supergeostrophic wind speeds.


In the ERA5 data, a tip jet event often appears near the Southern tip of Norway. An example is shown in figure 5.5B. Anticyclonic curvature is present here as well, but otherwise the flow is quite different. This is not surprising, though, since the surrounding topography (e.g. the presence of Denmark) is quite different than Greenland. Therefore, the theory of Moore and Renfrew cannot be translated one-to-one. It is worthwhile to note, though, that preliminary results from higher-resolution simulations (chapter 6) show pronounced jet activity originating from many smaller coastal extremities, such as the Isles of Zeeland and the easternmost edges of East Angles and The Southeast. This provides an incentive to further investigate the mechanics of tip jets.

The fourth and final event studied by Moore and Renfrew is a 'barrier wind' for which the opposite coast of Iceland seems to play an important role. Therefore the authors refer to a comparable study about wind speed-up in the Dover Strait (Capon, 2003), which is quite enlightening. Capon illustrates that both the orographic 'gap' and the roughness contrast between land and sea contribute to the jet formation in the Dover Strait, even though dimensional arguments suggest that a true gap flow (forced channelling through stable stratification in conjunction with the venturi effect) is unlikely. She also discusses potential contributions of the sea breeze and drainage flows. In the ERA5 data, the Dover Strait jet is frequently present. An example is shown in figure 5.5C. In this case, the entire English Channel is filled with the jet, which is forced by a northeasterly wind.

As described by Capon, a similar wind may form when the wind is from the southwest. If this jet reaches the IJmuiden met mast, it might be an important factor to explain the southwesterly jets identified in chapter 2. Further support for this jet is contained in preliminary results from chapter 6, where a pronounced jet area, apparently originating from the Dover Strait, shows up in the *averaged* low-level jet climatology of a high-resolution dataset. Also interesting is a comparable event in the ERA5 data that occurs near Skagerrak, the Strait between Denmark and Norway (figure 5.5D). Low-level jets are often present near the west coast of Denmark and the German Bight, though, and most of the time the 'true nature' of this jet is not so obvious as it seems in this single snapshot. These preliminary results thus suggest that the role of orography on North Sea low-level jets may be substantial, and call for further exploration of these effects.

## 5.4 Exploratory analysis of low-level jets in the ERA5 re-analysis data

In the preceding sections, I have surveyed a number of mechanisms that could be responsible for, or play a role in, the formation of low-level jets along the coastlines of the North Sea. The second part of this chapter is devoted to a more thorough analysis of low-level jets in the ERA5 data: a quest for evidence that can support or reject some of these hypothesis. I'll start with



**Figure 5.6:** A. Illustration of the situation in which a thermal low-pressure area develops over the Netherlands. B. Corresponding vertical profile of the (geostrophic)  $v$ -wind component. C. Histogram showing the ratio between low-level jet fall-off and thermal wind. The black lines delimit the range where the falloff and the thermal wind are within 70 % of each other. This figure contains all low-level jets in the ERA5 data within the study area.

an assessment of the potential role of baroclinity since, based on the literature, this seems to be the most important effect.

The situation for the Dutch coast is sketched in figure 5.6A. During the day, convective heating and the subsequent thermal expansion lead to a low-pressure area over land as compared to the adjacent sea-surface. In accordance with theory, the near surface wind speed tends to blow in an anticlockwise direction around the low-pressure area, that is, towards the south at the Dutch coast (a negative  $v$ -wind component). At the same time, the thermal wind relation dictates that the balanced geostrophic wind should change with height, specifically, the  $v$ -wind component should become more positive. The corresponding vertical profile is sketched in figure 5.6B. This profile corresponds exactly with the theoretical equilibrium situation sketched in figure 5.1B. Because the wind speed must vanish at the surface due to the effects of friction, the true wind profile is not equal to the geostrophic wind profile. Thermal wind is responsible for the fall-off above the jet maximum, and friction plays an analogous role below.

If this mechanism is indeed responsible for the formation of low-level jets, the thermal wind contribution should explain (at least a substantial portion of) the fall-off. Therefore, the ratio of thermal wind and fall-off is shown in figure 5.6C. The thermal wind is a vector, and only the component aligned with the surface geostrophic wind is shown, for this component is responsible for the fall-off. A negative ratio means that the thermal wind opposes the surface geostrophic wind. It is clear that this is often the case. For most low-level jets, the ratio is even close to one. This is clear evidence that the baroclinic mechanism plays an important role in the formation of low-level jets over the North Sea. The rest of this section is dedicated to attribution of the low-level jets to this or other mechanism(s).

#### 5.4.1 Intro – An animation is worth a thousand pictures

The starting point of this analysis is the spatial distribution of low-level jets in figure 4.1. This *mean* distribution can be thought of as overlaying all individual low-level jet events. The ERA5 data covers the North Sea domain, including the east coast of the UK, the west coast of Northern France, Belgium, and the Netherlands up to the German Bight and the North Sea coast of Denmark, and finally the southern tip of Norway. As discussed in the previous section (e.g. fig-

ure 5.5), all these areas have their own distinct features that may favour or modify the formation and characteristics of low-level jets.

A closer look at these low-level jets in isolation, for example in the form of a short animation, reveals a number of additional features. First of all, the low-level jets are very dynamic. They first appear at one or several locations, then they grow and become stronger; sometimes they drift through the domain as though they have become decoupled from their geographical origins; every now and then, several disparate jets merge, or a big event breaks down into smaller sections. Some events have a clear spatial structure or orientation, e.g. following a coastline or resembling a frontal structure.

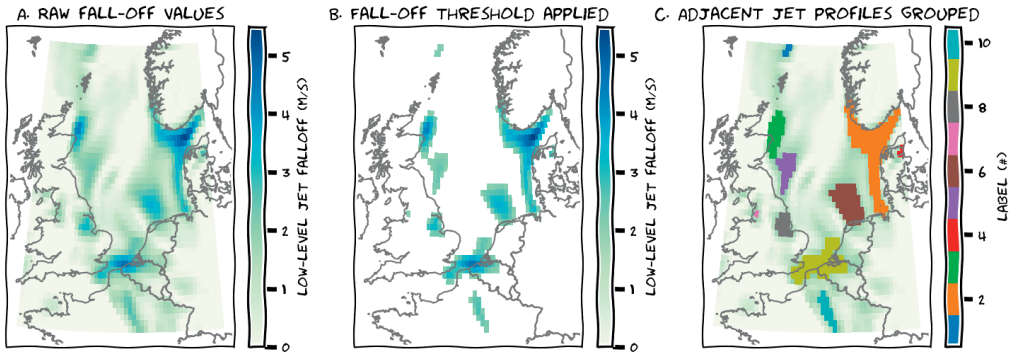
### 5.4.2 From jet profile to jet event

These observations about the low-level jet animation listed above all have one thing in common: they concern jet *events*, consisting of (many) adjacent grid cells exhibiting a jet-like structure in the vertical, rather than individual jet-shaped profiles. Usually, the fall-off decreases towards the edges of the jet event. This implies that using fixed falloff thresholds to classify some jets as *significant events* while discarding others is quite arbitrary, as both profiles can belong to the same event. The same logic implies that wind profiles that do not *currently* exhibit a jet-shaped structure may still be part of the same dynamic event taking place, but there is no obvious way to identify and track these points.

From a theoretical point of view these observations are hardly surprising, but they have some profound implications in terms of data processing. In order to characterize, and perhaps even categorize the low-level jets in the ERA5 data, it makes much more sense to consider the ‘whole’ rather than the individual jet profiles. The challenge is then, how to group coherent profiles together? The strategy used here is illustrated in figure 5.7. First, the fall-off of each and every jet-shaped profile is computed for the entire dataset. Subsequently, a fall-off threshold is applied to separate clusters that are only weakly connected. Stricter fall-off thresholds enable better separation of distinct events, but the downside is that weaker jets, possibly including ‘growth phase jets’, are disregarded altogether.

Another challenge presents itself when time-adjacency is considered as well. From a physical perspective, a jet event that evolves from one moment to the next is still the same event, even if it grows or migrates. It is not obvious how to deal with jet events that merge and split, though. One approach would be to change the ‘ID’ of the jet cluster(s) after each merge or split, thus counting them as separate event, possibly with a label (merged cluster; remnant). In practice, however, splits and merges occur so often that such book-keeping introduces more complexity than it solves. Therefore, only spatial linkage is considered here, as illustrated in figure 5.7.

A unique label is assigned to each jet event and subsequently, a number of jet characteristics is calculated on a group-by-label basis: (averaged) jet height, strength, and fall-off, as well as u- and v-wind components at jet height; (average) latitude, longitude, and distance to coast. Additionally, a number of diagnostic variables are computed: (bulk) Richardson number near the surface (computed with 2 m temperature, skin temperature and 10 m wind), the large-scale smoothed gradient magnitude of 2 m temperature field, and a vertically averaged temperature lapse rate (up to 2 km). Some variable are only time-dependent: day of year, hour of the day, time since sunset/sunrise, and the three Lamb Weather Type base components W, S and Z (for the large-scale westerly, southerly and cyclonic flow components, respectively, see Jones *et al.*,




**Figure 5.7:** Illustration of the jet event identification procedure. Panel A shows a snapshot of raw fall-off values at a single time. In B, a fall-off threshold is applied to obtain separated jet clusters. In C, each isolated cluster is given a unique number/label.

2013). Finally, the number of profiles that make up each jet event and the standard deviation of latitude and longitude are stored to convey some information about the spatial structure.

While I had good theoretical reasons to include all of these variables, not all of them proved helpful in characterizing the jet events. Time since sunrise/sunset, for example, correlates exactly with time of the year and day, and is therefore redundant from a characterization perspective. Standard deviation of latitude and longitude, on the other hand, (which was added in the hope to identify e.g. front-like jet events), are descriptive of only very few jets, while they add a lot of noise to most jet events. The strongest or most robust signal is obtained with a reduced set of parameters consisting of distance to coast, time of year and day, jet height, falloff and u- and v-wind components, as well as the surface (bulk) Richardson number. These variables for a subset of the jet events are shown in a comprehensive way in figure 5.8.


### 5.4.3 Clustering similar jet events together

Figure 5.8 is a ‘clustermap’ of 500 randomly sampled low-level jet events from the ERA5 data. While it contains relatively little data (the total dataset, based on a fall-off threshold of  $3 \text{ m s}^{-1}$ , contains approximately 200 000 jet events), it already illustrates some very interesting features. A hierarchical (Ward) clustering (illustrated with dendrograms) has been applied to both rows and columns, in order to reveal hidden structure in the data. In the column with the Bulk Richardson number ( $Ri_b$ ), for example, there is one distinct block with very high values (stable conditions). The column ‘d2c’ shows the distance to coast, where the sign designates land (negative) or sea (positive), showing that the distinctly stable jet events typically occur over land. Similarly, there are blocks with strong westerly jets (high u-component) over sea, and the strongest falloff (right top) is also observed offshore. Larger jet heights are typically found over land. The time of the day and time of the year have each been split in two components to better represent their cyclic nature. The hour of the day ( $H$ ) is represented by  $\sin(H)$  ( $= 1$  corresponds 6 UTC,  $= -1$  corresponds to 18 UTC) and  $\cos(H)$  ( $= 1$  corresponds to midnight,  $= -1$  corresponds to noon), and similarly for the time of the year (designated with  $M$  for Month for later use, but day of year is used in the computation). Hence, it appears that the stable, onshore low-level jet events occur predominantly during nighttime, while offshore jets occur more often during the day.



**Figure 5.8:** A ‘clustermap’ illustrating a random sample of 500 jet events out of the ERA5 data. Each row corresponds to a single event (some are numbered on the right), columns represent the variables used to characterize the jets (see text). The variables have been standardized (subtract mean, divide by standard deviation) in order to make them comparable. Hierarchical structure is illustrated by dendrograms. The first column with coloured clusters is obtained by cutting through the dendrogram at a fixed height (indicated by the blue line) and assigning all the branches below the cut to the same cluster.





**Figure 5.9:** Illustration of low-level jet characteristics as ‘revealed’ by a k-means cluster algorithm<sup>6</sup>. The specific cluster that is the focus of this figure is coloured in red, and other clusters have been plotted in the background in much lighter colours for reference. LWT-S, -W, and -Z denote Lamb Weather type components, BARO is for baroclinity, measured as a smoothed temperature gradient. MSL is mean sea level pressure, DSUNRISE and -SET is the time since sunset/sunrise. Lapse-T is the average temperature lapse-rate up to 2000 m and  $Ri_b$  is the bulk Richardson number.

While the figure above already provides some valuable insights, it is not very robust. In fact, I needed a few attempts to save a version of this figure that lent itself to the above illustration, and there are several reasons for this. First of all, the subset of 500 jet events is too small to adequately represent the full dataset, which consists of over 200 000 events (based on a falloff threshold of  $3 \text{ m s}^{-1}$ ). Secondly, there is such a wide variety of jet events that each combination of variables can exist, albeit that some are more likely than others. This makes it very difficult to identify distinct clusters. A third reason is that some variables included in the dataset - if they are not useful to describe a certain jet event - appear as noise in the multidimensional representation of the data. And finally, a fourth reason for the unstable behaviour is that the cluster method (Ward) used, although it produces nice clusters of more or less uniform size (useful for illustration), may not be the most suitable approach for this dataset; at least it doesn’t play well with the small random subsets. I’ve experimented with several alternative cluster algorithms, but I still haven’t found a method that suffers less from the aforementioned impediments.

<sup>6</sup>As opposed to the Ward method used for the clustermap in figure 5.8. Usually, these methods provide comparable results - but since the focus here is more an illustration of the general approach rather than the absolute results, any method would suffice.

There are several other ways to visualize the results of a cluster procedure. Perhaps the most straightforward method is to visualize the characteristics of each cluster separately (figure 5.9). The red cluster seems to consist of low-level jets that occur along the coastlines, preferably in winter and during the night. Furthermore, these jets are characterized by a strong baroclinity and the jets are predominantly directed towards the east. All other characteristics that are shown do not exhibit a marked affinity with this ‘type’ of jet. In the background, an ‘orange cluster’ can be recognized that occurs primarily near the coast during summer afternoons (mostly northerly winds; towards the south). Another ‘green’ cluster corresponds to jets that occur onshore during stably stratified conditions at night. Another onshore cluster is coloured in brown, and purple colours are mostly westerly jets originating from the UK east coast.

Again, some meaningful patterns seem to emerge, but this is not the whole story. First of all, many subjective choices are involved in obtaining this result, such as the desired number of clusters and which variables are supplied to the clustering algorithm. Too many clusters results in very similar clusters for most variables, too few results in a lot of noise in most variables. If location is the single input, then the resulting clusters will most likely show variations in location, but only if the location is highly correlated with another characteristic, jet direction for example, then this will show up in the output. In that sense, this exercise is a laborious form of multiple regression between jet characteristics.

Supplying many input variables, on the other hand, will give the clustering algorithm much freedom to exploit existing correlations. Unfortunately, the North Sea low-level jet climatology consists of a large number of jets with very different characteristics. Consequently, there appears to be a general lack of structure in the multidimensional space spanned by the characteristics listed above. The result is that the algorithm just divides that space in a number of subspace of approximately equal size. For example, if only location is supplied as input and only two clusters are sought, the resulting jets are geographically separated in a ‘southwest’ and a ‘northeast’ group. The same happens in (much) higher dimensions.

After experimenting with different numbers of clusters, various combinations of input variables, and several sample sizes (using cross-validation to check the results), a success formula that provides robust and meaningful results is still lacking. The biggest challenge might be to prove that the results are not ‘sought’, in the sense that they follow logically from subjective choices. Location, distance to coast, jet direction and time of year and day seem to be the most natural features to distinguish between jet mechanisms from a physical perspective, but if these features yield coherent clusters in terms of the very same features, is the result then physically meaningful or is this just a case of confirmation bias? If the answer to the last question is *yes*, it means the clustering approach has reached a dead end, and it is time to take a radically different approach.

#### 5.4.4 Back to basics: the hypothetico-deductive approach

Long-term progressive insight in scientific practice has inspired the so-called hypothetico-deductive method. The idea of this method is that one should not seek to confirm hypotheses, but rather to falsify them. I applied this method to the ERA5 low-level jet data to see whether the hypotheses of the previous section stand up to scrutiny. For the inertial oscillation, I deduced the following criteria:

1. If the low-level jet is indeed (caused/forced by) an inertial oscillation, then the Richardson number near the surface must be larger than 0, consistent with a stable stratification.
2. If the low-level jet is indeed (caused/forced by) an inertial oscillation, then the ageostrophic wind component should explain at least 70 % of the fall-off. Here, 70 % is a quite arbitrary threshold, which illustrates that also this method is not fully objective.
3. If the low-level jet is indeed (caused/forced by) an inertial oscillation, then the ageostrophic wind should rotate at a rate of (approximately)  $24^{\circ}\text{h}^{-1}$ , consistent with an inertial oscillation at  $52^{\circ}\text{N}$ . I used 15 and  $30^{\circ}$  as bounds.

The ERA5 data were confronted with these criteria, and the results are shown in figure 5.10A. The yellow outer circle encompasses all low-level jet events in the ERA5 data, the Venn diagram inside it shows the proportion of jets that meet one or more of the predictions. It is clearly visible that the strongest constraint is criterion 3, i.e. the ageostrophic wind should rotate clockwise. Still, 22 % of the low-level jets meet this criterion. About 50 % of the jets have an ageostrophic wind component that is comparable to the fall-off (criterion 2), and over 60 % meets the stability criterion (1). 15.4 % of the events does not meet any of the conditions. For these jets, the inertial oscillation hypothesis should clearly be rejected.


For the other events, a more nuanced interpretation is opportune. Figure 5.10B shows a slightly different rendering of the same data. The inner circle shows the jets that meet all criteria. For these events, the inertial oscillation hypothesis can clearly not be rejected. In other words, it is *very likely*. The other bands shows the events where two conditions are met (*likely*), where only one of the criteria is met (*slightly likely*), and where none of the criteria are met (*unlikely*). Of course, this again is a subjective interpretation, but I would argue that in this case it is within the leeway of expert judgement. Note that I did not alter the criteria based on preliminary results in an attempt to minimize the likelihood of rejection or whatsoever.

I performed a similar exercise for the thermal/baroclinic hypothesis, for which I employed the following criteria:


1. If the low-level jet is indeed caused by thermal forcing, then the thermal wind should have a substantial component that opposes the surface geostrophic wind. I defined ‘substantial’ as at least 0.3 times as strong as the surface geostrophic wind.
2. If the low-level jet is indeed caused by thermal forcing, then the distance to the coast should be within 100 km.
3. If the low-level jet is indeed caused by thermal forcing, then the thermal wind should explain at least 70 % of the fall-off. In fact, this is a slightly different form of prediction 1. It is stricter in the sense that the magnitude of the thermal wind should be comparable to the fall-off, but more permitting in the sense that the fall-off is not necessarily 0.3 times as strong as the jet maximum.

The results of testing these criteria against the ERA5 data is shown in figure 5.10C. 25 % of the jet events meet all criteria, which suggests that the thermal hypothesis has more potential to explain the jets than the inertial oscillation – although one could argue that the criteria for the inertial oscillation may have been stricter. However, support for one mechanism does not necessarily reject the other possibility. To investigate whether the two mechanisms act in concert, figure 5.10D shows a Venn diagram of the jets for which either or both of the hypotheses are very likely. Surprisingly, there is only little overlap between the two sets. Only 2.2 % of the jet


A. CRITERIA FOR AN INERTIAL OSCILLATION




B. LIKELINESS OF AN INERTIAL OSCILLATION



C. CRITERIA FOR A (COASTAL) BAROCLINIC JET



D. COINCIDENCE OF DOMINANT MECHANISMS



**Figure 5.10:** A. Venn Diagram illustrating the proportions of the ERA5 low-level jets that meet the criteria for an inertial oscillation. The outer circle denotes the total number of jet events in the data. B. A different rendering of the same data, suggesting that the inertial oscillation is more unlikely as fewer criteria are met. C. Similar to A but for the thermal/baroclinic hypothesis. D. Venn diagram illustrating the partial overlap between jet events that meet all criteria for both the inertial oscillation and the baroclinic mechanism. All spatial proportions in this figure are approximate; percentages are exact.

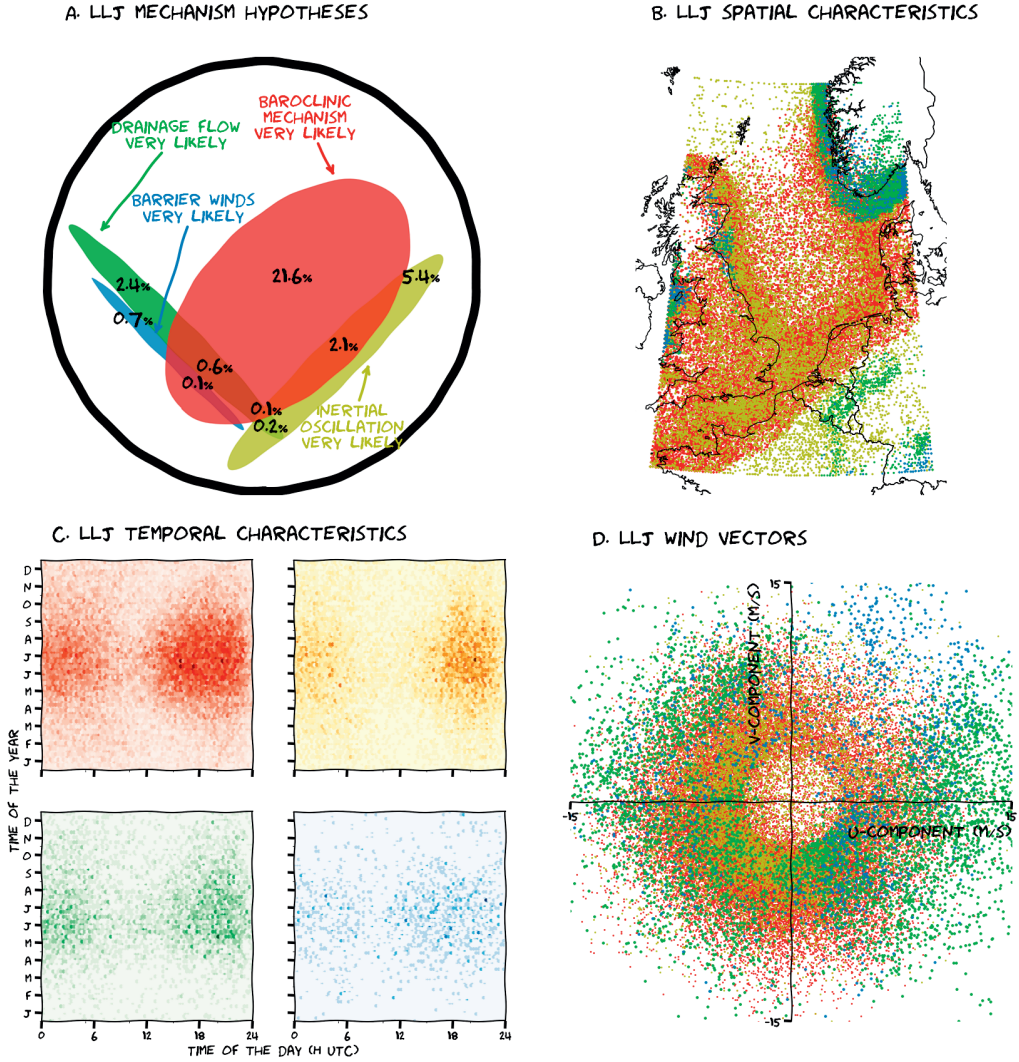
events simultaneously fulfil the criteria for an inertial oscillation and a baroclinic forcing. This figure may change if I also include the ‘likely’ events, but that again tends towards confirmation bias, so I’ve chosen to leave it here as is.

I *did*, however, test two other hypothesis with ‘quick and dirty’ predictions: if the Richardson number is larger than 0 and the jet had a substantial downslope component, a drainage flow (relatively cold air flowing downhill like a river) is ‘likely’. A barrier wind is ‘likely’ if the Brunt-Vaisala frequency is larger than  $0.01 \text{ s}^{-1}$  and the flow had a substantial component directed towards the slope (if any). Of course, more sophisticated criteria would result in a more reliable climatology, but orographic forcing is rather complex and a detailed treatment is beyond the scope of this study, which serves as a quick illustration rather than conclusive evidence. The joint distribution of all four low-level jet mechanisms are shown in figure 5.11A. Approximately 3.3 % of the jet events fulfil the criteria for the barrier wind, and 0.8 % for the drainage flow.

It is interesting to investigate the events for which all criteria were met, or rather, those for which a certain hypothesis could not be rejected. The spatial distribution of these events is shown in figure 5.11B. The thermal mechanism is clearly confined to the coastal regions, because that was one of the criteria. However, it seems that some of these jets are further offshore than the 100 km constraint. This is an artefact of the visualization, which is based on the mean latitude and longitude of the jet events whereas the distance to coast is based on the mean distance to coast. For low-level jets with a large spatial extent, this can introduce the apparent inconsistencies. The inertial oscillation is found over land, but also predominantly over sea. Drainage flows occur over land near the Ardennes and near the coast of Norway. In accordance with section 5.3.3, barrier winds also occur near the Norwegian coast. A few similar events are identified near the west coast of Scotland and in the south-west of Germany.

Temporal characteristics are shown in figure 5.11C. All events except for the barrier wind have two pronounced maxima: one during summer afternoons, the other in the summer early morning. The inertial oscillation seems to set in later in the evening than the thermal jets. This is in accordance with theory, which predicts that baroclinic jets appear in the course of the afternoon, whereas the inertial oscillation (at least over land) would set in later in the evening or in the night. A clear explanation for the jets in the morning is lacking, though. It could be that the jets are sustained throughout the night. Alternatively, the baroclinic mechanism may have reached its other extreme, leading to a renewed impulse for the manifestation of jet-shaped wind profiles. The theory for barrier winds does not imply a pronounced diurnal cycle, consistent with the results. There slightly higher jet frequency in mid-afternoon could point to a baroclinic contribution, but the diurnal cycle is indeed much less pronounced than for the other events. A seasonal cycle seems to be present, though. This is consistent with the seasonal cycle of sea surface temperature, leading to strong stratification and a smaller Froude number in spring and summer.

The jet direction is shown in figure 5.11D. The orographic events appear to be stronger than the other jets. The baroclinic jets are most frequently directed towards the south-west, in accordance with theory for the Dutch coast, and also towards the northwest, in line with the theory for the east coast of the UK (note that this was not explicitly required in the criteria). The inertial oscillations also have a predominant easterly component, which makes sense if they are formed on top of an internal boundary layer formed by offshore flow on the west coasts of the continent. It is also in analogy with Baas’ climatology over Cabauw (Baas *et al.*, 2009).



**Figure 5.11:** A. Similar to Figure figure 5.10D, but with inclusion of two orographic effects: barrier wind and drainage flow. B-D. Spatial distribution (B), temporal characteristics (C) and individual wind components (D) of the low-level jets attributed to the mechanisms in A.

For the east coast of the UK, however, a westerly component would be expected, which is less frequently observed.

It is possible that some ‘cross-contamination’ is contained in figure 5.11. Jet events for which the inertial oscillation hypothesis was rejected for its strict ageostrophic wind speed criterion might now be classified as baroclinic jets, even if there is a substantial contribution from inertial effects. These jets do not appear as overlap in figure 5.10D and 5.11A, because one of the criteria was not met. The similar temporal characteristics in figure 5.11C, for example, strongly suggest that there is a certain degree of interplay between the various mechanisms.

A unique aspect of the results presented in this section is that they concern the complete *climatology* of low-level jets in the study area. Previous studies focused on much more isolated events, considering only one or a few suitable cases. The relatively small study area with multiple, partly opposite coastlines of varying morphological characteristics further adds to the complexity of the observed phenomena. Consequently, it is not surprising that contributions of individual mechanisms can hardly be disentangled. Nevertheless, the results provide tenable evidence for a dominant role of thermal forcing, assisted, modified, and/or augmented by inertial and orographic effects.

These findings have important implications for wind energy applications. Many wind farm studies rely on simulations of isolated wind farms, and incorporation of a temperature equation is still more an exception than the rule. Although some of the effects revealed here may implicitly find their way into statistical representation of the wind climate, it is unlikely that a simulation without appropriate account for the complex embedding and thermal effects will provide reliable results. As such, this study provides a further incentive for the wind energy industry to move towards dynamic downscaling. This implies that coupling between numerical weather models (or reanalysis datasets) and large-eddy-simulations should be further developed and validated. These activities are already taking place (Muñoz-Esparza *et al.*, 2014, and e.g. in the DOWA project, see [dutchoffshorewindatlas.nl](http://dutchoffshorewindatlas.nl)), and hopefully this chapter leads to an increased awareness of the relevance of these developments.

## 5.5 Summary

Previous studies (chapters 2 and 4) have identified a substantial number of low-level jet events over the North Sea. This chapter aimed at understanding these events. It started with a survey of relevant processes described in the literature. This exercise revealed that coastal jets are often tightly linked to the thermal contrast between land and sea; the same mechanism that is responsible for the formation of the sea breeze. However, there are other mechanisms that may lead to jet-shaped wind profiles as well, such as the inertial oscillation and orographic effects.

To test whether these mechanisms could explain the low-level jets in the ERA5 data, I’ve shown two attempts to process or categorize these data. The first approach explored the use of cluster methods to group low-level jets with similar characteristics together. This was partially successful, although it proved challenging to obtain robust results and rule out the pitfall of confirmation bias. The second attempt was derived from the hypothetico-deductive model: subjective criteria about each of the aforementioned mechanisms were formulated. If the ERA5 jets did not meet these criteria, a hypothesis was rejected. In this way, I managed to obtain a set of jet events of which I could infer with reasonable certainty which mechanism played a role in their

formation. Most confirmation was found for the baroclinic mechanism, the thermal mechanism described above. Inertial oscillations also played a role, albeit less pronounced. This might partly be due to the stricter criteria employed for this hypothesis. Orographic effects appear to be important in the few areas with steep topography. One potentially relevant effect that is not included in this analysis is a gap flow or ‘channelling’ through the Dover Strait. For this effect, I refer to the study of Capon (2003).

The results of this chapter shed an important light on the nature of coastal winds in general, and offshore low-level jets in particular. In previous chapters, I labelled the low-level jet as an ‘anomalous wind event’. This seems to imply that it’s possible to isolate low-level jets from the ‘normal’ climatology and treat them separately. However, the mechanics revealed in this chapter suggest that the coastal low-level jets are just one manifestation of the flow, embedded in a complex interplay of temperature contrasts, interaction of local circulation with background wind, and orographic effects. This has a profound impact on how to deal with them in practice. Many inflow specifications for wind energy applications don’t account for thermal effects, yet it seems that thermal effects play a major role in the manifestation of the offshore wind field. It’s about time this awe-inspiring complexity propagates to the everyday practice of wind energy science.


This climatology may serve as a starting point for follow-up studies. I can imagine several paths along which research can proceed. Previous studies on the phenomenon of coastal jets have used observations for one or a few cases. With today’s abundant interest in the offshore environment, much more data is available than ever before, which should enable more case studies. It would be interesting to make composites of several similar events, such as in Moore and Renfrew (2005). Manually classifying a number of jet events, or using the criteria of this chapter to label certain events followed by an automated search for similar cases may provide alternative means to categorize a large dataset.

Other modelling studies used highly idealized set-ups, either in a conceptual model or using large-eddy simulation. This can help to better understand the basic mechanism and the interaction between the thermal wind adjustment and inertial oscillation. Alternatively, mesoscale models could be used to experiment with the impact of local topography in a more realistic setting. In either case, these studies would be confined to single cases. For an advanced climatology, an interesting direction would be to study the results from higher resolution downscaled reanalysis datasets. A sneak preview of such a dataset is presented in the next chapter.



# CHAPTER 6





## QUALITY OF WIND CHARACTERISTICS IN RECENT REANALYSES OVER THE NORTH SEA

---

NO AMOUNT OF SIMULATION CAN PREDICT EXACTLY WHAT WILL HAPPEN IN REALITY. BUT A GOOD SIMULATION IS ENORMOUSLY PREFERABLE TO BLIND TRIAL AND ERROR.

- RICHARD DAWKINS

In chapter 2 I introduced the concept of anomalous wind events and characterized a number of events based on observations over the North Sea. Reanalysis data provide a valuable source of data to extend this climatology spatially and temporally, but that information is meaningless unless the model performance is properly evaluated. In chapter 3, I validated wind forecasts of three weather models for a selection of 30 cases. This provided valuable insight into the general model performance, but anomalous wind events were not yet considered.

In the meantime, these three models have all been used to construct three brand new reanalysis datasets. In this chapter, I evaluate the representation of wind characteristics for each of these three reanalysis datasets. New methods are demonstrated to highlight edge cases in the evaluation of model performance. The results provide important lessons about successful down-scaling strategies. This chapter also forms a substantial corroboration of chapter 3.

This chapter is under review:

Kalverla, P. C., Holtslag, A. A. M., Ronda, R. J., and Steeneveld, G. J. (2019).

Quality of wind characteristics in recent reanalyses over the North Sea. *Under review*.

## 6.1 Introduction

Ongoing climate change and the demand for more sustainable energy production have raised the interest in wind energy resources. For instance The Netherlands produced 17.0 PJ of wind energy in 2011 (onshore and offshore) and the production has increased to 35.7 PJ in 2018. Offshore wind energy was responsible for  $\sim 10$  PJ of the increased energy production. Hence, it is clear that insight in the wind characteristics over potential wind park sites is crucial. Not only mean wind characteristics, but also insight in the special dynamics of the wind are crucial for resource and load assessment studies, i.e. low-level jets, wind ramps, extreme wind shear etc. (e.g. Smedman *et al.*, 1996). However, usually offshore wind observations are relatively scarce, especially at hub heights. As an alternative, reanalysis products are important sources of wind information, either as a direct source or to drive small-scale models for wind energy purposes (e.g. Olauson, 2018; Witha *et al.*, 2019a). Reanalysis datasets build upon the data-assimilation technique, i.e. determining most probable atmospheric state that is consistent with both theory and observations, by merging numerical weather prediction (NWP) model fields and observations.

Compared to mere observations, reanalysis datasets offer a better spatial coverage and usually a longer time frame. Therefore, reanalysis datasets have become crucial for wind energy applications (Olauson, 2018). However, NWP models are used to ‘fill in the gaps’, and these models are fundamentally limited in their representation of physical processes. Thus, reanalysis datasets are subject to uncertainty as well. In fact, as argued by Parker (2016), lacking uncertainty information may be their largest weakness. Cross-validation with observations that were not assimilated into the reanalysis product may provide an intuitive means to appreciate its value for practical purposes.

This paper evaluates three relatively new reanalysis datasets, i.e. (1) ERA-5 (C3S, 2017), a global reanalysis dataset produced by the European centre for medium-range weather forecasts (ECMWF) using their Integrated Forecasting System (IFS), (2) DOWA, the Dutch Offshore Wind Atlas produced by the Dutch national weather service KNMI using their regional NWP model HARMONIE (Bengtsson *et al.*, 2017), and (3) NEWA, the New European Wind Atlas (Lundtang Petersen *et al.*, 2014; Witha *et al.*, 2019b), produced by a consortium of European research institutes using the community Weather Research and Forecasting model (WRF; Powers *et al.*, 2017) for a multitude of partly spatially overlapping domains together covering most of the European continent.

Kalverla *et al.* (2019b) evaluated the performance of these three NWP models in operational forecast mode. Although (short) forecasts and reanalyses are not the same, systematic biases may point out model weaknesses noticeable in both products. All models tended to underestimate the wind speed by up to  $0.5 \text{ m s}^{-1}$ , with a typical root mean square error of up to  $2 \text{ m s}^{-1}$ . Stable boundary layers proved to be challenging conditions, despite recent efforts to improve the turbulent mixing formulation for these conditions (e.g. Tastula *et al.*, 2012; Sandu *et al.*, 2013; Valkonen *et al.*, 2014).

Moreover, Kalverla *et al.* (2019a) extensively compared the ERA-5 wind speed data against observations at multiple sites over the North Sea and found that the overall representation of wind speed was quite good, with a maximum root mean square error of  $1.5 \text{ m s}^{-1}$ . The superior performance of ERA-5 as compared to short forecasts in Kalverla *et al.* (2019b) is likely due to the

data-assimilation. As compared to the single location and relatively small number of forecasts (30 days) evaluated in Kalverla *et al.* (2019b), this provided more significant results.

To enable a climatological description of local wind structures that are generally concealed in summary statistics, Kalverla *et al.* (2017) introduced methods to systematically study various anomalous wind events. An anomalous event describes one type of local structure, e.g. the presence of a wind speed maximum in the vertical wind speed profile (a low-level jet, henceforth: LLJ) and the corresponding fall-off (the difference between the maximum wind speed and the subsequent wind speed minimum aloft), the difference in wind speed or direction between two neighbouring vertical levels (wind shear, wind veer) or between consecutive time slots (i.e. wind ramps). Even in the absence of a characteristic local structure, a wind event can be anomalous just because it is rare. Therefore, wind extremes, strong wind speeds with long return periods, were also included in Kalverla *et al.* (2017). In their validation of ERA5, Kalverla *et al.* (2019a) focused on one of these events: the low-level jet. They showed that the representation of LLJs in ERA-5 was mediocre: one-to-one correspondence was poor, the LLJs seemed to be vertically displaced (too high) and their magnitude underestimated; however, the climatological frequency representation of LLJ characteristics was reasonable.

We present a first evaluation of wind and anomalous wind events in the aforementioned three reanalyses datasets against observations from a prospective wind farm site at the North Sea, 85 km off the Dutch coast: met mast IJmuiden (MMIJ). MMIJ is located far enough from the coast and spans a long enough period of time to show reasonable agreement with the ERA-5 data (Kalverla *et al.*, 2019a). The relatively long time span of the MMIJ dataset allows (better than other platforms) for reliable statistics. However, MMIJ will only partly reflect effects of small-scale coastal processes, and may not completely obviate validation with other, near-shore observations. Further validation of DOWA against observations from other sites and with satellite data are reported in Duncan *et al.* (2018a,b). Initial validation and sensitivity studies that were performed for NEWA may be found in Witha *et al.* (2019b).

First section 2 briefly describes the datasets and the procedure to align the data spatiotemporally. Then section 3 and 4 present a general model evaluation for wind speed and direction, followed by section 5 which concerns the evaluation of anomalous events. Section 6 features a new spatial climatology of low-level jets based on DOWA. Conclusions and perspectives are discussed in section 7.

## 6.2 Description of the datasets

Three reanalysis datasets are used in this study. ERA-5 has been developed with a horizontal grid spacing of  $\sim 30$  km at mid-latitudes. Compared to ERA-5, DOWA assimilates additional regional observations and uses a dense grid spacing of 2.5 km. Every 3 h, data-assimilation is applied to initialise a new forecast cycle. The assimilated observations include ASCAT satellite sea-surface wind fields and MODE-S EHS aircraft wind profile measurements. NEWA also has a fine grid spacing of 3 km, but was produced with a slightly different procedure: it consists of 8-day runs in which some model fields are nudged towards the ERA-5 reanalysis data to prevent the simulations from drifting away from the synoptic situation. Table 6.1 further summarises relevant characteristics of the reanalysis data.

**Table 6.1:** Summary of the most important characteristics of the three reanalysis datasets discussed in this chapter, including references to further documentation.

	ERA5	DOWA	NEWA
Full name	ECMWF Retrospective Analysis 5th generation	Dutch Offshore Wind Atlas	NEW European Wind Atlas
Time span	1950-present	2008-2017	1989-2018
Horizontal resolution	31 km	2.5 km	27, 9, 3 km (1-way nests)
Vertical resolution	137 levels up to 0.01 hPa.	65 levels up to 10 hPa	61 levels up to 50 hPa
Hydostatic assumption	Yes	No	No
Domain	Global	The Netherlands, Southern North Sea, Wadden Sea	Outer domain Europe, with 10 regional sub-domains
Data-assimilation	4D-VAR with 12-hourly assimilation windows	3-hourly 3D-VAR	-
Data-sources	Satellites and in-situ observations (see documentation)	Routine observations, ASCAT (satellite), MODE-S EHS (aircraft)	-
Initialization	From analysis	From previous forecast with additional data-assimilation	ERA5
Strategy	Hourly surface analysis (2D optimal interpolation)	Concatenated 3-hour forecasts.	8-day forecasts including 24 h spin-up
Lateral boundaries	-	ERA5	ERA5
Nudging	-	-	Spectral nudging in outer domain, above PBL and level 20
Uncertainty	10-member ensemble at 62 km resolution	Uncertainty assessment report	Sensitivity experiments and multi-physics ensemble
Model	IFS Cycle 41r2	HARMONIE-AROME, Cycle 40h1.2.tg2	WRF V3.8.1 (with modifications)
Website	<a href="https://confluence.ecmwf.int/display/CKB/ERA5+data+documentation">confluence.ecmwf.int/display/CKB/ERA5+data+documentation</a>	<a href="http://www.dutchoffshorewindatlas.nl">www.dutchoffshorewindatlas.nl</a>	<a href="http://map.neweuropeanwindatlas.eu">map.neweuropeanwindatlas.eu</a>
Reference/documentation	C3S (2017); ECMWF (2016)	Bengtsson <i>et al.</i> (2017)	Witha <i>et al.</i> (2019b); Gonzalez Rouco <i>et al.</i> (2019)

For validation we use the MMIJ dataset, which spans four years of observations (2012-2015) at several altitudes up to 315 m, spaced approximately 25 m apart. Observations at 27, 58, and 90 m are from mast-mounted cup- and sonic anemometers and wind vanes, while the data beyond 115 m were obtained with an upward-pointing continuous-wave lidar. More details can be found in Werkhoven and Verhoef (2012) and Kalverla *et al.* (2017, 2019a). The large temporal extent, distance to shore and vertical measurement range makes this dataset optimally suited to characterize the wind climate. Data quality was found to be very good (e.g. Maureira Poveda and Wouters, 2015), with only a few gaps (figure 6.1a). The observation data is available at 10-min intervals, and were hourly-averaged to facilitate comparison with the reanalysis data. The reanalysis data were vertically (linearly) interpolated to, and temporally aligned with the observations to obtain four collocated datasets (figure 6.1b-d).

## 6.3 Evaluation of wind speed and the role of atmospheric stability

Time-averaged wind speed profiles (figure 6.2a) demonstrate a striking correspondence between DOWA and the observations, considering that MMIJ observations were not assimilated into the reanalysis product. ERA-5 underestimates the wind speed by  $\sim 0.5 \text{ m s}^{-1}$  through the whole profile while NEWA is nearly unbiased near the surface and reaches a slow bias of  $0.5 \text{ m s}^{-1}$  at 300 m. As a consequence, NEWA appears to underestimate the wind speed shear within the

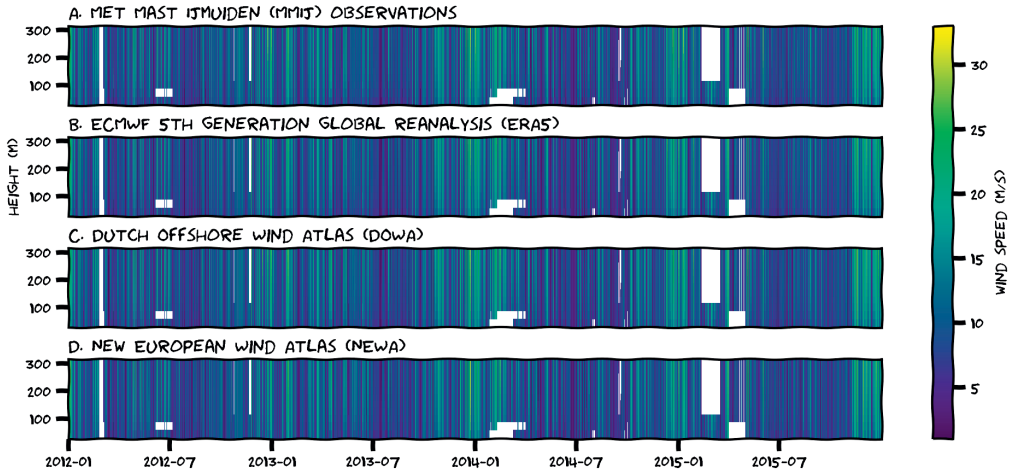


Figure 6.1: Time series of vertical wind speed profiles a) observed at Meteo Mast IJmuiden and b-d) three reanalysis datasets.

layer. Parsons *et al.* (2018) showed that ERA-5 on average shows a good skill for wind speed, but that ERA-5 underestimates the wind speed for very extreme sea states. Overall, the underestimated wind speed is consistent with Couto *et al.* (2019) who found a similar bias in front of the Portuguese coast. A more complete picture is obtained if we consider figure 6.2b, which shows both the bias and RMSE due to phase differences. Although DOWA is nearly unbiased, it does not exactly align with the observations, leaving an overall RMSE of  $\sim 1.5 \text{ m s}^{-1}$ . Only near the surface a small bias of  $0.1 \text{ m s}^{-1}$  is found. For ERA-5 and NEWA, the negative bias is clearly present, and its altitude-dependence is also apparent. For NEWA, the random errors are larger than for ERA-5 and DOWA, especially away from the surface.

Figure 6.3 reveals clearly that stability affects the bias in all datasets. For the most unstable conditions, NEWA overestimates the near surface wind by  $0.6 \text{ m s}^{-1}$  and this bias decreases aloft. While ERA-5 and DOWA show the same shape of the bias profile, they mainly underestimate the wind in the upper part of the profile. For all datasets the bias reduces and appears more uniform with height for moderate unstable stratification ( $-0.1 < Ri_b < -0.025$ ). The near neutral class contains the majority of datapoints ( $-0.025 < Ri_b < 0.0$ ) and herein DOWA is nearly unbiased while NEWA underestimates the wind speed, which increases with height. Surprisingly ERA-5 shows a slow bias of  $\sim 0.7 \text{ m s}^{-1}$  near the surface, though its bias decreases with height. The most prominent biases and sensitivities occur for stable conditions, as consistent with findings in Baas *et al.* (2016) for the HARMONIE model results over the North Sea. For near neutral conditions NEWA reveals a slow bias, which switches to a wind speed overestimation for  $Ri_b > 0.025$ , which increases for stronger stability, even to  $1.0 \text{ m s}^{-1}$  for ( $0.075 < Ri_b < 0.1$ ). For that class DOWA also overestimates the wind by  $\sim 0.5 \text{ m s}^{-1}$ , while ERA-5 represents this class rather well. The wind speed in NEWA is more accurate near the surface, which does not support a deficiency in the surface roughness formulation. Rather, it seems that too little momentum is transported downward to the surface. This could be a result of the large-scale nudging strategy employed in the NEWA. Above the boundary layer, momentum fields were nudged towards the ERA-5 values. If wind speed is underestimated in ERA-5, it is

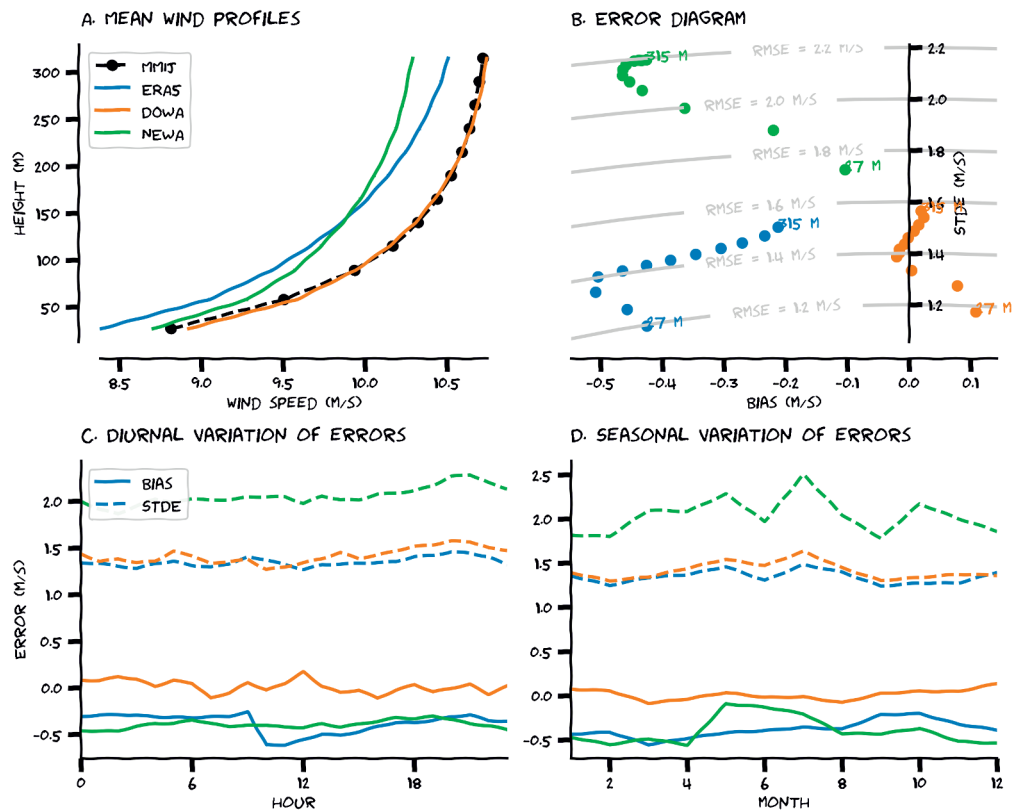


Figure 6.2: a. Observed and modelled time-averaged vertical wind speed profile at MM1J and b. Error diagram showing the mean (bias) and standard deviation (STDE) of the error distributions of three reanalysis datasets as compared to observations from met mast IJmuiden at multiple heights. c-d. Diurnal and seasonal evolution of the mean and standard deviation of the same error distributions.

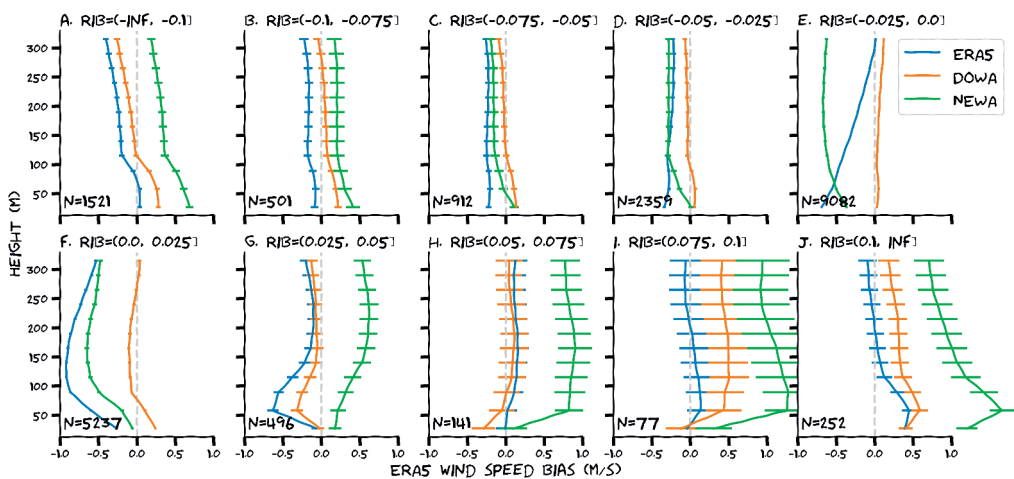


Figure 6.3: Vertical profile of the wind speed bias in all reanalysis datasets for different stability intervals (based on the observed bulk Richardson number). Top row: decreasingly unstable stratification. Bottom row: increasingly stable stratification.



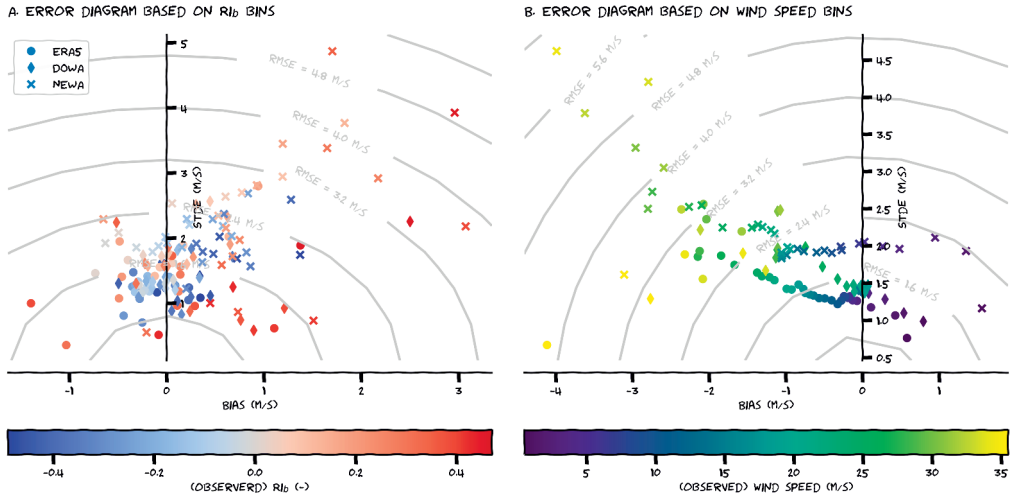
thus very plausible that this error propagates to NEWA. DOWA, in contrast, is completely free in the inner domain, except for the 3-hourly data-assimilation updates.

Examination of the seasonal cycle reveals that the slow bias in ERA-5 is present throughout the lowest 300 m, but is most prominent from September to February with a maximum negative bias at the surface. From March to August the bias profile shows a maximum at  $\sim 100$  m (not shown). The wind speed underestimation in ERA-5 might be explained by the surface roughness, or the Charnock parameter, which dictates the relation between wind and waves. The wind speed seems to be mainly underestimated near the surface, which may point to an overestimation of the surface roughness. However, a comparison of the modelled significant wave height against wave height observations at the nearby K13 platform indicates that ERA5 slightly underestimates the wave height by 9 % (not shown). Unravelling the wave height biases further in classes of atmospheric stability, we find the largest underestimation for  $Ri_b > 0.05$  of 33 %. For  $Ri_b < -0.05$  the wave heights are underestimated by 26 %. At near neutral conditions the wave heights in ERA-5 do show smaller biases. Although roughness and wind speed are interdependent over sea, it seems the wind speed underestimation is not triggered by an overestimated roughness here. The representation of atmospheric stability and the turbulent mixing under stable conditions are more likely explanations. Alternatively, the relatively coarse resolution of ERA-5 may induce a smoothing effect, especially for high wind speed events.

The smoothing effect described above can effectively suppress the random errors shown in figure 6.2. Phase differences such as a delayed front passage will lead to a double penalty. The inability to reproduce small-scale features thus prevents a double penalty, which explains why the ERA-5 data performs relatively well in terms of the standard deviation of the error (STDE, figure 6.2b). While one expects that higher-resolution models generally suffer more from this problem, DOWA performs similar to ERA-5, which is presumably the merit of frequent data-assimilation in DOWA (Duncan *et al.*, 2018a).

A remarkable discontinuity in the diurnal cycle of wind speed was revealed in the ERA-5 data in Kalverla *et al.* (2019a). At 10:00 UTC, the wind speed bias suddenly strengthens. To verify whether this artefact propagated to the other reanalysis datasets, the diurnal and seasonal cycle of the wind speed bias and STDE are shown in figure 6.2c-d. The discontinuity in the diurnal cycle only occurs for ERA-5. However, in the seasonal cycle, we find another remarkable feature in the NEWA data, i.e. a smaller bias in spring and early summer. The reason becomes clear upon inspection of figure 6.4, which shows the wind speed bias and STDE as function of (observed) wind speed and stability. Stable conditions lead to a substantial positive bias, while high wind speeds lead to a large negative bias. In other words, all models but especially NEWA tend to underestimate very strong winds, while they overestimate the wind speed during stable conditions. Since winds are generally stronger in winter and stable conditions occur more frequently in spring and summer, this helps to explain the seasonal cycle of the bias in NEWA.

The results in figure 6.4a are consistent with Kalverla *et al.* (2019b), who found a reduced model performance in stable conditions. Hence our current results are a substantial corroboration of this earlier result. Apparently, stable conditions are still challenging (Holtslag *et al.*, 2013; Sandu *et al.*, 2013; Steeneveld, 2014; Tsiringakis *et al.*, 2017), despite recent efforts to improve the turbulent mixing formulation (Riddaway, 2014; Bengtsson *et al.*, 2017; Olson *et al.*, 2019a,b). Furthermore, the results in figure 6.4b support the hypothesis that the slow speed bias results from a smoothing effect, as that would manifest itself most clearly for distinct wind speed maxima.

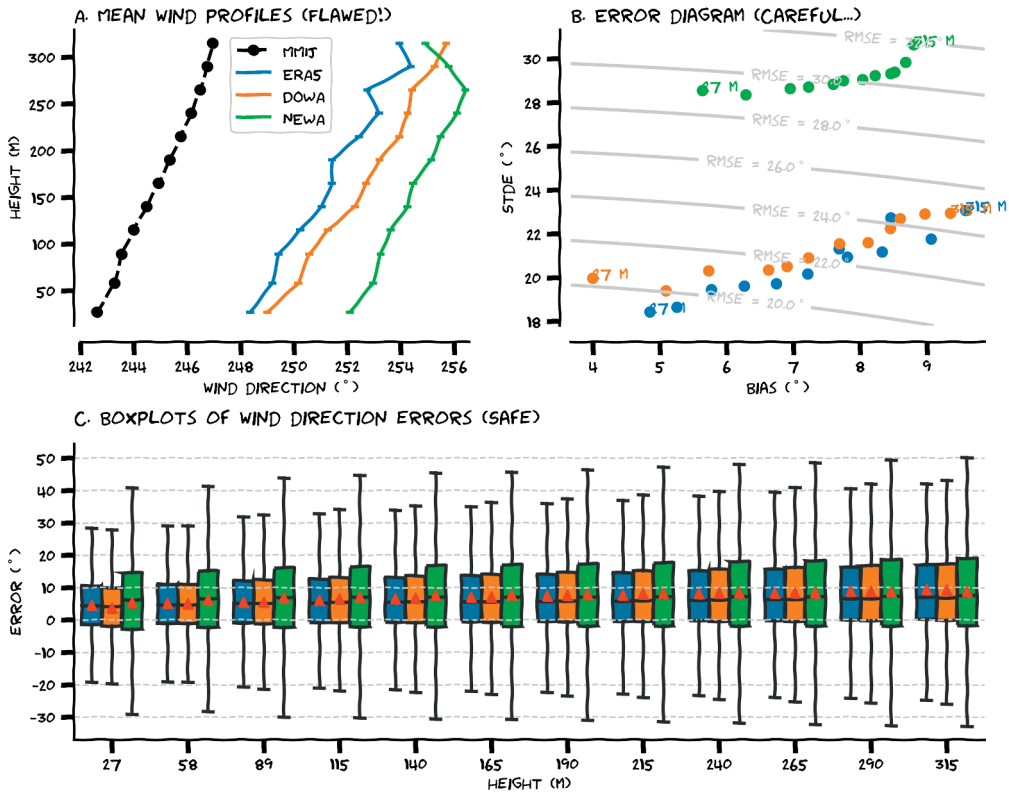


**Figure 6.4:** Error diagrams illustrating the performance of the three reanalysis datasets for mean wind speed below 300 m in the period 2012 to 2015 as function of a) Atmospheric stability and b) Wind speed itself. Observation data were used to aggregate the error statistics.

## 6.4 Evaluation of wind direction

Wind direction is crucial for offshore wind energy purposes for determining the directional shear on wind turbines and wind farm wake effects, understanding the model's representation of boundary-layer friction, and its representation of advection of onshore atmospheric phenomena towards e.g. offshore wind parks (Dörenkämper *et al.*, 2015; Wagner *et al.*, 2019). Here we discuss two methods to evaluate wind direction, i.e. the bias as the difference between the means of the wind direction of two samples (figure 6.5a), and the bias as the mean of wind direction differences (figure 6.5b). These definitions are equivalent for arithmetic means. Figure 6.5a suggests that the bias is more or less constant with height, figure 6.5b suggests it increases steadily. This apparent inconsistency stems from the use of vector means instead of arithmetic means in figure 6.5a – a common method to compute angular statistics. The difference between the vector mean and the arithmetic mean is greater when the angles are widely distributed (Jammalamadaka and SenGupta, 2001). Figure 6.5b shows that the standard deviation increases with height, thus the apparent wind veer with height might represent a statistical artefact rather than a physical effect.

To circumvent the pitfalls of circular statistics, the performance for wind direction can be inferred directly from the error distribution of the wind direction (figure 6.5c). Indeed, both the width and the mean of the error distribution increase with height, consistent with figure 6.5b. The positions of the means relative to the medians, and the upward shift of the 75-percentile as compared to the relatively constant location of the 25-percentiles indicate a changing skewness with height. From a physical point of view the wind in the reanalysis veers with respect to the observations, and this veering increases with height. These results confirm findings in Kalverla *et al.* (2019b) and previous literature, where the models' inability to represent a realistic wind veer with height was related to excessive mixing in stable conditions and to strong baroclinity (Brown *et al.*, 2005; Holtslag *et al.*, 2013; Sandu *et al.*, 2013). Despite the remaining biases, wind direction in ERA-5 and DOWA is better represented than in the 30 operational forecasts



**Figure 6.5:** Modelled and observed wind direction using a) circular statistics and b) mean of difference in wind direction. c) Boxplots of the wind direction error distributions for three reanalysis datasets as compared to MMIJ observations at multiple levels. Red triangles denote the arithmetic means. Outliers are not drawn, because that would require axes limits of up to  $\pm 180^\circ$ .


evaluated in Kalverla *et al.* (2019b), probably due to data-assimilation, or model improvements discussed in Riddaway (2014) and Bengtsson *et al.* (2017). Again, the relative wide error distributions in NEWA may be the result of a substantial double penalty, considering that NEWA consists of 8-day simulations without data-assimilation – as compared to the 3-hourly update cycles for DOWA.

## 6.5 Anomalous wind events in reanalysis products

Building upon the model representation for general wind characteristics, this section discusses the representation of anomalous events in these datasets. The methodology used to assess model performance is explained after the subsection about LLJs, using the LLJ data as illustration.

### 6.5.1 Low-level jets

Low-level jets are wind profiles with a wind speed maximum near the surface, as illustrated e.g. in Kalverla *et al.* (2019a). LLJs over MMIJ occur primarily in spring and early summer, often



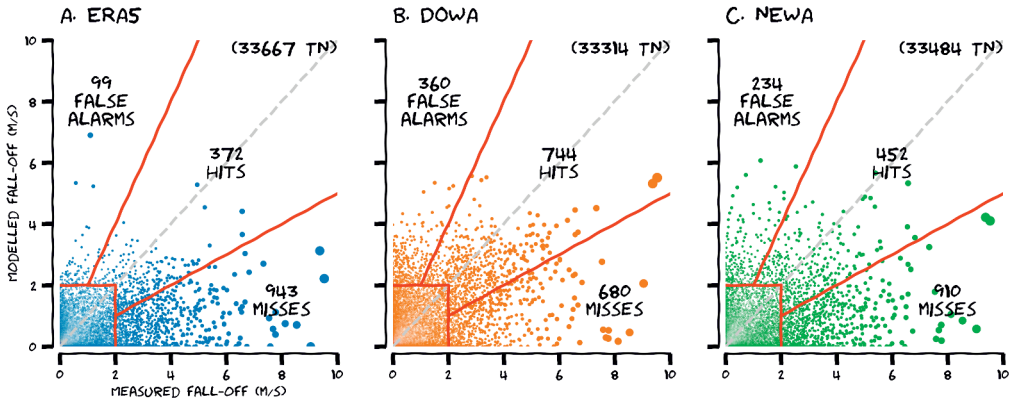
**Figure 6.6:** Seasonal (a) and diurnal (b) cycle of low-level jets, based on observations and reanalysis data. Low-level jets are defined as all hourly wind speed profiles with a maximum exceeding a fall-off threshold of  $2 \text{ m s}^{-1}$  at the location of met mast IJmuiden, expressed as percentage of the total number of wind profiles in that month/hour.

appear at the end of the afternoon and persist until the next morning. Wagner *et al.* (2019) studied the LLJ climatology at the FINO1 site in the German Bight and found LLJs occur at 14.5% of the time and 64.8% of the days, mostly from directions between east and south. They are formed by a variety of mechanisms, but baroclinic effects, orographic effects Capon (e.g. flow is forced through the Dover Strait, see 2003), and the combination with a stable boundary layer explain most of their characteristics. Recently, Kalverla *et al.* (2019a) demonstrated that LLJs are present in the ERA-5 data, although they tend to be located too high above the surface. Consequently, when the ERA-5 data are interpolated to observation heights, the amount of LLJs is grossly underestimated. However, the seasonality could still be faintly recognized at MMIJ.

To investigate whether the refined datasets improve upon the representation of LLJs, we study the seasonal and diurnal cycles for all three reanalysis datasets. (figure 6.6). The observations exhibit a pronounced seasonal cycle, the erratic nature of which has been discussed at length in Kalverla *et al.* (2019a). In the four-year observation period (2012 to 2015), May and July saw relatively many LLJ events compared to April and June.

Indeed, ERA-5 grossly underestimates the amplitude of the seasonal cycle, and completely misses the peak in May. The other two datasets, especially DOWA, demonstrate considerably better skill. The diurnal cycle is characterised by a distinct dip around noon, and peaks in the afternoon and the early morning, blended together in the night. The afternoon peak – presumably related to the adjustment of the sea-breeze – appears to be best represented in the reanalysis data. If two different mechanisms are responsible for LLJ formation, one of these mechanisms might be better resolved than the other. Alternatively, the formation mechanism might be relatively well represented, but the jets' propagation through the night proves challenging. Further investigation of individual LLJ events could provide a definitive answer in this matter, but that exercise is beyond the scope of this evaluation.

Often, a fall-off threshold is used to distinguish between 'real' LLJs and 'normal' conditions that happen to show a weak wind speed maximum by chance. In line with previous studies, a fall-off threshold of  $2 \text{ m s}^{-1}$  was used for figure 6.6. Alternatively, the absolute fall-off may be inspected directly. This is shown in figure 6.7, where the red box near the origin indicates



**Figure 6.7:** Scatter plots of absolute fall-off of low-level jets as represented in reanalysis datasets versus observed fall-off. A  $2 \text{ m s}^{-1}$  fall-off threshold is indicated in red. Outside this region, red lines indicate the region where reanalysis and observations agree on the absolute fall-off value to within a factor of 2. The dashed grey line indicates a 1:1 correspondence. Scatter points are scaled with distance from origin to expose the structure both in dense and sparse regions of the graph. TN is an abbreviation for true negatives, relating to the bottom left corner in each panel.

the fall-off threshold of  $2 \text{ m s}^{-1}$ . All points outside this box may be regarded as significant LLJ events – either observed, modelled, or (preferably) both.

### 6.5.2 Quantification of model performance for anomalous events

Usually model performance is expressed in summary statistics or a variety of skill scores. With the current data the signal will be dominated by non-significant events, and since the differences are subtle, mismatches in timing will lead to very low correlation coefficients (notice the dense clustering of scatter points along the zero lines of both axes of figure 6.7). It is reasonable to suppose that a stronger LLJ event is more likely to be picked up by the reanalysis data, and as it occurs, these most anomalous events are the main focus of this section.

To quantify model performance, we establish the following contingency ‘rules’: (1) If either observations or reanalysis data report on the presence of a significant LLJ event, and the absolute fall-off in both datasets is comparable to within a factor 2, then the model performance is satisfactory and the data point is counted as a hit. (2) If a significant LLJ is observed, but not present in the reanalysis data, or if it is present in the reanalysis data but much weaker than observed (less than half as strong), then this data point is regarded a miss. (3) If a significant LLJ is present in the reanalysis data, but it is not observed, or it is observed but the reanalysis overestimates its strength by at least a factor 2, then it is labelled as a false alarm. All other events are true negatives.

The contingency rules allow for estimating skill scores, such as the probability of detection, false alarm rate, or critical success index (CSI; see Schaefer 1990). The CSI, defined as

$$CSI = \frac{\text{hits}}{(\text{hits} + \text{misses} + \text{false alarms})}$$

is a simple and intuitive parameter to compare the performance of several models: the score increases if more events are correctly predicted, and it decreases as more events are missed or falsely predicted.

**Table 6.2:** Critical success index (CSI), symmetric extreme dependency score (SEDS) and frequency bias (FBIAS) for the representation of low-level jets in ERA5, DOWA and NEWA.

	ERA5	DOWA	NEWA
CSI	0.26	0.42	0.28
SEDS	0.67	0.73	0.65
FBIAS	0.36	0.78	0.50

An alternative and more robust (but less intuitive) statistic than the CSI is the symmetric extreme dependency score, defined as (Hogan *et al.*, 2009):

$$SEDS = \frac{\ln[(hits + false\ alarms)/n] + \ln[(hits + misses)/n]}{\ln(hits/n)}$$

where  $n$  is the total number of events (hits+misses+false alarms+true negatives). The SEDS varies between  $-1$  and  $1$ , where  $1$  indicates a perfect forecast, a random forecast would receive a skill score of  $0$ , and a forecast that actually deteriorates the quality of a random forecast tends to  $-1$ .

The CSI and SEDS both penalize phase errors, which is desirable in forecast verification. However, for climatological studies for resource assessment phase errors are not essential. An alternative score that measures the climatological model performance is the (frequency) bias:

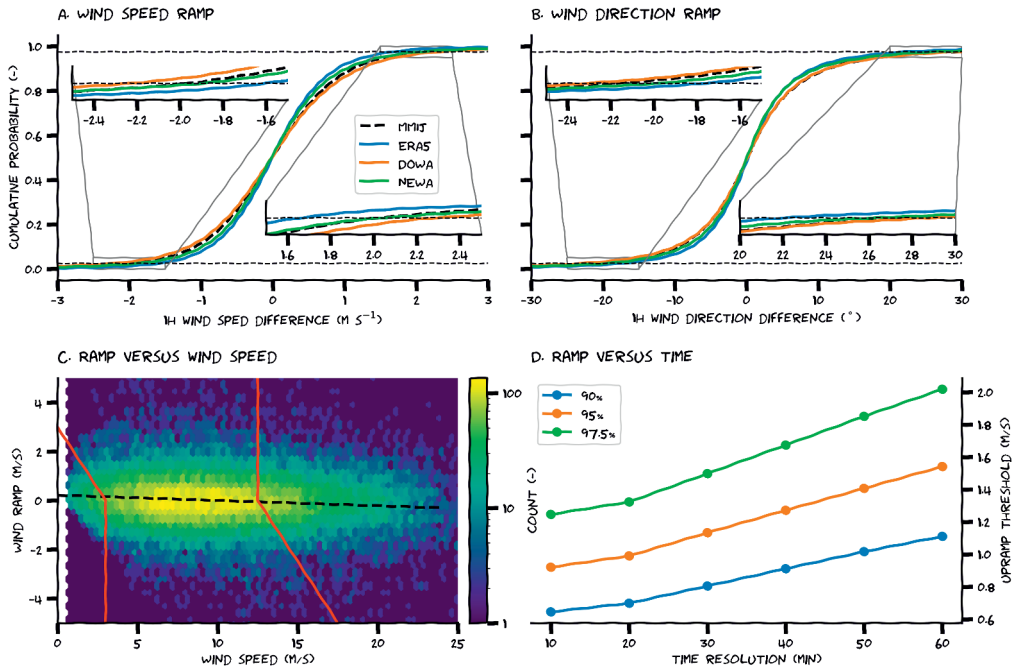
$$FBIAS = \frac{total\ predicted}{total\ observed} = \frac{hits + false\ alarms}{hits + misses}$$

A frequency bias of  $\sim 1$  means that the total amount of LLJ events is more or less correct, even if the timing is off.  $FBIAS < 1$ , represents an underestimation of the amount of LLJ events and vice versa. A downside of this score is that the model might be right for the wrong reason, although it seems unlikely that the model would randomly predict subtle features such as LLJs.

All reanalysis datasets underestimate the amount of significant LLJ events (table 6.2). In four years' time ( $\sim 35\,000$  hourly wind profiles), only 372 LLJs have been correctly captured in ERA-5, against 943 missed events. With 'only' 99 false alarms, this results in a frequency bias of  $0.36$ . DOWA performs better and picks up approximately twice as much LLJ events, reflected in a much higher frequency bias of  $0.78$ . The representation in NEWA is intermediate: more hits than ERA-5, but substantially more false alarms. Hence, (1) although NEWA does not seem to improve upon ERA-5 with respect to a general validation of the wind speed profiles, the increased resolution does favour the climatological LLJ representation and (2) especially DOWA improves upon the representation of the dynamical conditions in coastal areas.

### 6.5.3 Wind ramps

Wind ramps are rapid changes of the wind speed and/or direction in time. In climatologies for wind energy applications, the mean wind is often assumed to be stationary, or time-averaged statistics are considered. Therefore, wind ramps are interesting anomalous events that require additional, tailored evaluation. Kalverla *et al.* (2017) determined the wind speed and direction differences over various time intervals in the MMIJ dataset, and studied the frequency distributions to build some intuition about the magnitude of these difference. Naturally, the frequency



**Figure 6.8:** Cumulative probability distribution of 1-hourly a) wind speed and b) direction differences based on observations and reanalysis data. The dashed lines indicate the 2.5- and 97.5-percentiles as thresholds for the classification of typical up- and down-ramp events. c). Hexbin visualization of hourly MMJ data illustrating the frequency of wind speed ramps as function of wind speed. The dashed line represents a linear fit, the red lines denote the area where the wind ramp causes a disturbance within the cubic part of a conceptual power curve. d) Wind speed ramp as function of the time resolution of the underlying data, based on MMJ data resampled at different intervals using a moving average to mimic a smoothing effect. Three different up-ramp thresholds are used to assess the robustness of the relation.

distributions centres around zero, for in the absence of a long-term trend, increasing wind speeds must be balanced by equivalent decreases. For wind direction, this is not necessarily true, but it was found in practice. Kalverla *et al.* (2017) used the 5- and 95-percentiles to obtain site-specific characteristic up-ramp and down-ramp thresholds, and analysed the sensitivity to this threshold.

Here we use the 2.5- and 97.5-percentiles instead, to put even more emphasis on the most extreme conditions. The cumulative probabilities of hourly differences in wind speed and direction based on MMJ and the corresponding grid points in the reanalysis datasets are shown in (figure 6.8a-b). In general, only small differences appear between the datasets, though the distribution of wind speed differences in DOWA is slightly broader than observed, and the distribution is slightly too narrow in ERA-5 and DOWA. The differences between datasets are quantified through the 2.5- and 97.5-percentiles. A typical 1 h down-ramp at MMJ amounts to  $-2.0 \text{ m s}^{-1}$ , while ERA-5, DOWA and NEWA estimate the ramp intensity at  $-1.6$ ,  $-2.2$ , and  $-1.9 \text{ m s}^{-1}$ , respectively. Typical up-ramp values amount to  $2.0 \text{ m s}^{-1}$  according to observations, and  $1.7$ ,  $2.3$ , and  $2.0 \text{ m s}^{-1}$  for the respective reanalysis datasets. Thus, NEWA best captures the climatology of wind speed ramps. Before further quantification, some notes must be made about the evaluation of wind ramps.

Figure 6.8c shows the joint distribution of wind speed and subsequent hourly wind speed differences at MMIJ. Such a representation might be relevant for forecasting applications, where wind ramps within the cubic part of the power curve lead to the largest power fluctuations. A slightly negative correlation between wind speed and wind speed difference appears. This makes sense: if the wind speed is already high, it is more likely to decrease than increase.

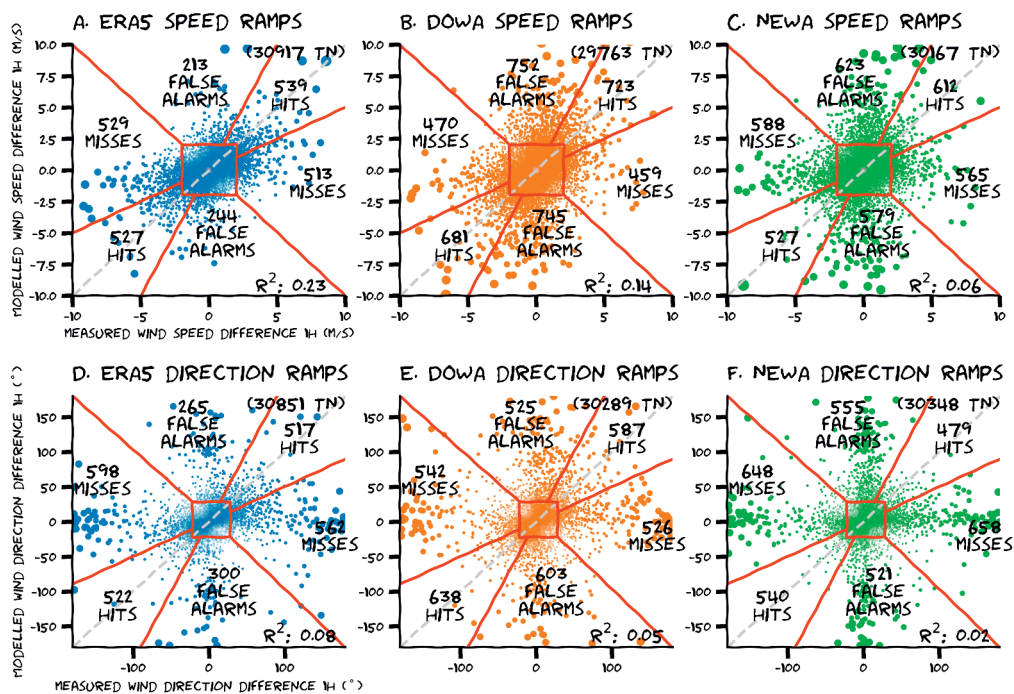
Ramps on relatively short time scales are probably more relevant for energy applications, but unfortunately ERA-5 and DOWA do not offer short timescales (NEWA is available every 30 minutes). To provide some intuition about the relevance of hourly ramps, figure 6.8d depicts the typical wind speed up-ramps as function of the time window for MMIJ. Herein, a moving average and a resampling (with corresponding time window) were subsequently applied to the MMIJ data. The 97.5-percentile increases almost linearly with the time window of the underlying data, except below 20 minutes, where the acceleration appears smaller. To show the robustness of this result, and also to illustrate the sensitivity to the ramp threshold, three different percentiles are shown in figure 6.8d. Both the typical ramp and the acceleration are smaller for lower ramp thresholds, which makes sense. The current results differ from Kalverla *et al.* (2017), since they only applied resampling (no moving average). The averaging appears to be responsible for the decreased acceleration between 10 and 20 minutes (not shown). These factors explain why the empirical square root relation between wind ramp magnitude and time window found in Kalverla *et al.* (2017) does not correspond with the present results.

Upon visual inspection, the distribution of wind direction differences is well-captured by DOWA and slightly underestimated in the other datasets (figure 6.8b). Typical 1 h direction up- and down-ramps at MMIJ are  $-22^\circ$  and  $28^\circ$ , respectively. ERA-5 underestimates both thresholds:  $-18^\circ$  and  $22^\circ$ . DOWA slightly overestimates them:  $-23^\circ$  and  $29^\circ$ . NEWA reports thresholds of  $-20^\circ$  and  $26^\circ$ , i.e. a small underestimation. The asymmetry of the distribution is probably related to frontal passages, which are accompanied by an abrupt wind veer. Arguably, wind direction ramps are much more harmful for wind turbines because they can lead to yaw misalignment (and therefore additional structural loads). For wind farms, they can lead to rapid power fluctuations – in that case, the emphasis is more on forecasting than on climatological representation.

While reanalyses are not actual forecasts, it is illustrative to inspect the 1:1 correspondence between observed and simulated wind ramp events (figure 6.9a-c). The  $R^2$  value corresponding to a linear fit is annotated for reference, but this parameter is mostly determined by the bulk of the data, while our main interest are the extreme cases. Hence, the contingency rules from the previous section are used. The red box in the middle corresponds to the typical up- and down-ramp thresholds as observed. In this case, the distribution of wind speed differences is two-sided, i.e. an additional possibility where the reanalysis would ‘predict’ an up-ramp while a down-ramp is actually observed, or vice-versa is present. Although this rarely occurs in practice, the possibility requires an additional rule to distinguish between false up-ramp alarms and missed down-ramp events (and vice versa). The rule employed here can be summarised as ‘whichever is greater’. The corresponding skill scores are provided in table 6.3.

The frequency bias again demonstrates that ERA-5 underestimates the amount of wind ramps. This is expected, since the relatively coarse horizontal grid spacing of 30 km seriously limits the model representation of small-scale structures that are responsible for ramp events. However, the number of false alarms is also limited (at least, relative to the other datasets). Consequently, the CSI and SEDS are highest for ERA-5. The false alarm rate is much higher in DOWA, which





**Figure 6.9:** Scatter plot of 1-hourly wind speed (a-c) and direction (d-f) difference as estimated from ERA-5, DOWA, and NEWA. In this figure, both positive and negative values are considered. The ramp thresholds in red have been based on the 2.5- and 97.5-percentiles of the observed ramps. Additional red lines have been inserted to separate missed up-ramp events from false down-ramp alarms and vice versa.

**Table 6.3:** Skill scores for wind ramps and wind shear ('shear' and 'veer' are used here to refer to the longitudinal and lateral components). Up- and down-ramps have been combined in the calculation of skill scores, although it would be possible to calculate them separately as well.

	ERA5	DOWA	NEWA	ERA5	DOWA	NEWA
	Speed ramps			Direction ramps		
CSI	0.42	0.37	0.33	0.38	0.36	0.30
SEDS	0.70	0.61	0.58	0.66	0.61	0.56
FBIAS	0.72	1.24	1.02	0.73	1.03	0.9
	Wind shear			Wind veer		
CSI	0.40	0.69	0.66	0.00	0.00	0.00
SEDS	0.73	0.88	0.86	-0.11	-0.05	-0.04
FBIAS	0.51	0.80	0.84	0.09	0.14	0.13

indicates an overestimation of ramp events ( $FBIAS = 1.24$ ). NEWA has almost no frequency bias, but the 1:1 correspondence with observations is particularly poor, as reflected by the high amount of misses and false alarms compared to the number of hits. In other words, a more realistic climatology of ramps in NEWA comes with a deterioration in timing of these events. While the climatology is more important during the resource assessment and planning phase, correct timing is obviously quite relevant for forecasting applications.

Model performance for wind direction ramps is visualised and quantified in figure 6.9d-f and table 6.3. The interpretation is analogous as for wind speed ramps: ERA-5 underestimates wind ramps, DOWA in this case performs best in a climatological sense, and especially NEWA struggles with the timing of events. Further investigation is needed to assess whether the physical characteristics of the wind ramps (both in wind speed and direction) are consistent between all four datasets.

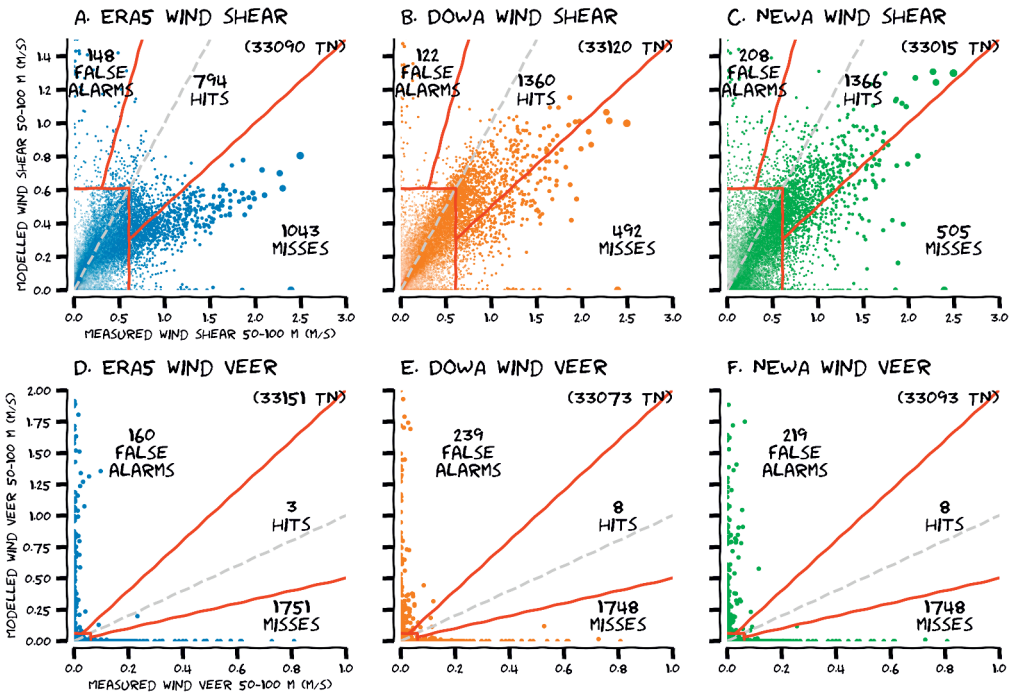
### 6.5.4 Extreme shear

It is illustrative to split the wind vector in a streamwise (or: longitudinal) and a normal (or: spanwise/lateral) component (Kalverla *et al.*, 2017). If the wind would turn without a change in magnitude, a substantial wind shear would remain hidden if only wind speed was analysed. Besides, the energy in the lateral wind component would be falsely regarded in load and power calculations. Thus, the analysis of extreme shear starts with an evaluation of the longitudinal wind component (aligned with the 115 m wind, approximately hub-height) in figure 6.10a-c, which is based upon accumulated wind shear between 50 and 100 m.

Clearly, the degree of wind shear is underestimated by all reanalysis datasets. Following the convention of Kalverla *et al.* (2017), the extreme shear threshold is defined as the 95-percentile. This value is  $0.61 \text{ m s}^{-1}$  according to the MMJ data, while ERA-5, DOWA and NEWA estimate it at 0.41, 0.49, and  $0.47 \text{ m s}^{-1}$ , respectively. Especially for ERA-5, where the timing is quite well-represented, the underestimation of the absolute wind shear is so large that most extreme events (exceeding the 95-percentile of the observations) are classified as missed events. Both DOWA and NEWA perform much better in this respect, with slightly better performance for DOWA.

It is interesting to explore the causes behind general underestimation of wind shear and the difference between the datasets. Earlier we found that the surface roughness in ERA-5 is slightly underestimated. Stable stratification suppresses turbulent mixing and thereby supports the development of strong shear (and LLJs). As shown in section 4 and in Kalverla *et al.* (2019b), NWP models still struggle to adequately represent these conditions. An alternative rendering of figure 6.10 in which the observed Richardson number was used to further categorise the data in the scatter points (not shown) revealed that the majority of the extreme shear cases are indeed characterised by stable stratification, while unstable and neutral cases with little wind shear make up the bulk of the data. This explains why the underestimation of extreme shear is not reflected in the mean wind profile in figure 6.2a.

The lateral wind shear, which may be interpreted as a measure of vertical wind veer, is evaluated in figure 6.10d-f. The 95-percentile threshold of the extreme accumulated shear over the layer 50-100m is very low,  $0.06 \text{ m s}^{-1}$ , and underestimated in all reanalysis datasets (0.001, 0.002, and  $0.001 \text{ m s}^{-1}$ ; but these values are hardly significant). Almost all extreme events are missed by the reanalysis data, and the fact that these misses are barely compensated by false alarms



**Figure 6.10:** Scatter plot of modelled and observed accumulated wind shear in the longitudinal (a-c) and lateral (d-f) wind components between 50 - 100m. The red lines delimiting the extreme/anomalous shear threshold are based on the 95-percentiles of the observations.

confirms that also in a climatological sense, the reanalyses provide a poor impression of lateral (extreme) shear. Since turbulence generally tends to destroy vertical gradients, it is likely that these errors are the result of excessive mixing, either due to inadequate representation of the physics, including sufficient resolution, or due to misrepresentation of atmospheric stability, as discussed previously.

### 6.5.5 Wind extremes


Finally, we address wind speed extremes as anomalous events because the conventional statistics may not adequately capture them. E.g., the Weibull fit is strongly determined by the bulk of the data, but especially rare events in the tail may be relevant for structural loads. In Kalverla *et al.* (2017), extreme value theory was therefore applied to estimate the 50-year extreme wind speed, based on the IJmuiden observations. Because four years is too short to select only annual maxima, the method of independent storms (Palutikof *et al.*, 1999) was used to select  $\sim 40$  unrelated events within the measurement period. Here this analysis is repeated for the reanalysis data (figure 6.11), though with hourly-averaged observations rather than 10-min observations. Thus, here we find a much lower estimates of the 50-year wind speed extreme. Since design standards are based on the 10-min. estimate (Burton *et al.*, 2011), this section mostly serves as model validation and intercomparison.

The smoothing effect of a relatively coarse model resolution has been used in section 6.3 to explain the underestimation of high wind speeds in ERA-5 and NEWA, since the latter is nudged

towards the ERA-5 momentum fields. A similar effect is introduced by the time-averaging of the observations. The impact is substantial: while an extreme value of  $42.7 \pm 2.4 \text{ m s}^{-1}$  was reported in Kalverla *et al.* (2017), here we find  $36.7 \pm 2.1 \text{ m s}^{-1}$ , a difference of  $6 \text{ m s}^{-1}$ ! Consequently, the reanalysis data should not be used directly to estimate wind extremes. Nevertheless, the difference between the three reanalysis datasets can be compared to the uncertainty related to the spatio-temporal characteristics. With a 50-year extreme value of  $35.9 \pm 2.5 \text{ m s}^{-1}$ , ERA-5 closely approaches the estimate based on hourly-averaged observations. DOWA actually overestimates it with a value of  $38.5 \pm 2.2 \text{ m s}^{-1}$ , and NEWA underestimates it at  $34.2 \pm 1.6 \text{ m s}^{-1}$ .

Thus, the difference between the three models is small compared to the impact of time averaging. A more systematic investigation of the 50-year extreme as function of the spatio-temporal characteristics of the underlying data could provide the additional information required to obtain reliable estimates of wind extremes from the reanalysis data. Coupling between weather models and large-eddy simulations has received much attention lately, and is rapidly becoming the new status quo (e.g. Muñoz-Esparza *et al.*, 2014; Sanz Rodrigo *et al.*, 2017b; Hewitt *et al.*, 2018). These developments make such a study possible in the near future.

Yet, even if reliable data about past extremes including information about their validity is available, there is another potential pitfall that should be considered when applying extreme value theory. The fundamental assumption is that all extreme events in the data are drawn from the same parent distribution. The parent distribution, in this case, is the long-term wind climate at MMIJ. Using only four years of data is already pushing the limits of this assumption, for not all physical extremes that may occur within the current climate may be represented in the subset. Moreover, since the theory is used to make predictions far into the future, the assumption that the climate does not change may be violated. In 2017, Ophelia set a record for the eastern-most Atlantic major hurricane. Scientists at the Dutch national weather service warn that such storms, which can get considerably stronger than other types of storms in this area, may occur more often as the ocean warms up by climate change (Haarsma *et al.*, 2013; Baatsen *et al.*, 2015; Dekker *et al.*, 2018). Hence, present results on wind extremes only act as an illustration of their uncertainty, and interested readers are strongly advised to turn their attention to the dedicated literature referred.




**Figure 6.11:** Gumbel plots for extreme value analysis based on three reanalysis datasets as compared to observations from met mast IJmuiden. Return times are shown on the top axis. Shaded regions represent uncertainty estimates of the mean plus or minus one standard deviation based on a Monte Carlo procedure.

## 6.6 A spatial low-level jet climatology based on the Dutch Offshore Wind Atlas

Now that the performance of the three reanalysis datasets has been properly evaluated, it is instructive to briefly highlight their potential regarding spatial analysis of anomalous events. To illustrate the refinement achieved by downscaling, the ten-year mean low-level jet frequency over the DOWA domain is superimposed on the ERA-5 data (figure 6.12). Especially the detailed orographic structure in the southeast of the domain stands out, but also the coastal morphology is represented much more truthfully.

Some striking features are revealed in figure 6.12. E.g., the concave eastern coastline of East Anglia clearly favours LLJ formation. A band of preferred LLJ occurrence which more or less follows the concave shape of this coastline appears. Furthermore, the impact of the Dover Strait is clearly visible in the climatology. Especially the easternmost extremity of Kent leaves a LLJ ‘wake’ towards the northeast. Some aspects of this specific jet are discussed in Capon (2003).

A fixed fall-off threshold of  $2 \text{ m s}^{-1}$  was used to produce figure 6.12. Alternatively, it would be possible to map the mean, median or 95-percentile of the absolute fall-off. That would provide additional information about the spatial distribution of LLJ characteristics. Moreover, this approach can be used for other anomalous events as well. A comprehensive spatial analysis of a variety of anomalous events, however, is left for future work.



**Figure 6.12:** Spatial climatology of low-level jets up to a height of 600 m as represented in the Dutch Offshore Wind Atlas, superimposed upon the same figure based on ERA-5 data.

## 6.7 Summary and conclusions

This study evaluates three state-of-the-art reanalysis datasets used in the wind energy industry, i.e. ERA-5, DOWA and NEWA against four years of high-quality wind profile observations over the North Sea. Exceptional performance was found for the Dutch Offshore Wind Atlas, which was nearly unbiased in terms of wind speed. ERA-5 demonstrates comparable root mean square errors ( $\sim 1.4 \text{ m s}^{-1}$ ), but it generally underestimates the wind speed, probably due to the smoothing effects due to its relatively coarse resolution. NEWA, despite its increased resolution, does not improve upon ERA-5, which seems to inherit the wind underestimation from ERA-5, and an increase of the random errors suggest that the model is considerably more sensitive to the double-penalty problem. The fact that DOWA performs much better in this respect, even though its resolution is comparable to NEWA, illustrates the impact of the modelling strategy and additional data-assimilation.

The wind in the reanalyses is typically veered with respect to the observations, and this veering increases with height. ERA-5 and DOWA performed very similar, while NEWA again exhibited a wider range of wind direction errors. Also a potential pitfall in using summary statistics for the evaluation of wind direction was illustrated.

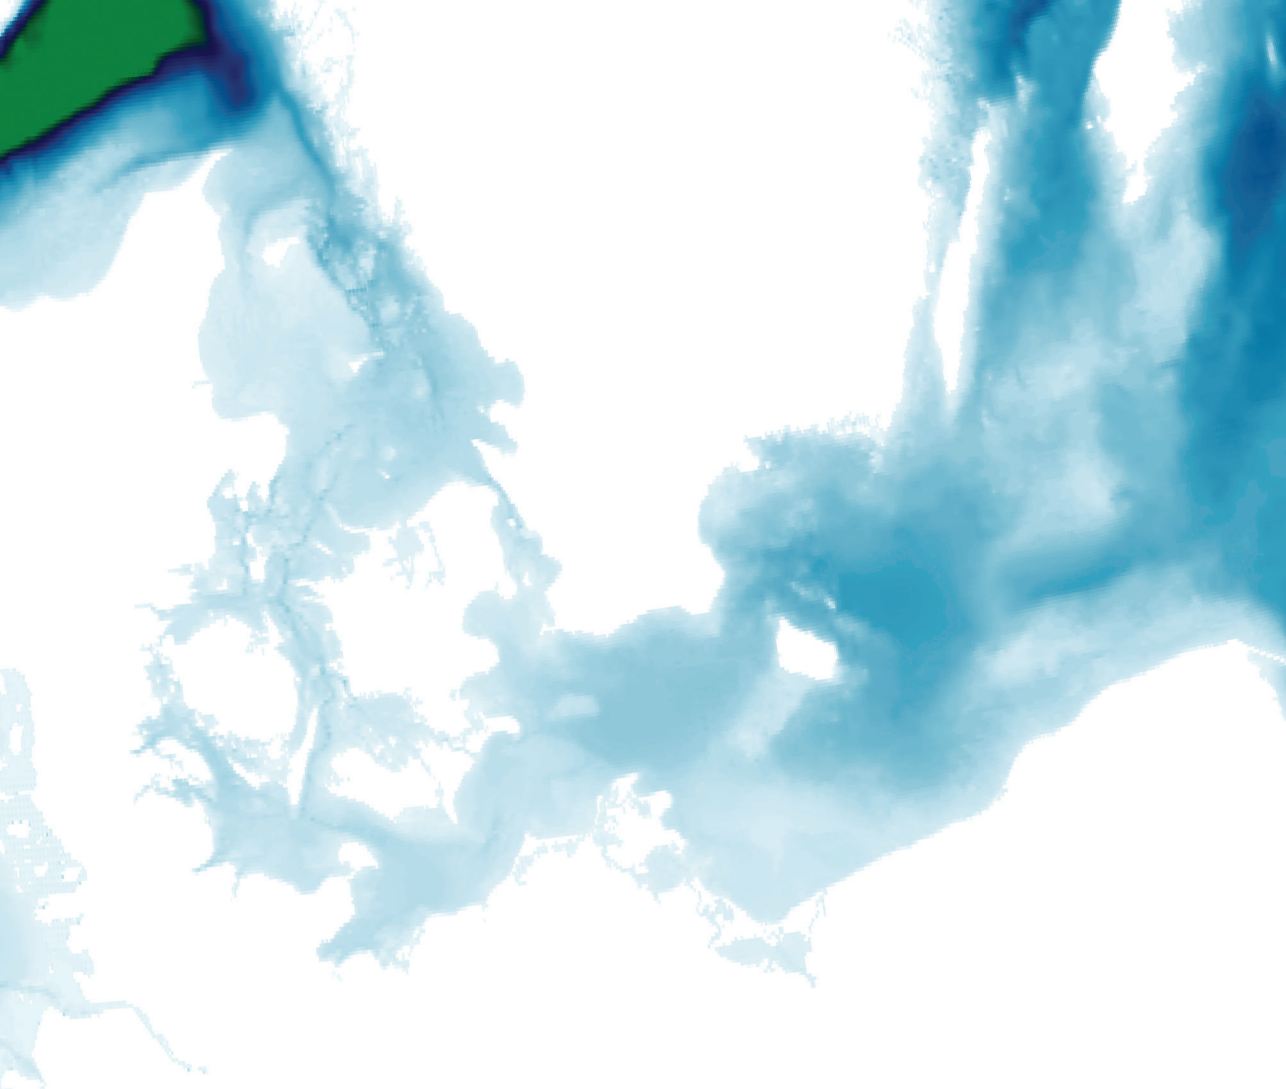
The reanalyses' representation of anomalous wind events was evaluated as well. Generally, the relatively high-resolution models are able to represent more fine-scale structures, but this comes at the cost of considerable mismatches in timing of the events. For low-level jets, DOWA outperforms the two other datasets. In a climatological sense, wind ramps are best represented in NEWA, but one-to-one correspondence is slightly better in DOWA. Extreme wind shear is best represented by the higher resolution models, though they still underestimate the vertical wind shear, which has been linked to deficiencies in the representation of stable conditions. The representation of lateral shear, or wind veer is very poor in all datasets. For wind extremes, the differences between the models are nullified by the uncertainties related to spatiotemporal characteristics of the underlying data and about changes in future climate.

Finally, a climatological map of low-level jet frequency based on the DOWA data was briefly discussed. Compared to ERA-5, the enhanced resolution reveals much more details of the low-level jet climatology: orographic features can clearly be distinguished. The high-resolution data reveal many interesting aspects of the low-level jet climatology such as the role of orography and coastal effects. This opens up a wealth of possibilities for further investigations, and it is advised that climatological maps of anomalous events are incorporated in future standards.



# CHAPTER 7





## GENERAL DISCUSSION: FROM IDEALIZED INFLOW FIELDS TO EMBEDDED EMULATIONS

---

WE ARE LONG WELL AWARE OF THE RISKS OF BELCHING OUT CO<sub>2</sub>. BUT WE SEEM PETRIFIED AS RABBITS CAUGHT IN THE HEADLIGHTS. RISKS ALONE WILL NOT GET US INTO ACTION. WE HAVE TO REFRAME. WE HAVE TO EXCHANGE THE FRAME OF RISK FOR THE FRAME OF OPPORTUNITY. AND FOR THIS WE NEED NEW 'IMAGINARIES'. IMAGES OF A FUTURE THAT CAN WORK. THAT GIVES US PERSPECTIVE AND HOPE.

VOICE OVER TEXT FROM: 2050 - AN ENERGETIC ODYSSEY



## **Em•u•late**

*tr.v.* **em•u•lat•ed, em•u•lat•ing, em•u•lates**

1. To strive to equal or excel, especially through imitation: *an older pupil whose accomplishments and style I emulated.*
2. To compete with successfully; approach or attain equality with.
3. *Computers* To imitate the function of (another system), as by modifications to hardware or software that allow the imitating system to accept the same data, execute the same programs, and achieve the same results as the imitated system.

Source: <https://www.thefreedictionary.com/emulation>, last accessed June 23, 2019

## 7.1 Introduction: a bird's-eye view

This thesis deals with the representation of offshore winds in energy applications. Uncertainty, as explained in chapter 1, enters estimates of annual energy production and turbine lifetime through imperfect knowledge of the (wind) climate, and through the way in which these source data are processed. Uncertainty quantification in the source data is straightforward: state-of-the-art reanalysis data in the Dutch Offshore Wind Atlas are nearly unbiased, but limited precision still introduces a root mean square error of about  $1.5 \text{ m s}^{-1}$  (chapter 6). Observations are limited mostly by their temporal coverage, which may introduce uncertainties of a similar magnitude (chapter 4). Despite recent advancements, conditions characterized by stable stratification still form the biggest challenge for weather models (chapters 3, 6). Consequently, wind shear, veer, and turbulence are not well captured in reanalysis data. Wind extremes are often underestimated and wind direction is typically veered with respect to observations.

While this answer is very clear-cut, it is not particularly satisfactory (certainly not for a Philosophy Doctorate). Further interpretation of the results requires a reflection on the second source of uncertainty: the way in which the source data are processed. In chapter 1, I already introduced the basics of the engineering models used for power calculations, wind farm optimization and load assessment. However, this flow modelling landscape is evolving rapidly and consequently, the lessons from this thesis are most valuable in the context of a *status futururus*. The following question adequately serves this purpose: *How may the lessons from this thesis contribute to the development of next-generation load- and power assessment procedures?*

## 7.2 Building a digital wind farm

Sanderse *et al.* (2011), and more recently Hewitt *et al.* (2018, bearing the same title as this section) provide an overview of the various (types of) models used for wind energy applications (also see figure 1.3). Generally, a trade-off exists between the level of geometric detail and the size of the model domain. On one end of the spectrum, high-resolution flow simulations very accurately describe the interaction between the fluid and the structure, but their spatial and temporal extent is constrained. On the other side, coarser models simulate the flow through an entire wind farm, but the level of detail is limited. Wind turbines in these models are represented on an abstract, or idealized level; they are *parametrized*. Advances in computer power and flow modelling are continuously bringing these two ends of the simulation spectrum closer together. Parametrizations in the coarser models are being improved by learning from the detailed simulations, and detailed simulations in turn are fed with more sophisticated/realistic inflow from the coarser models (e.g. Englberger and Dörnbrack, 2018). Extrapolating this trend into the future suggests the wind energy sector will increasingly realize the vision of true digital twins (Boschert and Rosen, 2016): virtual replicas of physical objects or systems that emulate the functioning of their real counterparts, to be used for advanced planning, monitoring, and real-time optimization.

A study that approaches this functionality may be found in Kanev *et al.* (2018). They used a three-dimensional model (parametrized turbulence, no temperature equation) to simulate the flow through a wind farm for all wind directions and a range of yaw angles to influence the direction of the wake. The simulated wind profiles at each turbine were compared with a pre-compiled load table. This set-up allowed for a mutual optimization of annual energy production

and turbine lifetime through well-informed active wake control. While still far from a digital replica, this study serves to show the direction in which research is headed. Obviously, future simulation systems can benefit from more realistic inflow fields. In fact, meteorological data sources may be integrated into the model chain.

Large-eddy simulations (LES) constitute a natural interface between meteorological input data and engineering models. Much effort has recently been invested in dynamic coupling between ‘conventional’ weather models and finer-scale large-eddy simulations, and in improving their efficiency (Muñoz-Esparza *et al.*, 2014; Schalkwijk *et al.*, 2015b,a; Heinze *et al.*, 2017; Sanz Rodrigo *et al.*, 2017b). Essentially, this enables wind farm simulations to be embedded within the full complexity of the regional or even global atmospheric circulation. Indeed, LES-downscaled wind and turbulence data for designated wind farm areas is one of the deliverables of the Dutch Offshore Wind Atlas project<sup>1</sup>. Along with this embedded view on wind farm simulations comes a different perspective on inflow specification.


### 7.3 Incentives for a paradigm shift

Contemporary procedures to estimate load and power stand out for their *local* perspective, which I like to characterize as a ‘bottom-up approach’ (figure 7.1). The interaction between wind and turbine(s) is of interest, and outside influences are preferably neglected or described statistically in a highly idealized fashion, which is reflected in terminology like ‘external conditions’, ‘inflow fields’, and ‘siting parameters’. One argument for a simplified representation of the wind climate is that commonly used engineering models cannot process more sophisticated inflow, for example in the case of two-dimensional wake models (Annoni *et al.*, 2018). The other reason is that a statistical representation of the wind climate in the form of a Weibull distribution or wind rose facilitates time extrapolation, substantially reducing the number of simulations required to estimate annual or lifetime estimates of power and fatigue loads.

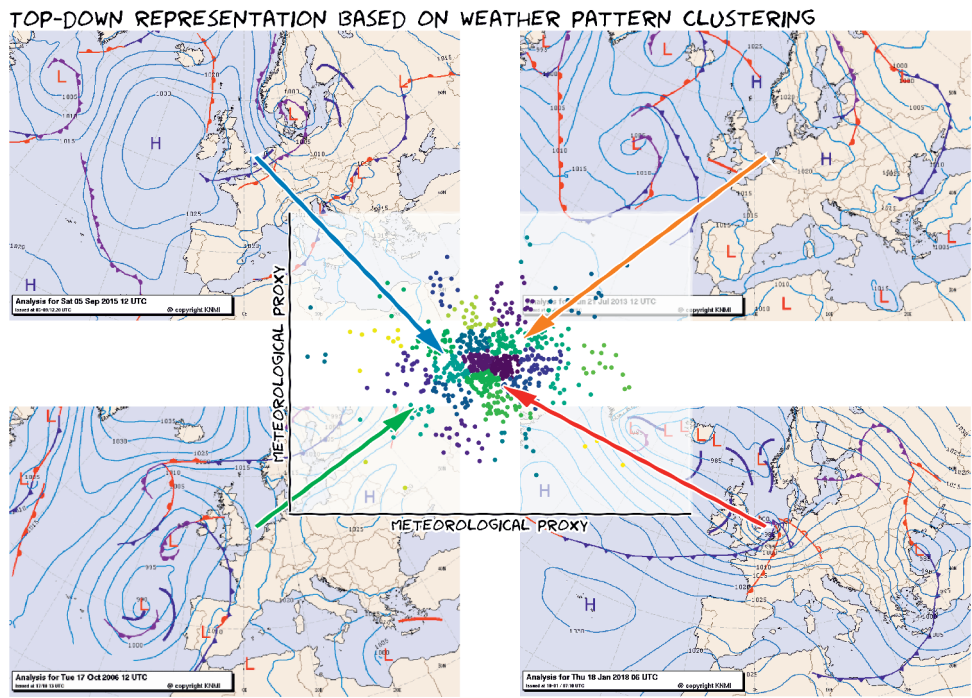
In chapter 1, I suggested that oversimplification may introduce uncertainties or inaccuracies. For the same wind speed at hub height, a low-level jet wind profile, for example, exhibits completely different shear and turbulence characteristics than a logarithmic wind profile. To obtain a first estimate of the impact of these ‘spherical cows’, I introduced the concept of anomalous wind events. The results presented in chapters 2, 4, and 6 demonstrate that especially low-level jets occur frequently over the North Sea, which invalidates the conceptual load assessment strategy illustrated in figure 7.1. Is it possible to conceive of a set of summary statistics that does justice to the complex physical nature of coastal flows, and that can be used for ‘inflow’ specification? As argued in chapter 5, these new siting parameters should at least include the joint probability of thermal effects, specifically density stratification and baroclinity, in conjunction with the ‘background’ wind. That also calls for methods or models that can convert these parameters to realistic inflow fields for subsequent simulations down the model chain. Englberger and Dörnbrack (2018) present a first step in that direction.

The availability of high-quality downscaled reanalysis datasets eventually begs the question whether summary statistics are actually still needed. While vertical profiles of wind speed itself are abundantly available, vertical profiles of the statistical Weibull parameters are reported in literature, parametrized and even accredited with physical meaning (e.g. Gryning *et al.*, 2016); doesn’t that defeat their purpose?

<sup>1</sup>dutchoffshorewindatlas.nl, last accessed June 24, 2019



**Figure 7.1:** Conceptual illustration of the bottom-up approach, emphasizing the use of a Weibull distribution in conjunction with traditional wind profile parametrization and synthetic turbulence to generate inflow fields for load simulations.



**Figure 7.2:** Illustration of the top-down approach, in which weather situation with similar large-scale structure are clustered together, because for similar large-scale situations, the local manifestation of the flow is expected to be similar as well.

One might argue that Weibull parameters are still useful for time extrapolation, but the embedded view on wind farm simulations as downscaled LES from weather models or reanalysis data inspires another perspective. Rather than describing the climatology in terms of *local* statistics, one might alternatively try to classify the *large-scale* situation in terms of characteristic weather patterns (figure 7.2). This ‘top-down’ perspective is reflected in the breadth of literature on synoptic classification, weather types or weather patterns, and ‘typical meteorological years’ (Philipp *et al.*, 2010). The basic premise is that similar weather patterns on a large scale yield similar manifestations of the turbulent wind field locally. In chapter 3, I already alluded to this idea for the selection of 30 more or less representative cases. I am still intrigued by the idea that the local weather can be described in terms of a small number of ‘modes’. Jiménez *et al.* (2009), for example, successfully related six dominant spatial wind patterns over Navarra (a province in Northern Spain) to the large-scale circulation and thermodynamic fields.

The top-down approach combines well with the characteristics of coastal flows revealed in chapter 5. The presence of an overlying high-pressure area, for example, favours clear skies and hence the build-up of strong horizontal temperature gradients. Moreover, high-pressure areas are usually characterised by relatively weak background winds, which facilitates the manifestation of local circulations. The exact positioning of the high-pressure area determines the background-wind direction and hence the interaction of local circulations with the background wind (figure 5.3). Thus, provided that the model chain is sophisticated enough, phenomena such as low-level jets will spontaneously emerge whenever a favourable weather situation presents itself. Obviously, it is not that simple in practice, but one may regard the perspective as an ideal and explore what consequences it entails.

## 7.4 Towards data-driven case selection

Is it not necessary, then, to bother about anomalous events within the top-down approach at all? Certainly it is! First of all, it is important to understand under which circumstances anomalous events may be expected. Chapter 5 stresses the importance of thermal effects in coastal meteorology. A case selection method (or, equivalently, weather type/synoptic classification) that would not include temperature in any form may therefore not yield a truly representative subset. Moreover, the atmospheric response to the large-scale forcing may be highly non-linear. A strong synoptic system tends to suppress or nullify the impact of more subtle features and small-scale processes, while a weak synoptic system leaves room for a whole range of subtle mesoscale phenomena. With an eye on uncertainty reduction, this calls for a denser sampling of such permitting flow regimes. Simple as it seems, this last statement constitutes a whole new paradigm, which I’ll refer to as *data-driven case selection*.

The data-driven approach starts with a limited number of cases, or inputs, evaluates the model (chain) and then analyses the output. Subsequently, new inputs are selected and the procedure is repeated. The selection of new inputs may be based on various factors: evaluation with available ‘ground truth’ observations; using information about the uncertainty of the inputs; or identifying regions where the response to the input is (highly) non-linear. Note that evaluation may happen at various levels of the model chain to add uncertainty information, but selection of new inputs is based on the final outcome of the model chain. For example, evaluation of weather models with wind observations can yield valuable information about the uncertainty of the inputs to the subsequent models in the chain. This exercise may yield a high degree of un-

certainty in wind direction for wind speeds below  $3 \text{ m s}^{-1}$ . However, if the quantity of interest is power, it is not worthwhile to add additional wind farm simulations, for the wind turbines won't operate under these conditions. In other words, the uncertainty does not propagate to the final quantity of interest.

Perhaps a data-driven methodology will finally obviate the concept of anomalous events, for the case that's most anomalous will by definition be the next addressed. This however stretches the definition of anomalous events from cases that introduce uncertainty because common assumptions are violated to cases that are relatively uncertain for any reason. That brings me to another reason why vigilance is required with the top-down approach.

## 7.5 A note on model performance

So far, I have just assumed that the accuracy of observations and model data is known and can be used to propagate uncertainty information through the model chain. Indeed, several error statistics have been reported in the previous chapters that may be used for such ends. These statistics should be interpreted with caution, though. Just like logarithmic wind profiles, instantaneous model fields can deviate substantially even if the average wind profile is well captured (figure 1.7). Therefore, model validation should not only focus on overall error statistics, but especially on the representation of extreme or anomalous events. Chapters 4 and 6 demonstrate novel methods to account for such events in the evaluation of model performance.

Advances in atmospheric modelling are pushing (global) weather models to ever finer scales. The European Centre for Medium-Range Weather Forecasts (ECMWF) has set the ambitious target of running a 5-km ensemble forecast by 2025 (Adrian, 2015). In chapter 3, model resolution was apparently not a decisive factor. Nevertheless, chapter 6 shows that the higher resolution models were able to better capture the climatology of anomalous events, and wind speed was underestimated in ERA5. This might still be a result of the relatively coarse resolution<sup>2</sup>. Alternatively, as suggested in chapter 3, the hydrostatic assumption might produce a similar damping effect. Finally, the surface roughness may play a role.

The increase in resolution envisioned by ECMWF will violate the hydrostatic assumption anyway, and therefore the centre is also investigating the use of non-hydrostatic solvers (Adrian, 2015). As shown in chapter 6, the downside of a higher-resolution or non-hydrostatic model is that it suffers more from the double-penalty problem, and therefore the precision may decrease, despite improvements in the physical representation of the flow. Chapter 6 clearly demonstrated that (additional) data-assimilation can reduce this effect. The downscaling strategy applied for the Dutch Offshore Wind Atlas proved very successful, and may be used as an example for future (downscaling) studies.

Despite recent efforts, weather models still struggle to adequately represent the vertical structure of the boundary layer, especially in stably stratified conditions (Holtslag *et al.*, 2013; Sandu *et al.*, 2013; Riddaway, 2014; Bengtsson *et al.*, 2017; Olauson, 2018, chapters 3, 6). Importantly, wind turning across the boundary layer is underestimated, leading to a veered model wind with respect to observations (Brown *et al.*, 2005; Svensson and Holtslag, 2009; Sandu *et al.*, 2013, chapters 3, 6). It is in the interest of both the meteorology and the wind energy communities

<sup>2</sup>Note that the operational IFS, as evaluated in chapter 3, runs on a higher resolution ( $\sim 15 \text{ km}$ ) than what was used for the production of ERA5 ( $\sim 30 \text{ km}$ ).

to equip offshore measurement platforms with additional instrumentation to measure density stratification and turbulent fluxes of heat and momentum. The data so obtained may be used for more in-depth understanding of the structure of the (stable) offshore boundary layer, for model verification, and eventually model improvement.

## 7.6 Establishing the status futurus

Undoubtedly, the development towards a *status futurus* will see many new challenges, and building up experience is crucial for their timely identification. As the first results from climatological large-eddy simulations become available, it is imperative that these data are thoroughly validated. Chapters 3, 4 and 6 provide a starting point for validation studies. Most importantly, evaluation should focus not only on summary statistics, but care should be taken to also assess the representation of anomalous events.

The list of anomalous events described in this thesis is certainly not exhaustive. Other events may be defined and investigated following a similar approach. Besides, this thesis focused mostly on low-level jets and I've started to classify them according to their physical governing mechanism. An analogous exercise may be performed for extremes of wind shear, wind veer, and wind ramps to obtain a set of physical events related to rapid fluctuations in wind speed or direction, or extreme shear.

In fact, this is already happening today. Cold air outbreaks have been shown to exhibit elongated horizontal structure with wind speed differences of  $10 \text{ m s}^{-1}$  on spatial scales of 10 km (de Roode *et al.*, 2019). These could lead to wind ramps much stronger than seen in this thesis. Furthermore, Hawbecker *et al.* (2017) studied the characteristics of a thunderstorm downbursts. With a typical time span of 10 to 15 minutes, these may be recognized in the ten-minute observations and thus the ramp climatology in chapter 2. Due to their distinct physical structure, they might also show up as low-level jets.

Downburst, however, may also be identified as a physical mechanisms responsible for extreme gusts. Contemporary guidelines for wind turbine design distinguish between ultimate loads and fatigue loads (IEC, 2005, 2009; Burton *et al.*, 2011). Extreme gusts are an example of the first: they may cause abrupt failure but its contribution to the accumulation of fatigue loads is probably not substantial. The opposite is true for low-level jets. In either case, it is relevant to define, characterize and understand these events, in order to design a suitable load case. Eventually, a complete set of reference or benchmark cases of physical events should come to replace the load cases in current design standards.

How could these physical load cases be represented within the downscaling framework discussed in this chapter? In 2014, Park *et al.* have already shown that it is possible to derive inflow fields for load simulations from LES of a stable boundary layer including a low-level jet. At the time, it was a proof of concept. Recently Hawbecker *et al.* (2018) derived a similar set of inflow fields from large-eddy simulations of a thunderstorm downburst, and they plan to use these for additional load simulations. Perhaps Baas and van den Brink (2013) could provide the starting point for downscaling a set of independent storms. This is a challenging endeavour, for it would require a rather large domain, or a very adequate coupling to the larger scales. Successful results, however, as discussed in chapter 6, will contribute substantially to under-

standing the relation between the spatio-temporal characteristics of the simulations and the corresponding wind extremes.

A procedure along the lines of Englberger and Dörnbrack (2018) could be used to generate LES-derived inflow fields for more ordinary conditions. They performed an LES for a complete diurnal cycle to cover a wide range of atmospheric conditions. However, chapter 5 (and also e.g. Floors *et al.*, 2015) suggests that in coastal areas the thermal contrast between land and sea might be just as important as the diurnal cycle. Thus, the range of conditions should be further extended by introducing that thermal contrast in the LES as well. One way to achieve this is to use a (stretched) domain with two different surface types and open boundaries in the cross-coast direction as in Antonelli and Rotunno (2007). Alternatively, an additional term to account for baroclinity could be added to the momentum equations as in Taylor and Ferrari (2010).

These are just a few steps in the direction of more realistic inflow fields for high-fidelity simulations. However, the status futurus also requires that the feedbacks of wind turbines and wind farms on the inflow fields themselves are understood and accounted for. Therefore, turbine or wind farm parametrizations will need to be implemented in weather models. Such parametrizations already exist (Fitch *et al.*, 2012; Volker *et al.*, 2015), and the Dutch national weather service is currently working on their implementation in operational forecasts<sup>3</sup>. A relevant research question that is just starting to be addressed is how large-scale implementation of off-shore wind farms feeds back to the atmosphere, and how this impacts on (regional) weather and climate.

## 7.7 Towards an energetic odyssey

*2050 - An Energetic Odyssey* is a fascinating animation produced by the IABR (International Architecture Biennale Rotterdam) to provide a concrete picture of what the energy transition entails. Focussing on the North Sea, it sketches the image of 25 000 10 MW wind turbines, connecting the electricity grids of the surrounding countries. Although it is an ambitious challenge, seeing all minor contributions blend into that bigger picture is first and foremost deeply inspiring.

As an extension of our imagination, model simulations play a pivotal role in the establishment and realization of such a sustainable vision. The output of realistic simulations is, however, prone to hasty interpretation unless the models are thoroughly validated and their limitations are properly understood. By providing an updated validation and a critical reflection on contemporary modelling approaches, this thesis contributes to that overall objective of meeting the human rights of every person within the means of our life-giving planet, with which I started the general introduction in chapter 1.

<sup>3</sup>“The main innovations in the [DOWA] project are [...] to include wind farm (cluster) wake effects in mesoscale weather model HARMONIE.” <https://www.knmi.nl/research/observations-data-technology/projects/dowa>, accessed July 8, 2019





# Bibliography

- Ackermann, T. and Söder, L. (2002). An overview of wind energy-status 2002. *Renewable and Sustainable Energy Reviews*, 6(1):67–127.
- Adams, E. (1997). Four ways to win the sea breeze game. *Sailing World*, 3:44.
- Adrian, G. (2015). Ecmwf strategy 2016-2025 - the strength of a common goal.
- Andreas, E. L., Claffy, K. J., and Makshtas, A. P. (2000). Low-Level Atmospheric Jets And Inversions Over The Western Weddell Sea. *Boundary-Layer Meteorology*, 97(3):459–486.
- Angevine, W. M., Tjernström, M., and Žagar, M. (2006). Modeling of the coastal boundary layer and pollutant transport in new england. *Journal of applied meteorology and climatology*, 45(1):137–154.
- Annoni, J., Fleming, P., Scholbrock, A., Roadman, J., Dana, S., Adcock, C., Porte-Agel, F., Raach, S., Haizmann, F., and Schlipf, D. (2018). Analysis of control-oriented wake modeling tools using lidar field results. *Wind Energy Science*, 3(2):819–831.
- Antonelli, M. and Rotunno, R. (2007). Large-eddy simulation of the onset of the sea breeze. *Journal of the Atmospheric Sciences*, 64(12):4445–4457.
- Archer, C. L. and Jacobson, M. Z. (2005). Evaluation of global wind power. *Journal of Geophysical Research: Atmospheres*, 110(D12).
- Baas, P., Bosveld, F. C., and Burgers, G. (2016). The impact of atmospheric stability on the near-surface wind over sea in storm conditions. *Wind Energy*, 19(2):187–198.
- Baas, P., Bosveld, F. C., Klein Baltink, H., and Holtslag, A. A. M. (2009). A Climatology of Nocturnal Low-Level Jets at Cabauw. *Journal of Applied Meteorology and Climatology*, 48(8):1627–1642.
- Baas, P. and van den Brink, H. W. (2013). Evaluation of harmonie simulations for 16 historical storms. *Contribution to the WP1 of the WT12017\_HB Wind Modelling Project*.
- Baatsen, M., Haarsma, R. J., Van Delden, A. J., and De Vries, H. (2015). Severe autumn storms in future western europe with a warmer atlantic ocean. *Climate dynamics*, 45(3-4):949–964.
- Banakh, V. A., Smalikho, I. N., Köpp, F., and Werner, C. (1995). Representativeness of wind measurements with a cw doppler lidar in the atmospheric boundary layer. *Applied Optics*, 34(12):2055–2067.
- Bardal, L. M. and Sætran, L. R. (2016). Wind Gust Factors in a Coastal Wind Climate. *Energy Procedia*, 94:417–424.
- Barstad, I. and Grønås, S. (2005). Southwesterly flows over southern norway – mesoscale sensitivity to large-scale wind direction and speed. *Tellus A*, 57(2):136–152.
- Bauer, P., Thorpe, A., and Brunet, G. (2015). The quiet revolution of numerical weather prediction. *Nature*, 525(7567):47.
- Bengtsson, L., Andrae, U., Aspelien, T., Batrak, Y., Calvo, J., *et al.* (2017). The HARMONIE-AROME Model Configuration in the ALADIN-HIRLAM NWP System. *Monthly Weather Review*, 145(5):1919–1935.
- Benjamin, S. G., Weygandt, S. S., Brown, J. M., Hu, M., Alexander, C. R., *et al.* (2016). A north american hourly assimilation and model forecast cycle: The rapid refresh. *Monthly Weather Review*, 144(4):1669–1694.
- Bhaganagar, K. and Debnath, M. (2014). Implications of stably stratified atmospheric boundary layer turbulence on the near-wake structure of wind turbines. *Energies*, 7(9):5740–5763.
- Blackadar, A. K. (1957). Boundary layer wind maxima and their significance for the growth of nocturnal inversions. *Bulletin of the American Meteorological Society*, 38(5):283–290.
- Bollmeyer, C., Keller, J. D., Ohlwein, C., Wahl, S., Crewell, S., *et al.* (2015). Towards a high-resolution regional reanalysis for the european cordex domain. *Quarterly Journal of the Royal Meteorological Society*, 141(686):1–15.
- Boschert, S. and Rosen, R. (2016). Digital twin – the simulation aspect. In *Mechatronic Futures*, pp. 59–74. Springer, Cham (CH).

- Breton, S.-P. and Moe, G. (2009). Status, plans and technologies for offshore wind turbines in Europe and North America. *Renewable Energy*, 34(3):646–654.
- Brown, A. R., Beljaars, A. C. M., Hersbach, H., Hollingsworth, A., Miller, M., and Vasiljevic, D. (2005). Wind turning across the marine atmospheric boundary layer. *Quarterly Journal of the Royal Meteorological Society*, 131(607):1233–1250.
- Burk, S. D. and Thompson, W. T. (1996). The Summertime Low-Level Jet and Marine Boundary Layer Structure along the California Coast. *Monthly Weather Review*, 124(4):668–686.
- Burton, T., Jenkins, N., Sharpe, D., and Bossanyi, E. (2011). *Wind Energy Handbook, Second Edition*. John Wiley & Sons, Ltd, Chichester, UK.
- Butz, C., Liechti, J., Bodin, J., and Cornell, S. E. (2018). Towards defining an environmental investment universe within planetary boundaries. *Sustainability Science*, 13(4):1031–1044.
- C3S (2017). ERA5: Fifth generation of ECMWF atmospheric reanalyses of the global climate. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). <https://cds.climate.copernicus.eu/cds-app#!/home>.
- Capon, R. A. (2003). Wind speed-up in the dover straits with the met office new dynamics model. *Meteorological Applications*, 10(3):229–237.
- Carroll, J., McDonald, A., Dinwoodie, I., McMillan, D., Revie, M., and Lazakis, I. (2016). Availability, operation and maintenance costs of offshore wind turbines with different drive train configurations. *Wind Energy*, 20(2).
- Carta, J. A., Velázquez, S., and Cabrera, P. (2013). A review of measure-correlate-predict (mcp) methods used to estimate long-term wind characteristics at a target site. *Renewable and Sustainable Energy Reviews*, 27:362 – 400.
- Carter, G. M., Dallavalle, J. P., and Glahn, H. R. (1989). Statistical forecasts based on the national meteorological center’s numerical weather prediction system. *Weather and Forecasting*, 4(3):401–412.
- Carvalho, D., Rocha, A., Gómez-Gesteira, M., and Santos, C. S. (2014). Offshore wind energy resource simulation forced by different reanalyses: Comparison with observed data in the Iberian Peninsula. *Applied Energy*, 134:57 – 64.
- Christakos, K., Cheliotis, I., Varlas, G., and Steeneveld, G.-J. (2016). Offshore Wind Energy Analysis of Cyclone Xaver over North Europe. *Energy Procedia*, 94:37–44.
- Christakos, K., Varlas, G., Reuder, J., Katsafados, P., and Papadopoulos, A. (2014). Analysis of a low-level coastal jet off the western coast of norway. *Energy Procedia*, 53:162 – 172. EERA DeepWind’ 2014, 11th Deep Sea Offshore Wind R&D Conference.
- Coelingh, J. P., van Wijk, A. J. M., Cleijne, J. W., and Pleune, R. (1992). Description of the North Sea wind climate for wind energy applications. *Journal of Wind Engineering and Industrial Aerodynamics*, 39(1):221–232.
- Coelingh, J. P., van Wijk, A. J. M., and Holtslag, A. A. M. (1996). Analysis of wind speed observations over the North Sea. *Journal of Wind Engineering and Industrial Aerodynamics*, 61(1):51–69.
- Colle, B. A. and Novak, D. R. (2010). The new york bight jet: climatology and dynamical evolution. *Monthly Weather Review*, 138(6):2385–2404.
- Colmenar-Santos, A., Perera-Perez, J., Borge-Diez, D., and dePalacio Rodríguez, C. (2016). Offshore wind energy: A review of the current status, challenges and future development in spain. *Renewable and Sustainable Energy Reviews*, 64:1 – 18.
- Cook, N. J. N. J. (1985). The designer’s guide to wind loading of building structures / Pt.1., Background, damage survey, wind data and structural classification. Structures. Wind loads. Effects. - For design (BNB/PRECIS).
- Couto, A., Silva, J., Costa, P., Santos, D., Simões, T., and Estanqueiro, A. (2019). Towards a high-resolution offshore wind atlas - the portuguese case.  
**URL:** [https://www.sintef.no/globalassets/project/eera-deepwind-2019/presentations/c2\\_teresasimoes\\_vf.pdf](https://www.sintef.no/globalassets/project/eera-deepwind-2019/presentations/c2_teresasimoes_vf.pdf)

- de Roode, S. R., Frederikse, T., Siebesma, A. P., Ackerman, A. S., Chylik, J., *et al.* (2019). Turbulent transport in the gray zone: A large eddy model intercomparison study of the constrain cold air outbreak case. *Journal of Advances in Modeling Earth Systems*, 11(3):597–623.
- Dearing, J. A., Wang, R., Zhang, K., Dyke, J. G., Haberl, H., *et al.* (2014). Safe and just operating spaces for regional social-ecological systems. *Global Environmental Change*, 28:227–238.
- Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., *et al.* (2011). The era-interim reanalysis: configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society*, 137(656):553–597.
- Dekker, M. M., Haarsma, R. J., de Vries, H., Baatsen, M., and van Delden, A. J. (2018). Characteristics and development of european cyclones with tropical origin in reanalysis data. *Climate dynamics*, 50(1-2):445–455.
- Desmond, C., Murphy, J., Blonk, L., and Haans, W. (2016). Description of an 8 MW reference wind turbine. *Journal of Physics: Conference Series*, 753(9):092013.
- Dodge, D. M. (2015). Wind Power—An Illustrated History of its Development.  
**URL:** <http://www.telosnet.com/wind/>
- Donkers, J., Brand, A., and Eecen, P. (2011). Offshore Wind Atlas of the Dutch part of the North Sea. *EWEA 2011 Annual Event, Brussels*.
- Dörenkämper, M., Optis, M., Monahan, A., and Steinfeld, G. (2015). On the offshore advection of boundary-layer structures and the influence on offshore wind conditions. *Boundary-Layer Meteorology*, 155(3):459–482.
- Dorninger, C., Abson, D. J., Fischer, J., and Wehrden, H. v. (2017). Assessing sustainable biophysical human–nature connectedness at regional scales. *Environmental Research Letters*, 12(5):055001.
- Dotzek, N., Emeis, S., Lefebvre, C., and Gerpott, J. (2010). Waterspouts over the North and Baltic Seas: Observations and climatology, predication and reporting. *Meteorologische Zeitschrift*, 19(1):115–129.
- Doyle, J. D. and Shapiro, M. A. (1999). Flow response to large-scale topography: The greenland tip jet. *Tellus A*, 51(5):728–748.
- Draxl, C., Hahmann, A. N., Peña, A., and Giebel, G. (2014). Evaluating winds and vertical wind shear from weather research and forecasting model forecasts using seven planetary boundary layer schemes. *Wind Energy*, 17(1):39–55.
- Drobinski, P., Rotunno, R., and Dubos, T. (2011). Linear theory of the sea breeze in a thermal wind. *Quarterly Journal of the Royal Meteorological Society*, 137(659):1602–1609.
- Duine, G.-J., Jones, C., Carvalho, L. M., and Fovell, R. G. (2019). Simulating sundowner winds in coastal santa barbara: Model validation and sensitivity. *Atmosphere*, 10(3).
- Duncan, J. B., Marseille, G. J., and Wijnant, I. L. (2018a). Dowa validation against ascatsatellite winds. Technical Report R11649, TNO, Petten (NL).
- Duncan, J. B., van der Werff, P. A., and Bot, E. T. G. (2018b). Understanding of the offshore wind resource up to high altitudes ( $\leq 315$  m). Technical Report R11592, TNO, Petten (NL).
- ECMWF (2016). *IFS Documentation – Cy41r2*, chapter II: Data Assimilation. ECMWF.
- Eggels, A. W., Crommelin, D. T., and Witteveen, J. A. S. (2018). Clustering-based collocation for uncertainty propagation with multivariate dependent inputs. *International Journal for Uncertainty Quantification*, 8(1). Particularly section 3.2.
- Englberger, A. and Dörnbrack, A. (2018). A numerically efficient parametrization of turbulent wind-turbine flows for different thermal stratifications. *Boundary-Layer Meteorology*, 169(3):505–536.
- EWEA (2013). *Where’s the money coming from? Financing offshore wind farms*. European Wind Energy Association.
- Fairall, C. W., Bradley, E. F., Hare, J. E., Grachev, A. A., and Edson, J. B. (2003). Bulk Parameterization of Air–Sea Fluxes: Updates and Verification for the COARE Algorithm. *Journal of Climate*, 16(4):571–591.
- Farrugia, R. N. (2003). The wind shear exponent in a Mediterranean island climate. *Renewable Energy*, 28(4):647–653.

- Fast, J. D. and McCorcle, M. D. (1990). A two-dimensional numerical sensitivity study of the great plains low-level jet. *Monthly Weather Review*, 118(1):151–164.
- Fitch, A. C., Olson, J. B., Lundquist, J. K., Dudhia, J., Gupta, A. K., Michalakes, J., and Barstad, I. (2012). Local and mesoscale impacts of wind farms as parameterized in a mesoscale nwp model. *Monthly Weather Review*, 140(9):3017–3038.
- Fleming, P. D. and Probert, S. D. (1984). The evolution of wind-turbines: An historical review. *Applied Energy*, 18(3):163–177.
- Floors, R., Peña, A., and Gryning, S.-E. (2015). The effect of baroclinicity on the wind in the planetary boundary layer. *Quarterly Journal of the Royal Meteorological Society*, 141(687):619–630.
- Foken, T. (2006). 50 Years of the Monin-Obukhov Similarity Theory. *Boundary-Layer Meteorol*, 119(3):431–447.
- Foley, A. M., Leahy, P. G., Marvuglia, A., and McKeogh, E. J. (2012). Current methods and advances in forecasting of wind power generation. *Renewable Energy*, 37(1):1 – 8.
- Fowler, L. D., Skamarock, W. C., Grell, G. A., Freitas, S. R., and Duda, M. G. (2016). Analyzing the Grell–Freitas Convection Scheme from Hydrostatic to Nonhydrostatic Scales within a Global Model. *Monthly Weather Review*, 144(6):2285–2306.
- Gallego, C., Cuerva, A., and Costa, A. (2014). Detecting and characterising ramp events in wind power time series. *Journal of Physics: Conference Series*, 555(1):012040.
- Gallego-Castillo, C., Cuerva-Tejero, A., and Lopez-Garcia, O. (2015). A review on the recent history of wind power ramp forecasting. *Renewable and Sustainable Energy Reviews*, 52:1148–1157.
- Garratt, J. R. (1977). Review of Drag Coefficients over Oceans and Continents. *Monthly Weather Review*, 105(7):915–929.
- Garratt, J. R. (1990). The internal boundary layer—a review. *Boundary-Layer Meteorology*, 50(1-4):171–203.
- Garratt, J. R. and Ryan, B. F. (1989). The structure of the stably stratified internal boundary layer in offshore flow over the sea. *Boundary-Layer Meteorology*, 47(1-4):17–40.
- Gaumont, M., Réthoré, P.-E., Ott, S., Peña, A., Bechmann, A., and Hansen, K. S. (2014). Evaluation of the wind direction uncertainty and its impact on wake modeling at the horns rev offshore wind farm. *Wind Energy*, 17(8):1169–1178.
- Gebauer, J. G., Fedorovich, E., and Shapiro, A. (2017). A 1d theoretical analysis of northerly low-level jets over the great plains. *Journal of the Atmospheric Sciences*, 74(10):3419–3431.
- Gebräad, P. M. O., Teeuwisse, F. W., Wingerden, J. W., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y. (2014). Wind plant power optimization through yaw control using a parametric model for wake effects—a cfd simulation study. *Wind Energy*, 19(1):95–114.
- Glahn, H. R. and Lowry, D. A. (1972). The use of model output statistics (mos) in objective weather forecasting. *Journal of Applied Meteorology*, 11(8):1203–1211.
- Gonzalez Rouco, F., García Bustamante, E., Hahmann, A. N., Karagili, I., Navarro, J., Tobias Olsen, B., Sile, T., and Witha, B. (2019). Report on uncertainty quantification (Deliverable D4.4).
- Grell, G. A. and Freitas, S. R. (2014). A scale and aerosol aware stochastic convective parameterization for weather and air quality modeling. *Atmospheric Chemistry and Physics*, 14(10):5233–5250.
- Grisogono, B. and Belušić, D. (2009). A review of recent advances in understanding the meso- and micro-scale properties of the severe bora wind. *Tellus A*, 61(1):1–16.
- Gryning, S.-E., Batchvarova, E., Brümmner, B., Jørgensen, H., and Larsen, S. (2007). On the extension of the wind profile over homogeneous terrain beyond the surface boundary layer. *Boundary-Layer Meteorol*, 124(2):251–268.
- Gryning, S.-E., Floors, R., Peña, A., Batchvarova, E., and Brümmner, B. (2016). Weibull wind-speed distribution parameters derived from a combination of wind-lidar and tall-mast measurements over land, coastal and marine sites. *Boundary-Layer Meteorology*, 159(2):329–348.

- Gutierrez, W., Araya, G., Kiliyanpilakkil, P., Ruiz-Columbie, A., Tutkun, M., and Castillo, L. (2016). Structural impact assessment of low level jets over wind turbines. *Journal of Renewable and Sustainable Energy*, 8(2).
- Gutierrez, W., Ruiz-Columbie, A., Tutkun, M., and Castillo, L. (2017). Impacts of the low-level jet's negative wind shear on the wind turbine. *Wind Energy Science*, 2(2):533–545.
- Haarsma, R. J., Hazeleger, W., Severijns, C., De Vries, H., Sterl, A., Bintanja, R., Van Oldenborgh, G. J., and van den Brink, H. W. (2013). More hurricanes to hit western europe due to global warming. *Geophysical Research Letters*, 40(9):1783–1788.
- Hahmann, A. N., Vincent, C. L., Peña, A., Lange, J., and Hasager, C. B. (2015). Wind climate estimation using wrf model output: method and model sensitivities over the sea. *International Journal of Climatology*, 35(12):3422–3439.
- Haiden, T., Bauer, P., Bidlot, J. R., Ferranti, L., Hewson, T., Prates, F., Richardson, D., and Vitart, F. (2014). Evaluation of ecmwf forecasts, including 2013-2014 upgrades. Technical Report TM742, ECMWF, Reading (UK).
- Hajer, M. A. and Pelzer, P. (2018). 2050—an energetic odyssey: Understanding ‘techniques of futuring’ in the transition towards renewable energy. *Energy Research & Social Science*, 44:222 – 231.
- Hansen, K. S., Barthelmie, R. J., Jensen, L. E., and Sommer, A. (2012). The impact of turbulence intensity and atmospheric stability on power deficits due to wind turbine wakes at Horns Rev wind farm. *Wind Energy*, 15(1):183–196.
- Harris, R. I. (1999). Improvements to the ‘Method of Independent Storms’. *Journal of Wind Engineering and Industrial Aerodynamics*, 80(1–2):1–30.
- Hawbecker, P., Basu, S., and Manuel, L. (2017). Realistic simulations of the july 1, 2011 severe wind event over the buffalo ridge wind farm. *Wind Energy*, 20(11):1803–1822.
- Hawbecker, P., Basu, S., and Manuel, L. (2018). Investigating the impact of atmospheric stability on thunderstorm outflow winds and turbulence. *Wind Energy Science*, 3(1).
- Heinze, R., Dipankar, A., Henken, C. C., Moseley, C., Sourdeval, O., et al. (2017). Large-eddy simulations over germany using icon: a comprehensive evaluation. *Quarterly Journal of the Royal Meteorological Society*, 143(702):69–100.
- Hess, P. and Brezowsky, H. (1969). *Katalog der grosswetterlagen Europas*. Deutscher Wetterdienst.
- Hewitt, S., Margetts, L., and Revell, A. (2018). Building a digital wind farm. *Archives of Computational Methods in Engineering*, 25(4):879–899.
- Hogan, R. J., O'Connor, E. J., and Illingworth, A. J. (2009). Verification of cloud-fraction forecasts. *Quarterly Journal of the Royal Meteorological Society*, 135(643):1494–1511.
- Högström, U. and Smedman-Högström, A.-S. (1984). The wind regime in coastal areas with special reference to results obtained from the Swedish wind energy program. In *Boundary Layer Structure*, pp. 351–373. Springer, Dordrecht (NL).
- Hoinka, K. P. and Castro, M. D. (2003). The iberian peninsula thermal low. *Quarterly Journal of the Royal Meteorological Society*, 129(590):1491–1511.
- Hollingsworth, A. (1994). Validation and diagnosis of atmospheric models. *Dynamics of Atmospheres and Oceans*, 20(3):227–246.
- Holton, J. R. (1967). The diurnal boundary layer wind oscillation above sloping terrain. *Tellus*, 19(2):200–205.
- Holtstlag, A. A. M. (1984). Estimates of diabatic wind speed profiles from near-surface weather observations. *Boundary-Layer Meteorol*, 29(3):225–250.
- Holtstlag, A. A. M., Svensson, G., Baas, P., Basu, S., Beare, B., et al. (2013). Stable atmospheric boundary layers and diurnal cycles: Challenges for weather and climate models. *Bulletin of the American Meteorological Society*, 94(11):1691–1706.
- Holtstlag, M., Bierbooms, W. A. A. M., and van Bussel, G. J. W. (2015). Validation of surface layer similarity theory to describe far offshore marine conditions in the Dutch North Sea in scope of wind energy research. *Journal of Wind Engineering and Industrial Aerodynamics*, 136:180 – 191.

- Holtslag, M. C., Bierbooms, W. A. A. M., and Bussel, G. J. W. (2016). Wind turbine fatigue loads as a function of atmospheric conditions offshore. *Wind Energy*, 19(10):1917–1932.
- IEC (2005). IEC 61400-1: Wind turbines - Part 1: Design requirements.
- IEC (2009). IEC 61400-3: Wind turbines - Part 3: Design requirements for offshore wind turbines.
- Irwin, J. S. (1979). A theoretical variation of the wind profile power-law exponent as a function of surface roughness and stability. *Atmospheric Environment* (1967), 13(1):191–194.
- Jammalamadaka, S. R. and SenGupta, A. (2001). *Topics in Circular Statistics*. World Scientific Press.
- Jiménez, P. A., González-Rouco, J. F., Montávez, J. P., García-Bustamante, E., and Navarro, J. (2009). Climatology of wind patterns in the northeast of the iberian peninsula. *International Journal of Climatology*, 29(4):501–525.
- Jones, P. D., Harpham, C., and Briffa, K. R. (2013). Lamb weather types derived from reanalysis products. *International Journal of Climatology*, 33(5):1129–1139.
- Jung, J. and Broadwater, R. P. (2014). Current status and future advances for wind speed and power forecasting. *Renewable and Sustainable Energy Reviews*, 31:762 – 777.
- Kalverla, P. C., Steeneveld, G.-J., Ronda, R. J., and Holtslag, A. A. M. (2017). An observational climatology of anomalous wind events at offshore meteoromast ijmuiden (north sea). *Journal of Wind Engineering and Industrial Aerodynamics*, 165:86 – 99.
- Kalverla, P. C., Duncan Jr., J. B., Steeneveld, G. J., and Holtslag, A. A. M. (2019a). Low-level jets over the north sea based on era5 and observations: together they do better. *Wind Energy Science*, 4(2):193–209.
- Kalverla, P. C., Steeneveld, G.-J., Ronda, R. J., and Holtslag, A. A. M. (2019b). Evaluation of three mainstream numerical weather prediction models with observations from meteorological mast ijmuiden at the north sea. *Wind Energy*, 22(1):34–48.
- Kanev, S. K., Savenije, F. J., and Engels, W. P. (2018). Active wake control: An approach to optimize the lifetime operation of wind farms. *Wind Energy*, 21(7):488–501.
- Kettle, A. J. (2014). Unexpected vertical wind speed profiles in the boundary layer over the southern North Sea. *Journal of Wind Engineering and Industrial Aerodynamics*, 134:149–162.
- Kiviluoma, J., Holttinen, H., Weir, D., Scharff, R., Söder, L., *et al.* (2016). Variability in large-scale wind power generation. *Wind Energy*, 19(9):1649–1665.
- Krogsæter, O. and Reuder, J. (2015). Validation of boundary layer parameterization schemes in the weather research and forecasting (wrf) model under the aspect of offshore wind energy applications—part ii: boundary layer height and atmospheric stability. *Wind Energy*, 18(7):1291–1302.
- Kuik, G. A. M. v., Peinke, J., Nijssen, R., Lekou, D., Mann, J., *et al.* (2016). Long-term research challenges in wind energy – a research agenda by the European Academy of Wind Energy. *Wind Energy Science*, 1(1):1–39.
- Lamb, H. and Frydendahl, K. (1991). *Historic Storms of the North Sea, British Isles and Northwest Europe*. Cambridge University Press, Cambridge (UK).
- Lange, B., Larsen, S., Højstrup, J., and Barthelmie, R. (2004). Importance of thermal effects and sea surface roughness for offshore wind resource assessment. *Journal of Wind Engineering and Industrial Aerodynamics*, 92(11):959–988.
- Lima, D. C. A., Soares, P. M. M., Semedo, A., and Cardoso, R. M. (2018). A global view of coastal low-level wind jets using an ensemble of reanalyses. *Journal of Climate*, 31(4):1525–1546.
- Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the atmospheric sciences*, 20(2):130–141.
- Lott, F. and Miller, M. J. (1997). A new subgrid-scale orographic drag parametrization: Its formulation and testing. *Quarterly Journal of the Royal Meteorological Society*, 123(537):101–127.
- Lundquist, J. K., DuVivier, K. K., Kaffine, D., and Tomaszewski, J. M. (2019). Costs and consequences of wind turbine wake effects arising from uncoordinated wind energy development. *Nature Energy*, 4(1):26.
- Lundtang Petersen, E., Troen, I., Ejlsing Jørgensen, H., and Mann, J. (2014). The new european wind atlas. *Energy Bulletin*, 17:34–39.

- Mahrt, L., Vickers, D., and Andreas, E. L. (2014). Low-level wind maxima and structure of the stably stratified boundary layer in the coastal zone. *Journal of Applied Meteorology and Climatology*, 53(2):363–376.
- Martini, M., Guanche, R., Losada, I. J., and Vidal, C. (2017). Accessibility assessment for operation and maintenance of offshore wind farms in the North Sea. *Wind Energy*, 20(4):637–656.
- Maureira Poveda, J. P. and Wouters, D. A. J. (2015). Wind Measurements at Meteorological Mast IJmuiden. Technical Report ECN-E-14-058, ECN, Petten (NL).
- McCaa, J., Cline, J. W., Shaw, W. J., Marquis, M., Berg, L. K., *et al.* (2017). Overview of the Wind Forecasting Improvement Project in Complex Terrain (WFIP2). Presented at the 97th American Meteorological Society Annual Meeting, Seattle, WA.
- Mehta, D., van Zuijlen, A. H., Koren, B., Holierhoek, J. G., and Bijl, H. (2014). Large eddy simulation of wind farm aerodynamics: A review. *Journal of Wind Engineering and Industrial Aerodynamics*, 133:1 – 17.
- Mérigaud, A. and Ringwood, J. V. (2016). Condition-based maintenance methods for marine renewable energy. *Renewable and Sustainable Energy Reviews*, 66:53–78.
- Miller, L. M., Gans, F., and Kleidon, A. (2011). Estimating maximum global land surface wind power extractability and associated climatic consequences. *Earth System Dynamics*, 2:1–12.
- Miller, S. T. K., Keim, B. D., Talbot, R. W., and Mao, H. (2003). Sea breeze: Structure, forecasting, and impacts. *Reviews of Geophysics*, 41(3).
- Miralles, D. G., Teuling, A. J., Van Heerwaarden, C. C., and De Arellano, J. V.-G. (2014). Mega-heatwave temperatures due to combined soil desiccation and atmospheric heat accumulation. *Nature geoscience*, 7(5):345.
- Mirocha, J., Kosović, B., and Kirkil, G. (2014). Resolved turbulence characteristics in large-eddy simulations nested within mesoscale simulations using the weather research and forecasting model. *Monthly Weather Review*, 142(2):806–831.
- Mirocha, J. D., Simpson, M. D., Fast, J. D., Berg, L. K., and Baskett, R. L. (2016). Investigation of boundary-layer wind predictions during nocturnal low-level jet events using the Weather Research and Forecasting model. *Wind Energy*, 19(4):739–762.
- Moeng, C.-H. and Sullivan, P. P. (1994). A Comparison of Shear- and Buoyancy-Driven Planetary Boundary Layer Flows. *Journal of the Atmospheric Sciences*, 51(7):999–1022.
- Moore, G. W. K. and Renfrew, I. A. (2005). Tip jets and barrier winds: A quikscat climatology of high wind speed events around greenland. *Journal of Climate*, 18(18):3713–3725.
- Muñoz-Esparza, D., Kosović, B., Mirocha, J., and van Beeck, J. (2014). Bridging the transition from meso-scale to microscale turbulence in numerical weather prediction models. *Boundary-Layer Meteorology*, 153(3):409–440.
- Murphy, A. H. (1988). Skill scores based on the mean square error and their relationships to the correlation coefficient. *Monthly Weather Review*, 116(12):2417–2424.
- Nunalee, C. G. and Basu, S. (2014). Mesoscale modeling of coastal low-level jets: implications for offshore wind resource estimation. *Wind Energy*, 17(8):1199–1216.
- Nygaard, N. G. (2014). Wakes in very large wind farms and the effect of neighbouring wind farms. *Journal of Physics: Conference Series*, 524(1).
- O'Brien, J. M., Young, T. M., O'Mahoney, D. C., and Griffin, P. C. (2017). Horizontal axis wind turbine research: A review of commercial cfd, fe codes and experimental practices. *Progress in Aerospace Sciences*, 92:1 – 24.
- Olauson, J. (2018). Era5: The new champion of wind power modelling? *Renewable Energy*, 126:322 – 331.
- Olson, J. B., Kenyon, J. S., Angevine, W., Brown, J. M., Pagowski, M., Sušelj, K., *et al.* (2019a). A description of the mynn-edmf scheme and the coupling to other components in wrf-arw. *NOAA Technical Memorandum OAR GSD-61*.
- Olson, J. B., Kenyon, J. S., Djalalova, I., Bianco, L., Turner, D. D., *et al.* (2019b). Improving wind energy forecasting through numerical weather prediction model development. *Bulletin of the American Meteorological Society*, 0(0):null.



- Oost, W. A., Komen, G. J., Jacobs, C. M. J., van Oort, C., and Bonekamp, H. (2001). Indications for a wave dependent charnock parameter from measurements during ASGAMAGE. *Geophysical Research Letters*, 28(14):2795–2797.
- Optis, M., Monahan, A., and Bosveld, F. C. (2014). Moving Beyond Monin-Obukhov Similarity Theory in Modelling Wind-Speed Profiles in the Lower Atmospheric Boundary Layer under Stable Stratification. *Boundary-Layer Meteorol*, 153(3):497–514.
- O'Neill, D. W., Fanning, A. L., Lamb, W. F., and Steinberger, J. K. (2018). A good life for all within planetary boundaries. *Nature Sustainability*, 1(2):88.
- Palutikof, J. P., Brabson, B. B., Lister, D. H., and Adcock, S. T. (1999). A review of methods to calculate extreme wind speeds. *Meteorological Applications*, 6(2):119–132.
- Parish, T. R. (2000). Forcing of the summertime low-level jet along the california coast. *Journal of Applied Meteorology*, 39(12):2421–2433.
- Park, J., Basu, S., and Manuel, L. (2014). Large-eddy simulation of stable boundary layer turbulence and estimation of associated wind turbine loads. *Wind Energy*, 17(3):359–384.
- Parker, W. S. (2016). Reanalyses and observations: What's the difference? *Bulletin of the American Meteorological Society*, 97(9):1565–1572.
- Parsons, M. J., Crosby, A. R., Orelup, L., Ferguson, M., and Cox, A. T. (2018). Evaluation of era5 reanalysis wind forcing for use in ocean response modeling.
- Peña, A., Floors, R., Sathe, A., Gryning, S.-E., Wagner, R., Courtney, M. S., Larsen, X. G., Hahmann, A. N., and Hasager, C. B. (2015). Ten Years of Boundary-Layer and Wind-Power Meteorology at Høvsøre, Denmark. *Boundary-Layer Meteorol*, 158(1):1–26.
- Peña, A., Hasager, C. B., Gryning, S.-E., Courtney, M., Antoniou, I., and Mikkelsen, T. (2009). Offshore wind profiling using light detection and ranging measurements. *Wind Energy*, 12(2):105–124.
- Petersen, E. L. and Troen, I. (2012). Wind conditions and resource assessment. *Wiley Interdisciplinary Reviews: Energy and Environment*, 1(2):206–217.
- Petersen, E. L., Troen, I., Jørgensen, H. E., and Mann, J. (2013). Are local wind power resources well estimated? *Environmental Research Letters*, 8(1):011005.
- Philipp, A., Bartholy, J., Beck, C., Erpicum, M., Esteban, P., *et al.* (2010). Cost733cat – a database of weather and circulation type classifications. *Physics and Chemistry of the Earth, Parts A/B/C*, 35(9):360 – 373.
- Pinson, P. and Hagedorn, R. (2012). Verification of the ecmwf ensemble forecasts of wind speed against analyses and observations. *Meteorological Applications*, 19(4):484–500.
- Plougonven, R. and Zhang, F. (2014). Internal gravity waves from atmospheric jets and fronts. *Reviews of Geophysics*, 52(1):33–76.
- Powers, J. G., Klemp, J. B., Skamarock, W. C., Davis, C. A., Dudhia, J., *et al.* (2017). The weather research and forecasting model: Overview, system efforts, and future directions. *Bulletin of the American Meteorological Society*, 98(8):1717–1737.
- Pryor, S. and Barthelmie, R. (2010). Climate change impacts on wind energy: A review. *Renewable and Sustainable Energy Reviews*, 14(1):430 – 437.
- Rácz, Z. and Smith, R. K. (1999). The dynamics of heat lows. *Quarterly Journal of the Royal Meteorological Society*, 125(553):225–252.
- Ranjha, R., Svensson, G., Tjernström, M., and Semedo, A. (2013). Global distribution and seasonal variability of coastal low-level jets derived from era-interim reanalysis. *Tellus A: Dynamic Meteorology and Oceanography*, 65(1):20412.
- Reen, B. P. and Stauffer, D. R. (2010). Data assimilation strategies in the planetary boundary layer. *Boundary-Layer Meteorology*, 137(2):237–269.
- REN21 (2016). *Renewables 2016 Global Status Report*. REN21 Secretariat (Paris).
- Riddaway, B. (2014). ECMWF Newsletter No.138 - Winter 2013/14.
- URL: <https://www.ecmwf.int/en/elibrary/14581-newsletter-no138-winter-2013-14>

- Rife, D. L., Pinto, J. O., Monaghan, A. J., Davis, C. A., and Hannan, J. R. (2010). Global distribution and characteristics of diurnally varying low-level jets. *Journal of Climate*, 23(19):5041–5064.
- Rife, D. L., Vanvyve, E., Pinto, J. O., Monaghan, A. J., Davis, C. A., and Poulos, G. S. (2013). Selecting representative days for more efficient dynamical climate downscaling: Application to wind energy. *Journal of Applied Meteorology and Climatology*, 52(1):47–63.
- Rockström, J., Steffen, W., Noone, K., Persson, Å., Chapin III, F. S., *et al.* (2009). A safe operating space for humanity. *Nature*, 461(7263):472.
- Rodrigues, S., Restrepo, C., Kontos, E., Pinto, R. T., and Bauer, P. (2015). Trends of offshore wind projects. *Renewable and Sustainable Energy Reviews*, 49:1114 – 1135.
- Rossby, C.-G. (1937a). On the mutual adjustment of pressure and velocity distributions in certain simple current systems, i. *Journal of Marine Research*, 1(1):15–28.
- Rossby, C.-G. (1937b). On the mutual adjustment of pressure and velocity distributions in certain simple current systems, ii. *Journal of Marine Research*, 1(3):239–263.
- Rotunno, R. (1983). On the linear theory of the land and sea breeze. *Journal of the Atmospheric Sciences*, 40(8):1999–2009.
- Sanderse, B., van der Pijl, S. P., and Koren, B. (2011). Review of computational fluid dynamics for wind turbine wake aerodynamics. *Wind Energy*, 14(7):799–819.
- Sandu, I., Beljaars, A., Bechtold, P., Mauritsen, T., and Balsamo, G. (2013). Why is it so difficult to represent stably stratified conditions in numerical weather prediction (nwp) models? *Journal of Advances in Modeling Earth Systems*, 5(2):117–133.
- Sanz Rodrigo, J., Churchfield, M., and Kosovic, B. (2017a). A methodology for the design and testing of atmospheric boundary layer models for wind energy applications. *Wind Energy Science*, 2(1):35–54.
- Sanz Rodrigo, J., Chávez Arroyo, R. A., Moriarty, P., Churchfield, M., Kosović, B., Réthoré, P.-E., Hansen, K. S., Hahmann, A., Mirocha, J. D., and Rife, D. (2017b). Mesoscale to microscale wind farm flow modeling and evaluation. *Wiley Interdisciplinary Reviews: Energy and Environment*, 6(2).
- Schaefer, J. T. (1990). The critical success index as an indicator of warning skill. *Weather and forecasting*, 5(4):570–575.
- Schalkwijk, J., Jonker, H. J. J., Siebesma, A. P., and Bosveld, F. C. (2015a). A year-long large-eddy simulation of the weather over cabauw: An overview. *Monthly Weather Review*, 143(3):828–844.
- Schalkwijk, J., Jonker, H. J. J., Siebesma, A. P., and Van Meijgaard, E. (2015b). Weather forecasting using gpu-based large-eddy simulations. *Bulletin of the American Meteorological Society*, 96(5):715–723.
- Schröter, J. S., Moene, A. F., and Holtslag, A. A. M. (2013). Convective boundary layer wind dynamics and inertial oscillations: the influence of surface stress. *Quarterly Journal of the Royal Meteorological Society*, 139(676):1694–1711.
- Schulz, C., Letzgus, P., Lutz, T., and Krämer, E. (2017). Cfd study on the impact of yawed inflow on loads, power and near wake of a generic wind turbine. *Wind Energy*, 20(2):253–268.
- Sedefian, L. (1980). On the vertical extrapolation of mean wind power density. *Journal of Applied Meteorology*, 19(4):488–493.
- Shapiro, A., Fedorovich, E., and Rahimi, S. (2016). A unified theory for the great plains nocturnal low-level jet. *Journal of the Atmospheric Sciences*, 73(8):3037–3057.
- Skamarock, W. C. and Klemp, J. B. (2008). A time-split nonhydrostatic atmospheric model for weather research and forecasting applications. *Journal of Computational Physics*, 227(7):3465 – 3485.
- Smeaton, J. (1759). Xviii. an experimental enquiry concerning the natural powers of water and wind to turn mills, and other machines, depending on a circular motion. *Philosophical transactions of the Royal society of London*, 51:100–174.
- Smedman, A.-S., Högström, U., and Bergström, H. (1996). Low level jets – a decisive factor for off-shore wind energy siting in the baltic sea. *Wind Engineering*, 20(3):137–147.
- Smedman, A.-S., Tjernström, M., and Högström, U. (1989). Analysis of the turbulence structure of a marine low-level jet. *Boundary-Layer Meteorology*, 66(1-2):105–126.

- Smedman, A.-S., Tjernström, M., and Högström, U. (1993). Analysis of the turbulence structure of a marine low-level jet. *Boundary-Layer Meteorology*, 66(1):105–126.
- Smith, S. D. (1988). Coefficients for sea surface wind stress, heat flux, and wind profiles as a function of wind speed and temperature. *Journal of Geophysical Research*, 93(C12):15467–15472.
- Snyder, B. and Kaiser, M. J. (2009). A comparison of offshore wind power development in europe and the u.s.: Patterns and drivers of development. *Applied Energy*, 86(10):1845 – 1856.
- Sørensen, J. N. and Larsen, G. C. (2018). Towards the north sea wind power revolution. *Wind Energy Science Discussions*, 2018:1–27.
- Staffell, I. and Pfenninger, S. (2016). Using bias-corrected reanalysis to simulate current and future wind power output. *Energy*, 114:1224 – 1239.
- Steele, C. J., Dorling, S. R., von Glasow, R., and Bacon, J. (2015). Modelling sea-breeze climatologies and interactions on coasts in the southern north sea: implications for offshore wind energy. *Quarterly Journal of the Royal Meteorological Society*, 141(690):1821–1835.
- Steenefeld, G.-J. (2014). Current challenges in understanding and forecasting stable boundary layers over land and ice. *Frontiers in Environmental Science*, 2:41.
- Steffen, W., Richardson, K., Rockström, J., Cornell, S. E., Fetzer, I., *et al.* (2015). Planetary boundaries: Guiding human development on a changing planet. *Science*, 347(6223):1259855.
- Stensrud, D. J. (1996). Importance of Low-Level Jets to Climate: A Review. *Journal of Climate*, 9(8):1698–1711.
- Storm, B. and Basu, S. (2010). The WRF Model Forecast-Derived Low-Level Wind Shear Climatology over the United States Great Plains. *Energies*, 3(2):258–276.
- Stull, R. (2000). *Meteorology for scientists and engineers*. Brooks/Cole, Pacific Grove (CA).
- Stull, R. B. (1988). *An Introduction to Boundary Layer Meteorology*. Springer Netherlands, Dordrecht (NL).
- Sun, W.-Y. and Sun, O. M. (2015). Bernoulli equation and flow over a mountain. *Geoscience Letters*, 2(1):7.
- Svensson, G. and Holtslag, A. A. M. (2009). Analysis of model results for the turning of the wind and related momentum fluxes in the stable boundary layer. *Boundary-Layer Meteorology*, 132(2):261–277.
- Talbot, C., Bou-Zeid, E., and Smith, J. (2012). Nested mesoscale large-eddy simulations with wrf: Performance in real test cases. *Journal of Hydrometeorology*, 13(5):1421–1441.
- Tastula, E.-M., Vihma, T., and Andreas, E. L. (2012). Evaluation of polar wrf from modeling the atmospheric boundary layer over antarctic sea ice in autumn and winter. *Monthly Weather Review*, 140(12):3919–3935.
- Taylor, J. R. and Ferrari, R. (2010). Buoyancy and wind-driven convection at mixed layer density fronts. *Journal of Physical Oceanography*, 40(6):1222–1242.
- Tijm, A. B. C., Holtslag, A. A. M., and van Delden, A. J. (1999). Observations and modeling of the sea breeze with the return current. *Monthly Weather Review*, 127(5):625–640.
- Tran, T., Tran, H., Mansfield, M., Lyman, S., and Crosman, E. (2018). Four dimensional data assimilation (fdda) impacts on wrf performance in simulating inversion layer structure and distributions of cmaq-simulated winter ozone concentrations in uintah basin. *Atmospheric Environment*, 177:75 – 92.
- Truewind, A. (2008). AWS Truewind’s final report for the alberta forecasting pilot project. *Wind Power Forecasting PILOT Project*, 66.
- Tsiringakis, A., Steeneveld, G. J., and Holtslag, A. A. M. (2017). Small-scale orographic gravity wave drag in stable boundary layers and its impact on synoptic systems and near-surface meteorology. *Quarterly Journal of the Royal Meteorological Society*, 143(704):1504–1516.
- Valkonen, T., Vihma, T., Johansson, M. M., and Launiainen, J. (2014). Atmosphere-sea ice interaction in early summer in the antarctic: evaluation and challenges of a regional atmospheric model. *Quarterly Journal of the Royal Meteorological Society*, 140(682):1536–1551.
- Van de Wiel, B. J. H., Moene, A. F., Steeneveld, G. J., Baas, P., Bosveld, F. C., and Holtslag, A. A. M. (2010). A conceptual view on inertial oscillations and nocturnal low-level jets. *Journal of the atmospheric sciences*, 67(8):2679–2689.

- Van Delden, A. (1992). The dynamics of meso-scale atmospheric circulations. *Physics Reports*, 211(6):251 – 374.
- Van Ulden, A. P. and Holtslag, A. A. M. (1985). Estimation of Atmospheric Boundary Layer Parameters for Diffusion Applications. *Journal of Climate and Applied Meteorology*, 24(11):1196–1207.
- Van Wijk, A. J. M., Beljaars, A. C. M., Holtslag, A. A. M., and Turkenburg, W. C. (1990). Evaluation of stability corrections in wind speed profiles over the North Sea. *Journal of Wind Engineering and Industrial Aerodynamics*, 33(3):551–566.
- Vautard, R., Thais, F., Tobin, I., Bréon, F.-M., De Lavergne, J.-G. D., Colette, A., Yiou, P., and Ruti, P. M. (2014). Regional climate model simulations indicate limited climatic impacts by operational and planned european wind farms. *Nature communications*, 5:3196.
- Verbong, G. and Geels, F. (2007). The ongoing energy transition: Lessons from a socio-technical, multi-level analysis of the dutch electricity system (1960-2004). *Energy Policy*, 35(2):1025–1037.
- Vieira, M., Snyder, B., Henriques, E., and Reis, L. (2019). European offshore wind capital cost trends up to 2020. *Energy Policy*, 129:1364 – 1371.
- Volker, P. H. J., Badger, J., Hahmann, A. N., and Ott, S. (2015). The explicit wake parametrisation v1. 0: a wind farm parametrisation in the mesoscale model wrf. *Geoscientific Model Development Discussions*, 8:3481–3522.
- Vorpahl, F., Schwarze, H., Fischer, T., Seidel, M., and Jonkman, J. (2013). Offshore wind turbine environment, loads, simulation, and design. *Wiley Interdisciplinary Reviews: Energy and Environment*, 2(5):548–570.
- Wagner, D., Steinfeld, G., Witha, B., Wurps, H., and Reuder, J. (2019). Low level jets over the southern north sea. *Meteorologische Zeitschrift*, pp. –.
- URL:** <http://dx.doi.org/10.1127/metz/2019/0948>
- Wallace, J. M. and Hobbs, P. V. (2006). *Atmospheric Science: An Introductory Survey*. Academic Press, Amsterdam (NL).
- Werkhoven, E. J. and Verhoef, J. P. (2012). Offshore Meteorological Mast IJmuiden - Abstract of Instrumentation Report. Technical Report ECN-Wind Memo-12-010, ECN.
- Westerhellweg, A., Canadillas, B., and Neumann, T. (2010). Direction dependency of offshore turbulence intensity in the german bight. In *DEWEK 2010 10th German wind energy conference*, pp. 66–70.
- Wharton, S. and Lundquist, J. K. (2012). Atmospheric stability affects wind turbine power collection. *Environmental Research Letters*, 7(1):014005.
- Whiteman, C. D. (2000). *Mountain meteorology: fundamentals and applications*. Oxford University Press.
- Wieringa, J. (1973). Gust factors over open water and built-up country. *Boundary-Layer Meteorology*, 3(4):424–441.
- Wilks, D. S. (2006). *Statistical methods in the atmospheric sciences*, volume 91. Academic press.
- Witha, B., Dörenkämper, M., Frank, H., García-Bustamante, E., González-Rouco, F., Navarro, J., Schneider, M., Steeneveld, G.-J., and Svensson, N. (2019a). The newa ferry lidar benchmark: Comparing mesoscale models with lidar measurements along a ship route.
- Witha, B., Hahmann, A., Sile, T., Dörenkämper, M., Ezber, Y., García-Bustamante, E., González-Rouco, J. F., Leroy, G., and Navarro, J. (2019b). WRF model sensitivity studies and specifications for the NEWA mesoscale wind atlas production runs.
- Zack, J. W. (2007). Optimization of wind power production forecast performance during critical periods for grid management. In *Proceedings of the European Wind Energy Conference EWEC, Milano (IT)*, volume 8.
- Zemba, J. and Friehe, C. A. (1987). The marine atmospheric boundary layer jet in the coastal ocean dynamics experiment. *Journal of Geophysical Research: Oceans*, 92(C2):1489–1496.



# Acknowledgements (Dankwoord)

I JUST REALIZED. I'VE BEEN REMISS. FORGIVE ME. I FORGOT TO THANK YOU FOR LOOKING OUT FOR ME YESTERDAY.

- NATHAN ALGREN, THE LAST SAMURAI.

Betekenis ontstaat door interactie. Ik kan schrijven en schrappen, zeggen en zwijgen, doen en laten wat ik wil; maar als niemand het leest, hoort, merkt of mist – wat heeft het dan voor waarde?

Dit hoofdstuk is een ode aan de mensen om mij heen. Aan iedereen die de afgelopen jaren direct of indirect heeft bijgedragen aan de totstandkoming van het boekje. Maar belangrijker nog: aan de ontwikkeling van de wetenschapper en de persoon die ik nu ben. Jullie hebben veel voor mij betekend.

Een productieve werkdag begint bij een prettige thuissituatie. Els, jij maakt mijn leven elke dag een stukje leuker. Je steunt me waar nodig en helpt me de balans bewaren tussen werk en ontspanning. Je bent veel slimmer (en nuchterder) dan ik en ik vind het mooi om te zien hoe we van elkaar leren en ons samen ontwikkelen.

Ook in Zwolle en Zevenbergen is het altijd prettig thuiskomen. De afstand van 'Zwolle' tot de academische wereld lijkt soms best groot, maar desondanks is er steeds de interesse en het meeleven. Bovendien kan ik altijd terecht voor advies en voorbeeld op sociaal en persoonlijk gebied. De term 'koude kant' heb ik nog nooit horen vallen buiten de kringen van mijn schoonfamilie. Ik prijs me des te gelukkiger dat ik hem überhaupt nog nooit heb gevoeld.

Ik hecht veel waarde aan een aantal vriendschappen die over vele jaren zijn opgebouwd, maar door toenemende afstand en 'afleiding' toch wel een beetje onder druk zijn komen te staan. Dat geldt voor Jonathan en Dennis, volleyballers van WaHo en huisgenoten van Dijkgraaf. Mooi dat we elkaar nog steeds weten te vinden!

I'll switch to English now, for that's generally understood in my professional surroundings. First of all that concerns the MAQ staff, who welcomed me in their group and provided me with critical and constructive feedback. In hindsight I think I should have done more to benefit from your expertise and devotion, which I have come to admire. A special thanks to our amazing secretaries, Caroline and Sandra, and of course I cannot forget Kees, our technician. The term 'support staff' does no justice to your contributions to the group. Besides your valuable work, your warm personalities also make the world a better place.

During my PhD I had the opportunity to supervise several thesis students. I enjoyed working with you, for I've seen some nice results and a lot of individual scientific progress. At the same time, you were like a mirror to me. Each of you worked very differently, and with some of you the supervision came more naturally than with others. This forced me to reflect on my own behaviour, both as a supervisor and a student. As such, you've all contributed to my personal and professional development.

That brings me to the role of my own supervisors. You gave me a lot of freedom and helped me cultivate my own ideas instead of imposing yours. I, in turn, was often critical; an attitude

that generally benefits research but that, if not properly kept in check, can also affect personal relations. I hope I didn't cause you too much trouble, although the evolution of Gert-Jan's hair colour over the past four years suggests otherwise...

Gert-Jan, you work tirelessly and super efficiently. You taught me to navigate the academic world, and you made sure we didn't deviate too much from the 'PhD roadmap'. I hope you'll find a suitable balance between your work and the new developments in your personal life. Reinder, you're always in for a chat and I'm often impressed by your ability to concisely and eloquently express the essence of a situation. You know your way around models and data, and I think MAQ is lucky to have you back. Your new position may be the most important one for the functioning and future development of the group. Bert, in Dutch I'd characterize your style as 'leiderschap met een lach'. You know how to clear the air and make people feel at ease, which formed the basis for a pleasant collaboration over the past four years. I took our meetings quite serious, and your attitude helped me to lighten up and see things in perspective.

My fellow PhD students at MAQ all contributed to a comfortable work environment, and the necessary distractions. Aris, thanks for the nice discussions during the trip to Copenhagen. At the time of writing, I'm looking forward to the next one. Imme, you're a very sociable person and I enjoyed our regular conversations in the mornings, before the office would become too crowded for a casual chat. Arjan, thanks for offering me a place to sleep after certain MAQ activities that rendered me unable to drive home. I'll miss our philosophical discussions and the walks around the campus. Lars, it was a pleasure to collaborate with you on the atmospheric practical modules, even though my contributions were minor compared to your amazing work. Good luck with your PhD in the states!


The following people have contributed directly to the content of this thesis. Fellow PhDs and their supervisors in the EUROS project, especially in work package 1. Several reviewers who took the time to evaluate my manuscripts, provided constructive feedback, and therein demonstrated benevolence rather than criticism. James, who initiated a pleasant collaboration with ECN, which resulted in the publication of chapter 4. Wim Bierbooms, who introduced us to the world of wind energy engineering. Sukanta for valuable feedback and enthusiasm. Various researchers I met at the conferences I attended. Mark, with whom a valuable exchange of ideas eventually led to a manuscript that is soon to be submitted. Scientists at KNMI where I attended the wind group meetings. Ine, as organiser of the wind group meetings and 'industry partner' in the EUROS project, you have contributed to the content of the thesis; but I would like to add that I've experienced your involvement in my project as very supportive.

I would like to acknowledge another group of people, who I don't know personally but without whom this research would not have been possible, or at least it wouldn't have been as easy. They are the developers of all the software packages I used. My work heavily relied on the (scientific) python ecosystem. One of the reasons this language is so successful is that there is a very active online community, and answers to almost every question can be found on stackoverflow. I've come to recognize some of the developers in the answers listed there, and I hope one day I can give something back to this amazing community.

As the defence is drawing near, I would like to name a few people who preceded and inspired me. The first defence I ever witnessed was that of Evan, and at the time I hardly realized what an amazing defence it was. I also admire the work of Jente, who does an amazing job when it comes to the communication of science. Chiel, although I didn't see your defence, I find your

scientific attitude very inspiring. And Gert-Jan Duine, it was a pleasure to work with you in France. Without that experience, I might not have seen myself as a potential PhD candidate.

I started the acknowledgements with the realization that there's no point in writing if nobody reads. Therefore, I want to thank you, the reader, for not letting my work go unnoticed. I hope it inspires new ideas and contributions, and that it strengthens the appreciation of meteorological processes in wind energy engineering applications. If there's nothing in it for you content wise, perhaps my figures can stir up some debate. I'd love for us to exchange some thoughts on the way in which scientific research is to be practised and communicated.



Source: xkcd.com.





## About the author

A broad interest in the natural sciences brought Peter to Wageningen, where he enrolled in the BSc programme “Soil, Water, Atmosphere”. He quickly discovered that the atmosphere, with its strong modelling and physics components, appealed most to him. He thus continued his studies in Wageningen, specializing in meteorology. After obtaining his master’s degree, he had the opportunity to start a PhD project on energy meteorology; an interesting topic that connected his education with his ambition to contribute to a sustainable society. This thesis describes the results of that project. Although he liked doing research, Peter felt that there was much to gain from advances in digital technologies. He now works at the Netherlands eScience center, where he hopes to make those technologies more easily accessible for the benefit of academic research.



Peter C. Kalverla (1992)

BSc Soil, Water,  
Atmosphere (2013)

MSc Earth and  
Environment (2015)

PhD Energy  
Meteorology (2019)

## Publications in academic journals

- Kalverla, P. C.**, Duine, G. J., Steeneveld, G. J., and Hedde, T. (2016). Evaluation of the Weather Research and Forecasting model in the Durance Valley complex terrain during the KASCADE field campaign. *Journal of Applied Meteorology and Climatology*, 55(4), 861-882.
- Verkley, W. T. M., **Kalverla, P. C.**, and Severijns, C. A. (2016). A maximum entropy approach to the parametrization of subgrid processes in two-dimensional flow. *Quarterly Journal of the Royal Meteorological Society*, 142(699), 2273-2283.
- Kalverla, P. C.**, Steeneveld, G. J., Ronda, R. J., and Holtslag, A. A. M. (2017). An observational climatology of anomalous wind events at offshore meteorological mast IJmuiden (North Sea). *Journal of Wind Engineering and Industrial Aerodynamics*, 165, 86-99.
- Kalverla, P. C.**, Steeneveld, G. J., Ronda, R. J., and Holtslag, A. A. M. (2019). Evaluation of three mainstream numerical weather prediction models with observations from meteorological mast IJmuiden at the North Sea. *Wind Energy*, 22(1), 34-48.
- Kalverla, P. C.**, Duncan Jr, J. B., Steeneveld, G. J., and Holtslag, A. A. M. (2019). Low-level jets over the North Sea based on ERA5 and observations: together they do better. *Wind Energy Science*, 4(2), 193-209.
- Kalverla, P. C.**, Holtslag, A. A. M., Ronda, R. J., and Steeneveld, G. J. (2019). Quality of wind characteristics in recent reanalyses over the North Sea. *Under review*.
- Schelbergen, M., **Kalverla, P. C.**, Schmehl, R., and Watson, S. J. (2019). Clustering wind profile shapes to estimate airborne wind energy production. *In preparation*.



*Netherlands Research School for the  
Socio-Economic and Natural Sciences of the Environment*

# D I P L O M A

*for specialised PhD training*

The Netherlands research school for the  
Socio-Economic and Natural Sciences of the Environment  
(SENSE) declares that

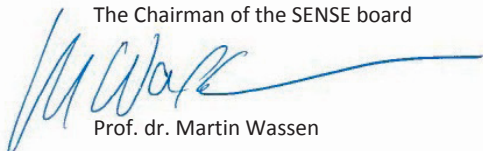
***Peter Christiaan Kalverla***

born on 9 January 1992 in Zwolle, The Netherlands

has successfully fulfilled all requirements of the  
educational PhD programme of SENSE.

Wageningen, 13 November 2019

The Chairman of the SENSE board



Prof. dr. Martin Wassen

the SENSE Director of Education



Dr. Ad van Dommelen

*The SENSE Research School has been accredited by the Royal Netherlands Academy of Arts and Sciences (KNAW)*



K O N I N K L I J K E N E D E R L A N D S E  
A K A D E M I E V A N W E T E N S C H A P P E N



The SENSE Research School declares that **Peter Christiaan Kalverla** has successfully fulfilled all requirements of the educational PhD programme of SENSE with a work load of 41.8 EC, including the following activities:

#### SENSE PhD Courses

- o Environmental research in context (2015)
- o Research in context activity: 'Initiating and co-organizing a symposium on energy meteorology entitled "From Model Field to Power Yield" (11 May 2017, Wageningen), and writing accessible summaries for a dedicated website and the Meteorologica popular science journal (2017)'

#### Other PhD and Advanced MSc Courses

- o Parametrization of subgrid physical processes, European Centre for Medium-range Weather Forecasts (2015)
- o ESCAPE Young Scientist Summer School, organised by The European H2020 FET project ESCAPE in Copenhagen (2019)
- o Supervising thesis students, Wageningen Graduate Schools (2017)
- o Orientation on teaching for PhD candidates, Wageningen Graduate Schools (2017)
- o Mobilizing your scientific network, Wageningen Graduate Schools (2016)

#### Management and Didactic Skills Training

- o Supervising six BSc student and one Msc student with thesis (2016-2018)
- o Assisting in the BSc course 'Meteorology and Climate' and BSc/Msc course 'Atmospheric modelling' (2016-2018) and various smaller teaching contributions
- o Guest lecture at TU Delft in the context of the course 'Siting conditions for offshore wind energy'(2019)
- o PhD Representative within chair group MAQ (2019)

#### Selection of Oral and Poster Presentations

- o An offshore climatology of anomalous wind events. 16th Annual meeting of the European Meteorological society, 12-16 September 2016, Trieste, Italy
- o *Validation of North Sea wind and weather in the WRF meteorological model*. Wind Energy Science Conference, 26-29 June 2017, Copenhagen, Denmark
- o *On the development and application of weather models for wind energy*. EUROS public seminar at Offshore Energy 2017, 11 October 2017, Amsterdam, The Netherlands
- o Invited speaker: *An offshore climatology of anomalous wind events*. KNMI wind day, 9 February 2018, de Bilt, The Netherlands
- o Outstanding poster award: *A North Sea climatology of anomalous wind events*. 18th Annual meeting of the European Meteorological society, 3-7 September 2018, Budapest, Hungary

SENSE Coordinator of PhD education

Dr. ir. Peter Vermeulen

The research described in this thesis is part the of the EUROS program, which is supported by NWO domain Applied and Engineering Sciences under grant number 14185 and partly funded by the Ministry of Economic Affairs.

Data analysis and numerical simulations were performed on the supercomputing facilities of SURFsara, sponsored by the Netherlands Organisation for Scientific Research (NWO) Physical Science Division (project number SH-312-15).

Financial support from Wageningen University for printing this thesis is gratefully acknowledged.

Cover design by Peter C. Kalverla and ProefschriftMaken

Printed on FSC certified paper by ProefschriftMaken/DigiForce



