

Low-level jets over the North Sea based on ERA5 and observations: together they do better

Peter C. Kalverla¹, James B. Duncan Jr.², Gert-Jan Steeneveld¹, and Albert A.M. Holtslag¹

¹Wageningen University, Meteorology and Air Quality Section, PO Box 47, 6700AA Wageningen, The Netherlands

²ECN part of TNO, Wind Energy, Westerduinweg 3, 1755LE Petten, The Netherlands

Correspondence: Peter Kalverla (peter.kalverla@wur.nl)

Abstract. Ten years of ERA5 reanalysis data are combined with met-mast and 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 large spatiotemporal 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.

For both measurement and model data, the representation of wind speed is biased. The limited temporal extent of observations leads to a wind speed bias on the order of $\pm 1 \text{ m s}^{-1}$ as compared to the long-term mean. In part due to data-assimilation strategies that cause abrupt discontinuities in the diurnal cycle, ERA5 also exhibits a wind speed bias of approximately 0.5 m s^{-1} . Representation of low-level jets in ERA5 is poor in terms of a one-to-one correspondence, and the jets appear vertically displaced ('smeared out'). However, climatological characteristics such as the shape of the seasonal cycle and the affinity with certain circulation patterns are represented quite well, albeit with different magnitudes. We therefore experiment with various methods to adjust modelled low-level jet rate to the observations or, vice versa, to correct for the erratic nature of the short observation periods using long-term ERA5 information. While quantitative uncertainty is still quite large, the presented results provide valuable insight into North Sea low-level jet characteristics. These jets occur predominantly for circulation types with an easterly component, with a clear peak in spring, and concentrate along the coasts at heights between 50-200 m. Further, it is demonstrated that these characteristics can be used as predictors to infer the observed low-level jet rate from ERA5 data with reasonable accuracy.

Copyright statement. TEXT

20 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*

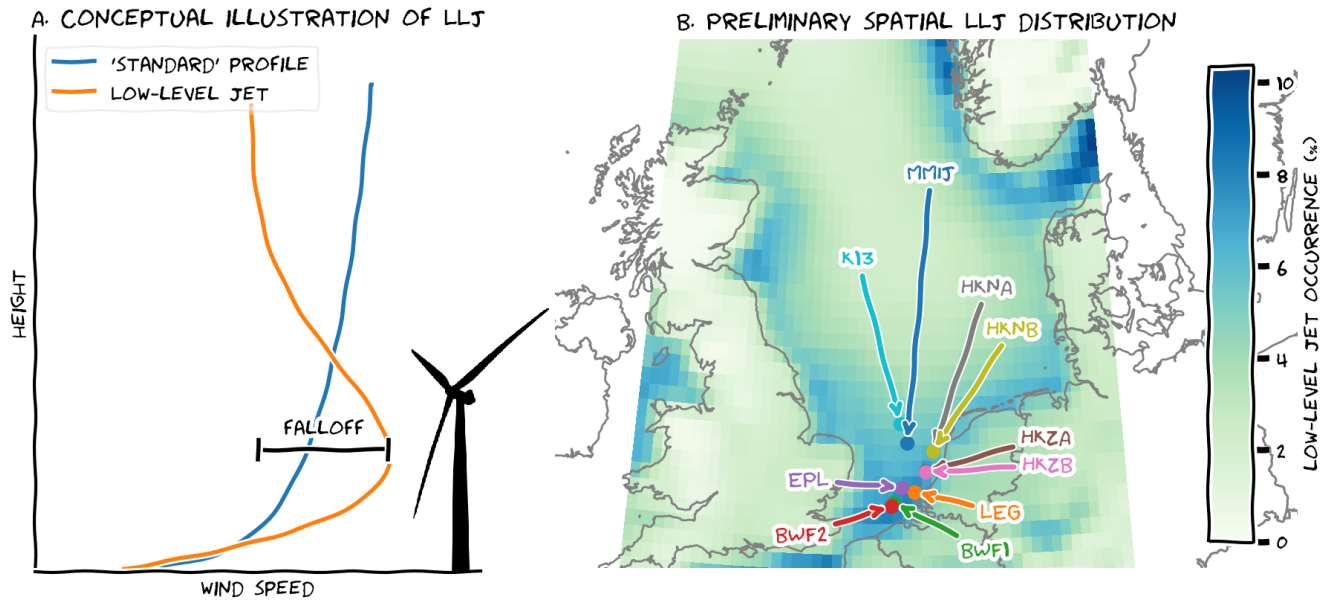


Figure 1. A. Example low-level jet profile as compared to the ‘standard’ logarithmic wind profile. B. Preliminary spatial distribution of annual low-level jet occurrence based on 10 years of ERA5 data up to 500 m. Overlaid are the location of the 10 measurement platforms used in this analysis: Met Mast IJmuiden (MMIJ), 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. Color coding is consistent across all figures.

parameterizations 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 parameterization 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 1A)¹. 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

¹In line with one of the core messages of the paper, i.e. the climatological characteristics of low-level jets are to be seen with some uncertainty, all figures in this manuscript have been rendered in a sketchy style

thermal contrast and differences in surface roughness between land and sea (Nunalee and Basu, 2014; Mahrt et al., 2014, e.g.). 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 (Moore and Renfrew, 2005, e.g.). Concerning 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
5 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 4-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, et cetera.
10 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 2; Copernicus Climate Change Service (C3S), 2017) reanalysis data and an extended set of observations.

Preliminary results based on 10 years of data in the lower 500 m of the atmosphere (Figure 1B) shows that ERA5 provides
15 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 ~ 1 year) 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
20 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
25 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
30 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.

A. DATA AVAILABILITY PER PLATFORM



B. MEASUREMENT/MODEL LEVELS




Figure 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 color coding in A highlights episodes of high (yellow/green) and low (blue) wind speed.

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.

To facilitate transparency and reproducibility, a series of Jupyter notebooks is available as supplementary material to this paper. 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.

2 A brief description of both datasets and their shortcomings

Observations are available from seven sites (Figure 1B). Three of these sites had two LiDARs operating simultaneously and one site (MMIJ) also featured a 90m met mast. The temporal span of measurements ranges from six months to over four years. 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 (Copernicus Climate Change Service (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
5 retrieved on a 0.3° by 0.3° 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 1B).

Before analyzing the morphology of these jets, we illustrate the limitations of both datasets concerning the representation
10 of wind speed. Figure 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 neighboring grid points). The full lines represent all 10 years of ERA5 data², 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
15 to coast). Dissimilarity between the ERA5 10-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 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 data span more than 4 years, some of it can be discarded in order to ensure an
20 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
25 reasonable accuracy. This procedure is known as measure-correlate-predict (MCP, 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 3C. In this diagram (co-opted from Kalverla et al., 2019), 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
30 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 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

²Some lines are exactly on top of each other because they are in the same grid point. Both are plotted, though, to preserve color coding




Figure 3. A. Averaged wind speed profiles for each measurement location, based on 10 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. Color coding is the same in all subplots, so C can serve as legend.

also contribute to overall error statistics. Based on the manufacturer information and previous validation (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 Supplementary material).

The observed biases exhibit a strong diurnal variation. During the night (Figure 3B), the bias is roughly between 0 and -0.5 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 12-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.

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 falloff are analysed. Falloff, as indicated in Figure 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 falloff threshold of 2 m s^{-1} . Figure 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 falloff 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 4A) is based on 10 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 falloff threshold of 2 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 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 falloff threshold and the total number of jets plotted.

Figure 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 10 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 falloff threshold) vanish, and not just those above 300 m. In order to classify a wind profile as a jet, falloff 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 falloff threshold in the ERA5 data as compared to the observations. Indeed, a more quantitative comparison in the form of a contingency table, based on one-to-one (1:1) jet correspondence between the two datasets, shows a very low critical success index (~ 0.2) and probability of detection (~ 0.2 , see Supplementary material). In other words, only 20% of low-level jets are correctly represented by ERA5. Does that imply that ERA5 is useless? No! Figure 4A indicates that potentially relevant information was filtered out. Even though the falloff is typically much smaller (to the extent that it falls below the falloff threshold), the height distribution of the ERA5 jets seems similar to the observations (also see Section 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

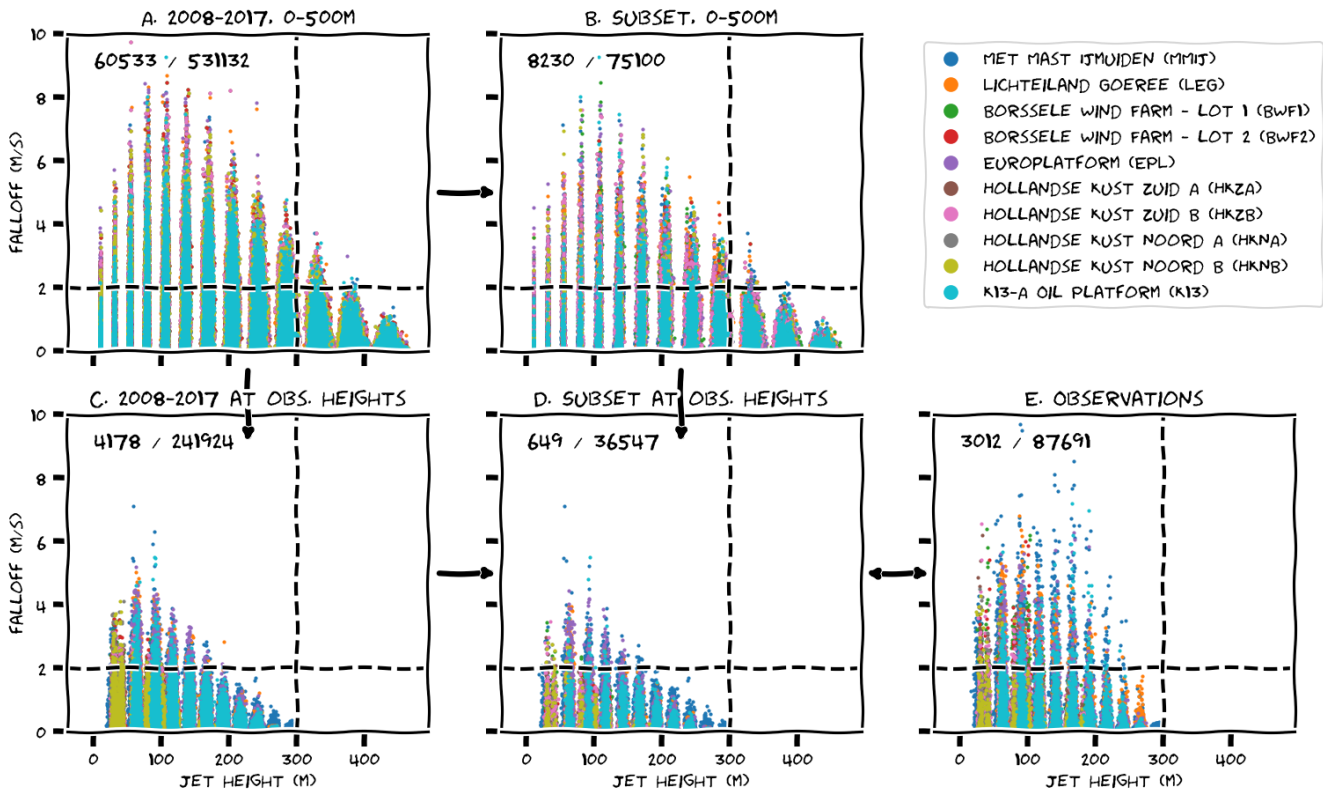


Figure 4. Scatter plots of falloff versus jet height for various representations of model data, and observations. In A and B, the jet height is represented by discrete model levels. Since these are specified in terms of pressure rather than height, they can exhibit small height variations in time. In 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.

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 ~ 0.2 ; probability of detection ~ 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 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 5 displays probability distributions of jet strengths (panel A) and jet heights (panel B) for various representations of the ERA5 data and observations³. It shows that the jet height and strength distributions are sensitive to the range limitation.

5 The median observed jet strength is about 8 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 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
10 seem to be somewhat closer to the surface than jets further offshore (not shown).

Three different representations of the observations are included in Figure 5. The first one is based on the 10-minute data. The second is based solely on the data of each full hour; in other words, we discarded 5/6th 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
15 the 10-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.

5 Datasets agree: most jets in spring and summer

Figure 6 displays the seasonal cycle of low-level jets and, in a similar fashion as Figure 4, how this cycle is subject to time
20 and range limitations. Over 10 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 10 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
25 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 (D) leads to a
30 very uninformative climatology, because the monthly low-level jet rates are all (close to) zero except for some unrepresentative

³Obviously, it is physically impossible to have a jet strength or height below zero. This is an artifact 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.

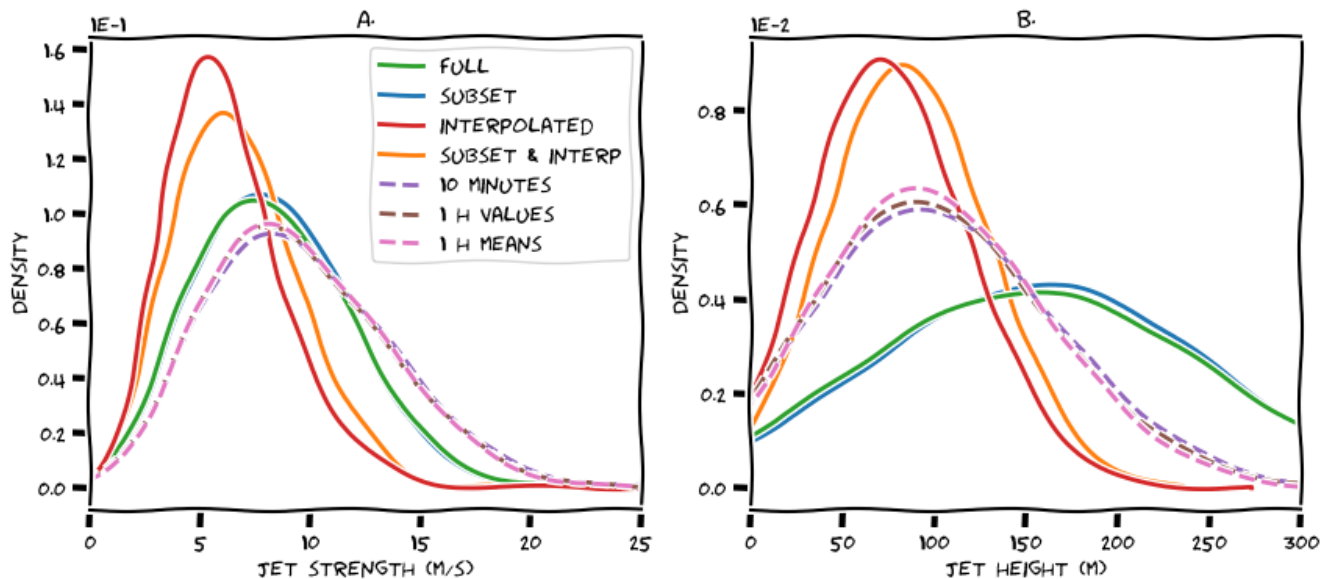


Figure 5. Kernel density estimates of the probability distribution of jet strength (A) and jet height (B) for various representations of the ERA5 data (full lines) and the observations (dashed lines), aggregated over all stations.

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 distill 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 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.

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

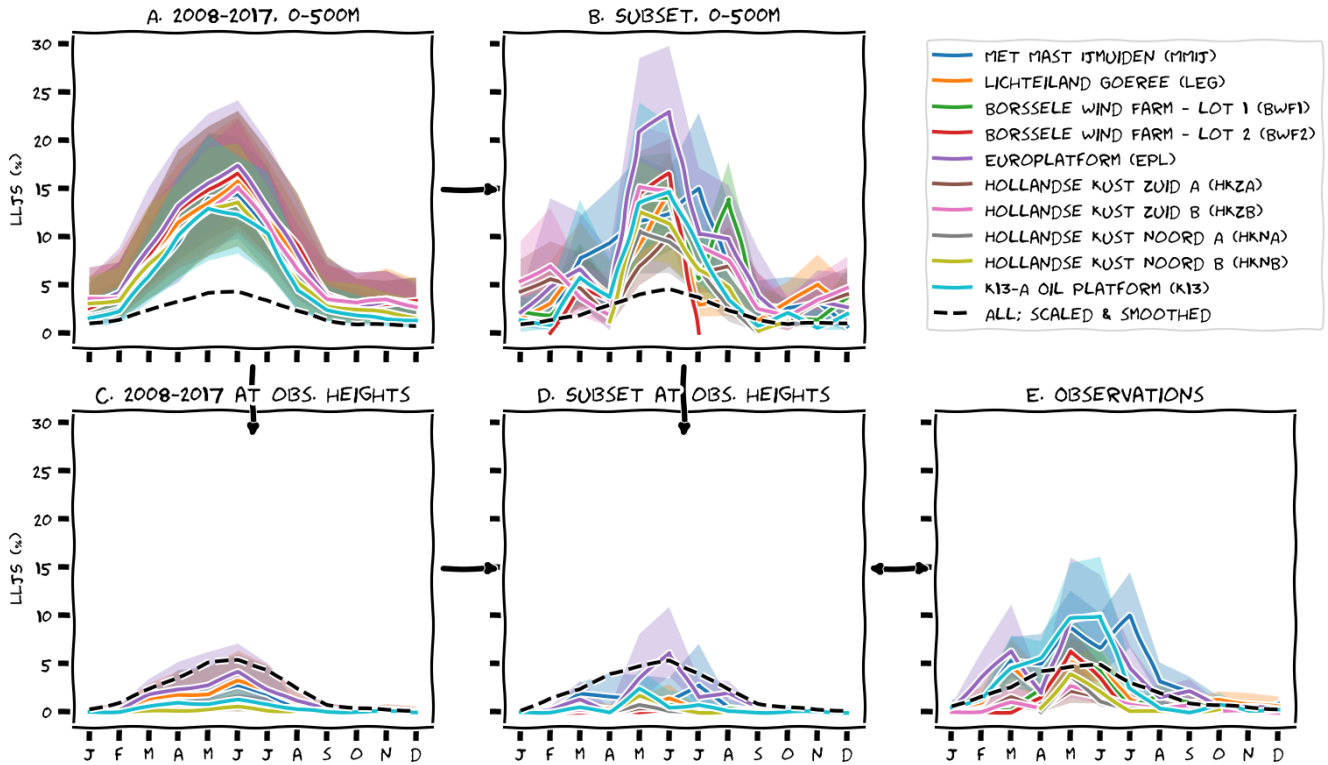


Figure 6. Seasonal cycle for various representations of model data, and observations. Shading is the sensitivity to $\pm 0.5 \text{ m s}^{-1}$ for the LLJ falloff threshold. The dashed lines represents an aggregated seasonal cycle of all platforms, smoothed with a rolling average of 3 months (2 at the 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.

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 1 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 in SI 4/6). For MMIJ the measurement period spanned more than 4 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 to 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 6. We compare panels B and E, and seek a fixed scaling factor that minimizes the difference between each pair of monthly observed and simulation LLJ frequencies. Denoting the *monthly mean* low-level jet frequency in ERA5 and collocated observations with x and y , respectively, an optimized scaling factor can be found by solving for a in $y = ax$ (using linear least squares regression). We do this for each platform individually, and also for their combined signal.

The results are illustrated in Figure 7A. The lighter colors 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 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 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 ~ 2 . However, as shown in Figure 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 3 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 (Suppl. Inf. 4/6).

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.

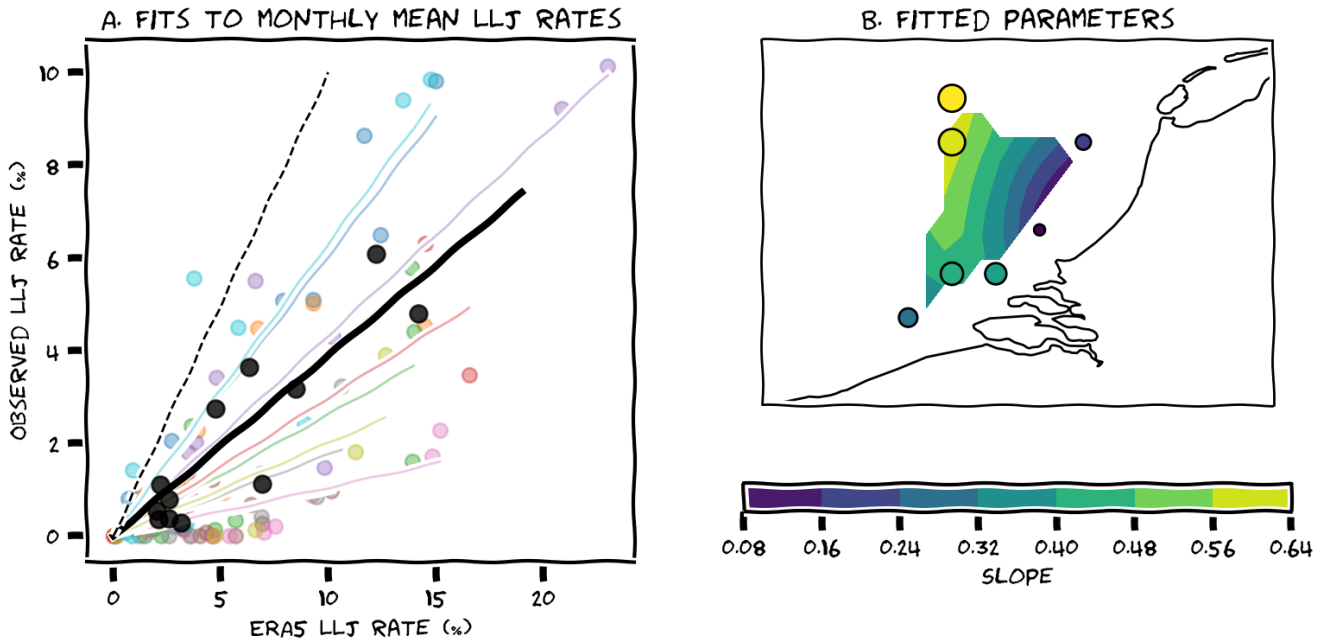


Figure 7. A. Illustration of linear regression between monthly low-level jet rates in the ERA5 data (subset, up to 500m) and the observations. Black line and scatter points represent aggregated data of all platforms, while the other colors correspond to fits for individual platforms. Dashed black line indicates a 1:1 correspondence. B. Spatial distribution of the obtained fit parameters for each individual platform. Like the color coding, marker size is scaled with the slope of the regression.

7 Other jet characteristics and their scaling potential

7.1 Diurnal cycle and stability

After analyzing in-depth the seasonal cycle of low-level jets, we now briefly consider some other variables that describe relevant characteristics of the low-level jet climatology, starting with the diurnal cycle. Figures 8A-C are again similar to Figure 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 1 and 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 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 8 shows the relation between low-level jet occurrence and atmospheric stability (expressed by the bulk Richardson number based on the ERA5 surface data: 2-m temperature, skin temperature, and 10-m 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 artifact 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 (~ 1.0), but qualitatively they seem to be less robust (not shown).

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 (Jones et al., 2013, LWT;) to perform this analysis. To derive these weather types we used the ERA5 mean sea level pressure on a 5° latitude/longitude grid of 16 points as laid out in the appendix of Jones et al. (2013), but centered over the area of interest. The method distinguishes three main groups: those with a dominant cyclonic (counterclockwise, 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, NorthEast, 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 9 shows the average low-level jet rate per weather type in the North Sea domain, based on 10 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 counterclockwise, and the middle full row is characterized by the absence of rotation. Notice how

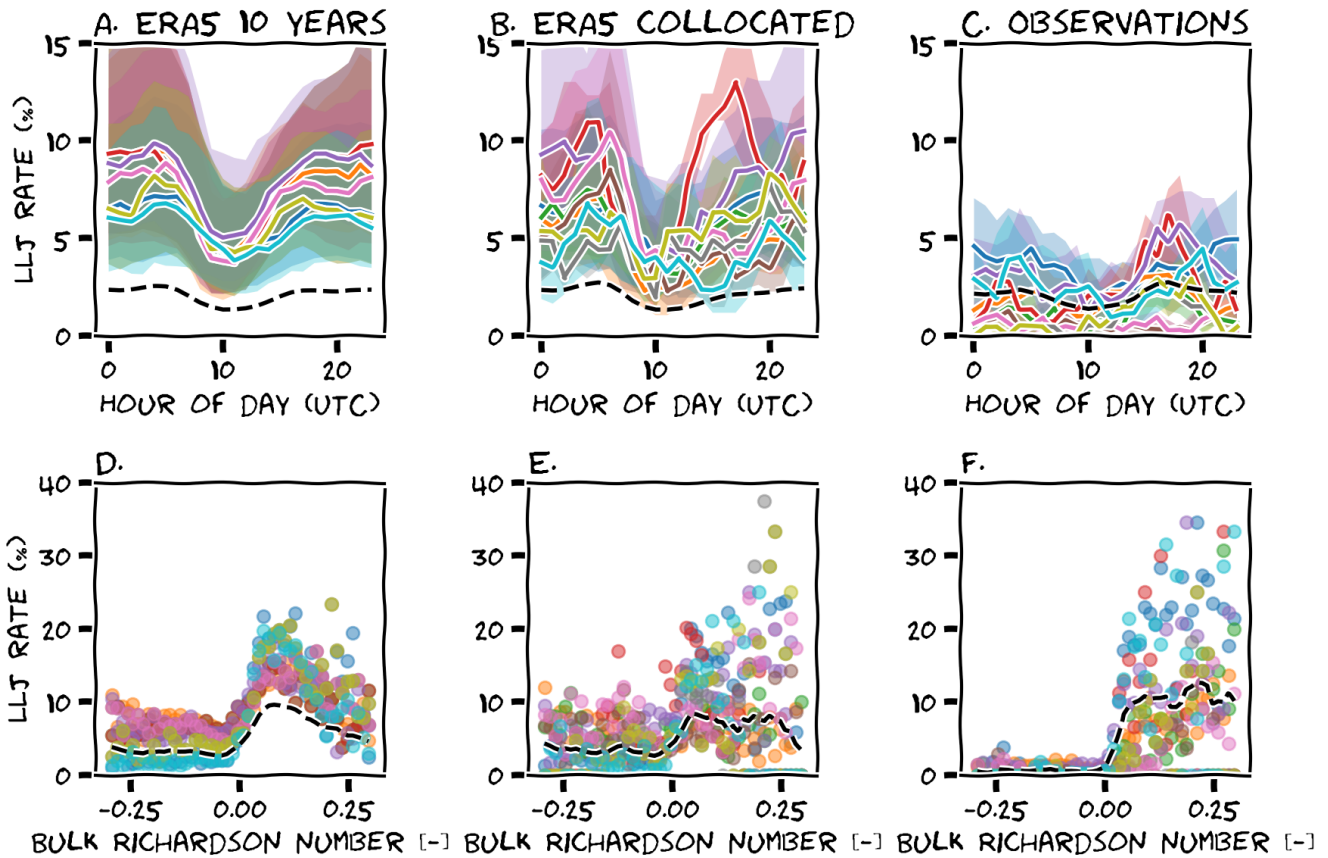


Figure 8. Average low-level jet rate for each hour of the day (A, B, C) and as function of the Bulk Richardson number (D, E, F), for the full (10 years) of ERA5 data up to 500m (A, D), a subset of this data collocated with the observations (B, E), and the observed data (C, F). Like in Figure 6, the black dashed lines represented a scaled and smoothed aggregated signal based on all platforms.

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

5 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 9 is actually meaningful, the absolute low-level jet rate (as indicated by the color bar) is still off by a factor of ~ 2 .

10 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, Ch. 6.5.2) and is similar to the Measure-Correlate-Predict methods mentioned in Section 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 supplementary material to facilitate reproducibility.

The general idea is illustrated in Figure 10A: we have a short timeseries 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 β_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

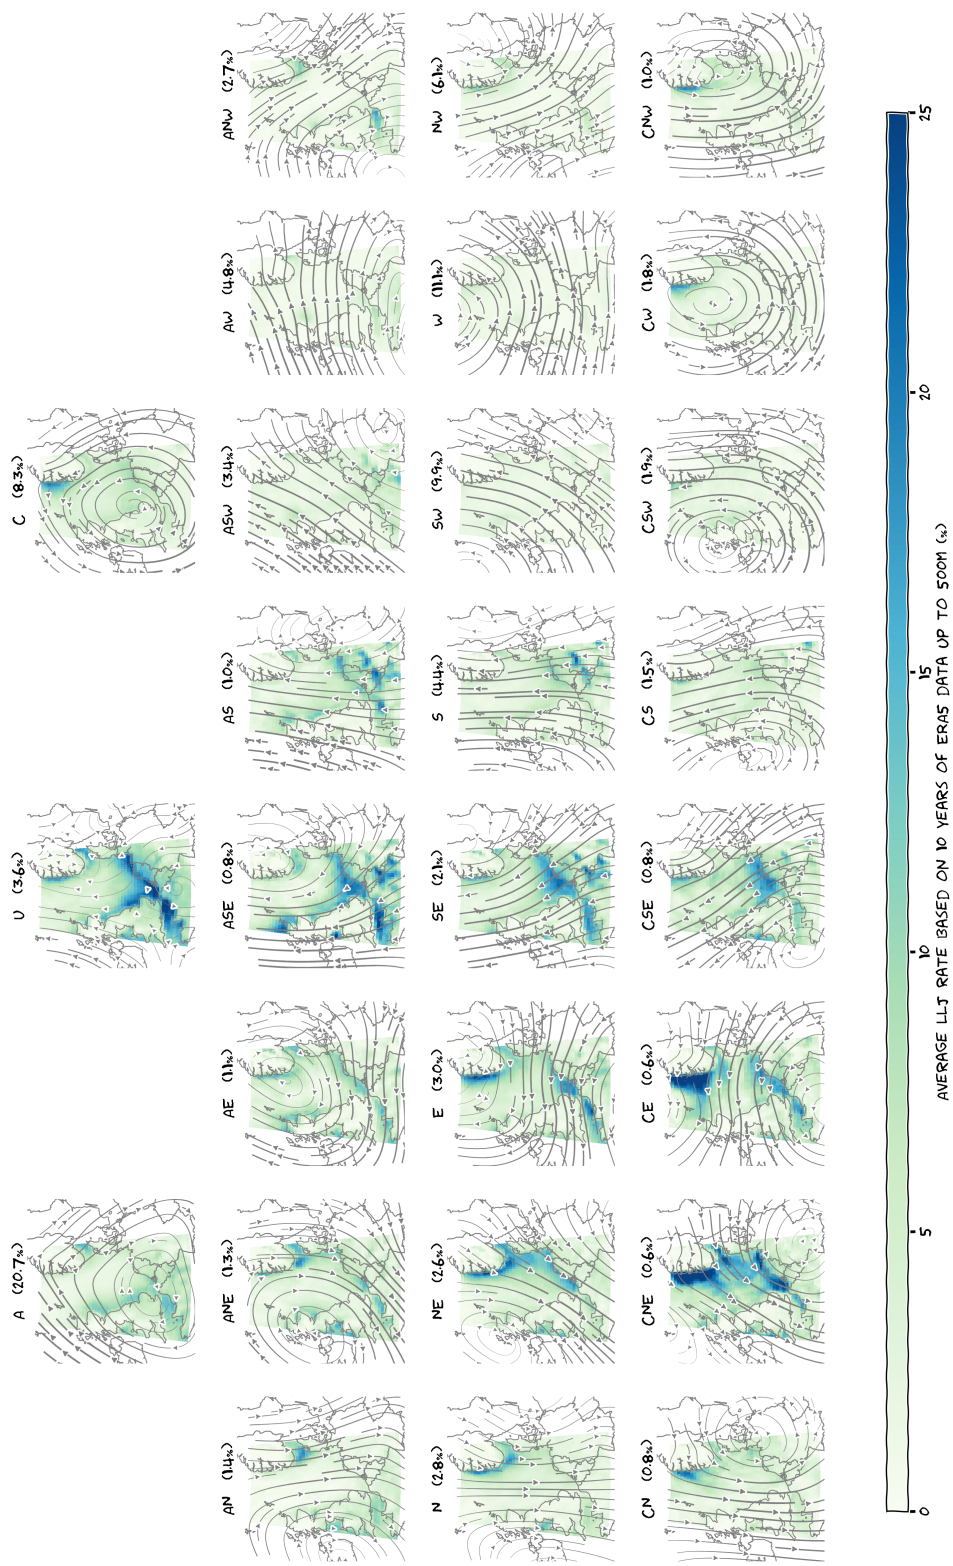


Figure 9. Spatial distribution of low-level jet rate in ERA5 data. As explained in the text, the values shown here overestimate available observations and should be interpreted with caution. A=Anticyclonic; C=Cyclonic U=Undefined; N, NE, E, etc are 8 wind direction sectors; combinations of a direction 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.

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-meter temperature and sea-surface temperature, which provided better results than the bulk Richardson number. We also
5 experimented with various training algorithms to determine the coefficients β_i of the logistic model (intermediate results can be found in Supplementary Material). In the end, we settled on a stochastic gradient descend algorithm.

First, we took only half of the MMIJ dataset (a bit more than 2 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 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
10 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 10B, orange
15 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 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%-point. This result confirms that the model generalizes well to new input data.

20 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%-point RMSE (but probably better). The observed seasonal cycle averaged over these four years of training data (Figure 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.
25 The result (Figure 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 6, we can conclude that we have adjusted the erratic nature of the short-term observations (Figure 6E), resulting in a seasonal cycle similar to that shown in Figure 6A, but with reduced amplitude. Compared to this final result, the crude amplitude adjustment with which we started in Section 6 now appears far too strong.

30 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.

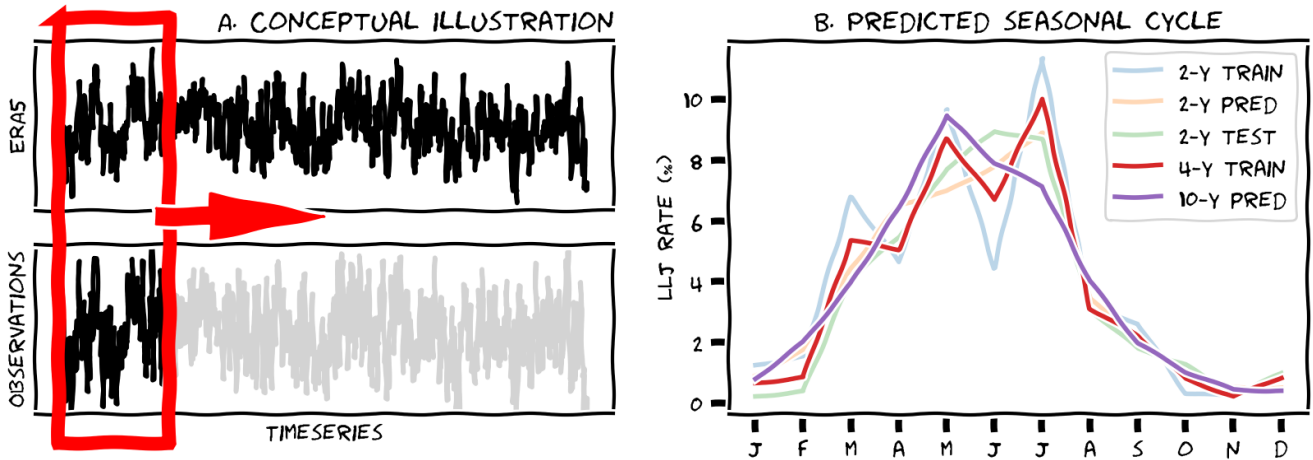


Figure 10. A. Illustration of the MCP/MOS/ML procedure in which a (logistic) model is trained with observation data and then used to predict long-term characteristics. B. Illustration of the MMIJ seasonal cycle of low-level jets based on: two years of observed data (train), two years of predicted (pred) and observed data (test), four years of observed data (train) and ten years of predicted data (pred).

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.

5 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^{-1} . 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
 10 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
 15 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^{-1} , 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 of about 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%-point 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 parameterizations 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.

Code and data availability. The ERA5 data were generated by ECMWF as part of the Copernicus Climate Change Service and will in the future be available through the Climate Data Store at <https://cds.climate.copernicus.eu/#!/home>. Observations were distributed by ECN part of TNO by order of the Dutch Ministry of Economic Affairs. They can be accessed at <https://windopzee.net/en/home/>. A series of Jupyter notebooks to facilitate reproducibility is available as supplementary material.

Appendix A: LiDAR data

Vertically pointing LiDAR provides efficient and non-intrusive measurement of ABL winds. Compared to traditional meteorological masts, LiDAR typically expand 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 spatiotemporal behavior. LiDAR type used included the WINDCUBE v2 pulsed LiDAR (only at LEG) and the Zephyr 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 10-min average vertical wind profile. A summary of the LiDAR measurement heights and data collection periods for all sites is provided in Figure 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 10-min 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 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 (σ_{TI}) greater than the mean (μ_{TI}) TI value (i.e. $TI \geq \mu_{TI} + 10\sigma_{TI}$); μ_{TI} and σ_{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 10-min mean wind speed exceeded 4 m s^{-1} .

Specific quality control measures were also applied to the LiDAR wind data. Any 10-min 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 $1e-5$ or greater than 100 – backscatter served as a proxy for CNR for data reported by the ZephIR 300s LiDAR.

Prior analyses (e.g. Poveda and Wouters, 2015) demonstrate that the ZephIR 300s LiDAR can incorrectly measure wind direction by 180° . Analyses of wind data at MIIJ 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.

Author contributions. Data analysis and preparation of the manuscript have been performed by the first author, under the supervision of the third and fourth author. Processing of LiDAR data was performed by second author, who also wrote the text for the appendix.

Competing interests. The authors declare that no competing interests are present

10 *Acknowledgements.* The analysis was performed on the high performance computing facility offered by the Dutch Science Organization NWO (grant number SH-312-15). This work is part of the EUROS project (NWO/TTW research grant number STW-14158). We thank two anonymous reviewers for their constructive feedback on an earlier version of the manuscript.

References

- Banakh, V. A., Smalikho, I. N., Köpp, F., and Werner, C.: Representativeness of wind measurements with a cw Doppler lidar in the atmospheric boundary layer, *Appl. Opt.*, 34, 2055–2067, <https://doi.org/10.1364/AO.34.002055>, <http://ao.osa.org/abstract.cfm?URI=ao-34-12-2055>, 1995.
- 5 Bhaganagar, K. and Debnath, M.: Implications of Stably Stratified Atmospheric Boundary Layer Turbulence on the Near-Wake Structure of Wind Turbines, *Energies*, 7, 5740–5763, <https://doi.org/10.3390/en7095740>, <http://www.mdpi.com/1996-1073/7/9/5740>, 2014.
- Blackadar, A. K.: Boundary Layer Wind Maxima and Their Significance for the Growth of Nocturnal Inversions, *Bulletin of the American Meteorological Society*, 38, 283–290, <https://doi.org/10.1175/1520-0477-38.5.283>, <https://doi.org/10.1175/1520-0477-38.5.283>, 1957.
- Bollmeyer, C., Keller, J. D., Ohlwein, C., Wahl, S., Crewell, S., Friederichs, P., Hense, A., Keune, J., Kneifel, S., Pscheidt, I., Redl, S., and
10 Steinke, S.: Towards a high-resolution regional reanalysis for the European CORDEX domain, *Quarterly Journal of the Royal Meteorological Society*, 141, 1–15, <https://doi.org/10.1002/qj.2486>, <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.2486>, 2015.
- Carta, J. A., Velázquez, S., and Cabrera, P.: 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, <https://doi.org/https://doi.org/10.1016/j.rser.2013.07.004>, <http://www.sciencedirect.com/science/article/pii/S1364032113004498>, 2013.
- 15 Carter, G. M., Dallavalle, J. P., and Glahn, H. R.: Statistical Forecasts Based on the National Meteorological Center’s Numerical Weather Prediction System, *Weather and Forecasting*, 4, 401–412, [https://doi.org/10.1175/1520-0434\(1989\)004<0401:SFBOTN>2.0.CO;2](https://doi.org/10.1175/1520-0434(1989)004<0401:SFBOTN>2.0.CO;2), [https://doi.org/10.1175/1520-0434\(1989\)004<0401:SFBOTN>2.0.CO;2](https://doi.org/10.1175/1520-0434(1989)004<0401:SFBOTN>2.0.CO;2), 1989.
- Christakos, K., Varlas, G., Reuder, J., Katsafados, P., and Papadopoulos, A.: Analysis of a Low-level Coastal Jet off the Western Coast of Norway, *Energy Procedia*, 53, 162 – 172, <https://doi.org/https://doi.org/10.1016/j.egypro.2014.07.225>, <http://www.sciencedirect.com/science/article/pii/S1876610214011023>, eERA DeepWind’ 2014, 11th Deep Sea Offshore Wind R&D Conference, 2014.
- 20 Copernicus Climate Change Service (C3S): ERA5: Fifth generation of ECMWF atmospheric reanalyses of the global climate, Copernicus Climate Change Service Climate Data Store (CDS), <https://cds.climate.copernicus.eu/cdsapp#!/home>, 2017.
- Dörenkämper, M., Optis, M., Monahan, A., and Steinfeld, G.: On the Offshore Advection of Boundary-Layer Structures and the Influence on Offshore Wind Conditions, *Boundary-Layer Meteorology*, 155, 459–482, <https://doi.org/10.1007/s10546-015-0008-x>, <https://doi.org/10.1007/s10546-015-0008-x>, 2015.
- 25 ECMWF: IFS Documentation – Cy41r2, chap. II: Data Assimilation, ECMWF, 2016.
- Englberger, A. and Dörnbrack, A.: A Numerically Efficient Parametrization of Turbulent Wind-Turbine Flows for Different Thermal Stratifications, *Boundary-Layer Meteorology*, 169, 505–536, <https://doi.org/10.1007/s10546-018-0377-z>, <https://doi.org/10.1007/s10546-018-0377-z>, 2018.
- 30 Gebraad, P. M. O., Teeuwisse, F. W., Wingerden, J. W., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y.: Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study, *Wind Energy*, 19, 95–114, <https://doi.org/10.1002/we.1822>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1822>, 2014.
- Glahn, H. R. and Lowry, D. A.: The Use of Model Output Statistics (MOS) in Objective Weather Forecasting, *Journal of Applied Meteorology*, 11, 1203–1211, [https://doi.org/10.1175/1520-0450\(1972\)011<1203:TUOMOS>2.0.CO;2](https://doi.org/10.1175/1520-0450(1972)011<1203:TUOMOS>2.0.CO;2), [https://doi.org/10.1175/1520-0450\(1972\)011<1203:TUOMOS>2.0.CO;2](https://doi.org/10.1175/1520-0450(1972)011<1203:TUOMOS>2.0.CO;2), 1972.
- 35 Gutierrez, W., Ruiz-Columbie, A., Tutkun, M., and Castillo, L.: Impacts of the low-level jet’s negative wind shear on the wind turbine, *Wind Energy Science*, 2, 533–545, <https://doi.org/10.5194/wes-2-533-2017>, <https://www.wind-energ-sci.net/2/533/2017/>, 2017.

- Holton, J. R.: The diurnal boundary layer wind oscillation above sloping terrain, *Tellus*, 19, 200–205, <https://doi.org/10.3402/tellusa.v19i2.9766>, <https://doi.org/10.3402/tellusa.v19i2.9766>, 1967.
- Holtslag, A. A. M., Svensson, G., Baas, P., Basu, S., Beare, B., Beljaars, A. C. M., Bosveld, F. C., Cuxart, J., Lindvall, J., Steeneveld, G. J., Tjernström, M., and Van De Wiel, B. J. H.: Stable Atmospheric Boundary Layers and Diurnal Cycles: Challenges for Weather and Climate Models, *Bulletin of the American Meteorological Society*, 94, 1691–1706, <https://doi.org/10.1175/BAMS-D-11-00187.1>, <https://doi.org/10.1175/BAMS-D-11-00187.1>, 2013.
- Jones, P. D., Harpham, C., and Briffa, K. R.: Lamb weather types derived from reanalysis products, *International Journal of Climatology*, 33, 1129–1139, <https://doi.org/10.1002/joc.3498>, <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/joc.3498>, 2013.
- Kalverla, P., Steeneveld, G.-J., Ronda, R., and Holtslag, A. A.: Evaluation of three mainstream numerical weather prediction models with observations from meteorological mast IJmuiden at the North Sea, *Wind Energy*, 22, 34–48, <https://doi.org/10.1002/we.2267>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.2267>, 2019.
- Kalverla, P. C., Steeneveld, G.-J., Ronda, R. J., and Holtslag, A. A.: An observational climatology of anomalous wind events at offshore meteomast IJmuiden (North Sea), *Journal of Wind Engineering and Industrial Aerodynamics*, 165, 86 – 99, <https://doi.org/https://doi.org/10.1016/j.jweia.2017.03.008>, <http://www.sciencedirect.com/science/article/pii/S0167610516307061>, 2017.
- Lima, D. C. A., Soares, P. M. M., Semedo, A., and Cardoso, R. M.: A Global View of Coastal Low-Level Wind Jets Using an Ensemble of Reanalyses, *Journal of Climate*, 31, 1525–1546, <https://doi.org/10.1175/JCLI-D-17-0395.1>, <https://doi.org/10.1175/JCLI-D-17-0395.1>, 2018.
- Mahrt, L., Vickers, D., and Andreas, E. L.: Low-Level Wind Maxima and Structure of the Stably Stratified Boundary Layer in the Coastal Zone, *Journal of Applied Meteorology and Climatology*, 53, 363–376, <https://doi.org/10.1175/JAMC-D-13-0170.1>, <https://doi.org/10.1175/JAMC-D-13-0170.1>, 2014.
- Moore, G. W. K. and Renfrew, I. A.: Tip Jets and Barrier Winds: A QuikSCAT Climatology of High Wind Speed Events around Greenland, *Journal of Climate*, 18, 3713–3725, <https://doi.org/10.1175/JCLI3455.1>, <https://doi.org/10.1175/JCLI3455.1>, 2005.
- Nunalee, C. G. and Basu, S.: Mesoscale modeling of coastal low-level jets: implications for offshore wind resource estimation, *Wind Energy*, 17, 1199–1216, <https://doi.org/10.1002/we.1628>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1628>, 2014.
- Olauson, J.: ERA5: The new champion of wind power modelling?, *Renewable Energy*, 126, 322 – 331, <https://doi.org/https://doi.org/10.1016/j.renene.2018.03.056>, <http://www.sciencedirect.com/science/article/pii/S0960148118303677>, 2018.
- Park, J., Basu, S., and Manuel, L.: Large-eddy simulation of stable boundary layer turbulence and estimation of associated wind turbine loads, *Wind Energy*, 17, 359–384, <https://doi.org/10.1002/we.1580>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.1580>, 2014.
- Peña, A., Hasager, C. B., Badger, M., Barthelmie, R. J., Bingöl, F., Cariou, J.-P., Emeis, S., Frandsen, S. T., Harris, M., Karagali, I., et al.: Remote sensing for wind energy, *DTU Wind Energy*, 2015.
- Petersen, E. L., Troen, I., Jørgensen, H. E., and Mann, J.: Are local wind power resources well estimated?, *Environmental Research Letters*, 8, 011 005, <http://stacks.iop.org/1748-9326/8/i=1/a=011005>, 2013.
- Poveda, J. M. and Wouters, D.: Wind measurements at meteorological mast IJmuiden, *ECN*, 2015.
- Ranjha, R., Svensson, G., Tjernström, M., and Semedo, A.: Global distribution and seasonal variability of coastal low-level jets derived from ERA-Interim reanalysis, *Tellus A: Dynamic Meteorology and Oceanography*, 65, 20 412, <https://doi.org/10.3402/tellusa.v65i0.20412>, <https://doi.org/10.3402/tellusa.v65i0.20412>, 2013.

- Reen, B. P. and Stauffer, D. R.: Data Assimilation Strategies in the Planetary Boundary Layer, *Boundary-Layer Meteorology*, 137, 237–269, <https://doi.org/10.1007/s10546-010-9528-6>, <https://doi.org/10.1007/s10546-010-9528-6>, 2010.
- Rife, D. L., Pinto, J. O., Monaghan, A. J., Davis, C. A., and Hannan, J. R.: Global Distribution and Characteristics of Diurnally Varying Low-Level Jets, *Journal of Climate*, 23, 5041–5064, <https://doi.org/10.1175/2010JCLI3514.1>, <https://doi.org/10.1175/2010JCLI3514.1>, 5 2010.
- Sedefian, L.: On the Vertical Extrapolation of Mean Wind Power Density, *Journal of Applied Meteorology*, 19, 488–493, [https://doi.org/10.1175/1520-0450\(1980\)019<0488:OTVEOM>2.0.CO;2](https://doi.org/10.1175/1520-0450(1980)019<0488:OTVEOM>2.0.CO;2), [https://doi.org/10.1175/1520-0450\(1980\)019<0488:OTVEOM>2.0.CO;2](https://doi.org/10.1175/1520-0450(1980)019<0488:OTVEOM>2.0.CO;2), 1980.
- Shapiro, A., Fedorovich, E., and Rahimi, S.: A Unified Theory for the Great Plains Nocturnal Low-Level Jet, *Journal of the Atmospheric Sciences*, 73, 3037–3057, <https://doi.org/10.1175/JAS-D-15-0307.1>, <https://doi.org/10.1175/JAS-D-15-0307.1>, 10 2016.
- Smedman, A.-S., Tjernström, M., and Högström, U.: Analysis of the turbulence structure of a marine low-level jet, *Boundary-Layer Meteorology*, 66, 105–126, <https://doi.org/10.1007/BF00705462>, <https://doi.org/10.1007/BF00705462>, 1993.
- Staffell, I. and Pfenninger, S.: Using bias-corrected reanalysis to simulate current and future wind power output, *Energy*, 114, 1224 – 1239, <https://doi.org/https://doi.org/10.1016/j.energy.2016.08.068>, <http://www.sciencedirect.com/science/article/pii/S0360544216311811>, 15 2016.
- Steele, C. J., Dorling, S. R., von Glasow, R., and Bacon, J.: 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, 1821–1835, <https://doi.org/10.1002/qj.2484>, <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.2484>, 2015.
- Tran, T., Tran, H., Mansfield, M., Lyman, S., and Crosman, E.: 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, <https://doi.org/https://doi.org/10.1016/j.atmosenv.2018.01.012>, <http://www.sciencedirect.com/science/article/pii/S1352231018300189>, 20 2018.
- Van de Wiel, B. J. H., Moene, A. F., Steeneveld, G. J., Baas, P., Bosveld, F. C., and Holtslag, A. A. M.: A Conceptual View on Inertial Oscillations and Nocturnal Low-Level Jets, *Journal of the Atmospheric Sciences*, 67, 2679–2689, <https://doi.org/10.1175/2010JAS3289.1>, <https://doi.org/10.1175/2010JAS3289.1>, 25 2010.
- Wharton, S. and Lundquist, J. K.: Atmospheric stability affects wind turbine power collection, *Environmental Research Letters*, 7, 014 005, <http://stacks.iop.org/1748-9326/7/i=1/a=014005>, 2012.
- Wilks, D. S.: *Statistical methods in the atmospheric sciences*, vol. 91, Academic press, 2006.