Norwegian hindcast archive (NORA3) - A validation of offshore wind resources in the North Sea and Norwegian Sea

Ida Marie Solbrekke[1], Asgeir Sorteberg[2], and Hilde Haakenstad[3]

[1]Geophysical Institute, Bergen Offshore Wind Centre (BOW), University of Bergen, Allegaten 70, 5020 Bergen, Norway
[2]Geophysical Institute, Bjerknes Centre for Climate Research (BCCR), Bergen Offshore Wind Centre (BOW), University of Bergen, Norway
[3]Geophysical Institute, University of Bergen, Norwegian Meteorological Institute, Norway

Correspondence: Ida Marie Solbrekke (ida.solbrekke@uib.no)

Abstract. A new high-resolution (3 km) numerical mesoscale weather simulation for the time-period 2004-2016 is validated for offshore wind power purposes for the North Sea and Norwegian Sea. The NORwegian hindcast Archive (NORA3) is a dynamically downscaled data set, forced with state-of-the-art atmospheric reanalysis as boundary conditions. An in depth validation of the simulated wind climatology towards the observed wind climatology has been carried out to determine whether

5 NORA3 can serve as a wind resource data set in the planning-phase of future offshore wind power installations. Special emphasis is placed on evaluating offshore wind power-related metrics and the impact of simulated wind speed deviations on the estimated wind power and the related variability. The general conclusion of the validation is that the NORA3 data is rather well suited for wind power estimates, but gives slightly conservative estimates on the offshore wind metrics. In other words, wind speeds in NORA3 are typically 5 % (0.5 ms^{-1}) lower than observed wind speeds, giving an underestimation of offshore

10 wind power of 10 %-20 % (equivalent to an underestimation of 3 percentage point in the capacity factor), for a selected turbine type and hub height. The model is biased towards lower wind power estimates due to overestimation of the wind speed events below typical wind speed limits of rated wind power (u <11-13 ms^{-1}) and underestimation of high wind speed events (u >11-13 ms^{-1}). The hourly wind speed and wind power variability are slightly underestimated in NORA3. However, the number of hours with zero power production caused by the wind conditions (around 12 % of the time) is well captured, while the duration

15 of each of these events is slightly overestimated, leading to 25-year return values for zero-power duration being too high for the majority of the sites. The model performs well in capturing spatial co-variability in hourly wind power production, with only small deviations in the spatial correlation coefficients among the sites. The observation-based decorrelation length was estimated to be 426 km, whereas the model-based length was 16 % longer.

1 Introduction

20 Exploiting the Norwegian continental shelf for offshore wind power purposes is highly relevant due to the excellent wind climate (Zheng et al., 2016) and the recent increase in political engagement. In June 2020 the Norwegian government decided to open the country's first two offshore areas at the Norwegian continental shelf, "Utsira Nord" and "Sørlige Nordsjøen II",

for concessions to build and operate large wind power installations (Regjeringen, 2020). In this context, the ability to map the spatial and temporal wind power potential is crucial for selecting the best areas for wind power production.

25   Observational sites in the North sea and the Norwegian sea are sparse and their numbers are insufficient to map the regional wind power potential. The lack of observational data makes it challenging for stakeholders and decision makers to choose new sites to open for offshore wind power concessions. Apart from using satellite data on surface winds, the only way to map the total wind power potential for a large offshore area is to use data from high-resolution numerical weather prediction (NWP) models that provides data near a typical hub height.

30   Several studies have mapped the wind energy potential of the North Sea and/or the Norwegian Sea using simulated data from the mesoscale Weather Research and Forecasting Model (WRF) (Berge et al., 2009; Byrkjedal and Åkervik, 2009; Byrkjedal et al., 2010; Skeie et al., 2012; Hasager et al., 2020). Berge et al. (2009) investigated how well the WRF model captured the offshore wind conditions in the North Sea from 2004-2007. After comparison of the simulated data with observations from oil- and gas platforms the authors conclude that WRF model is a reliable tool for characterizing the average wind conditions in the region in question. The model was verified using observations from offshore sites in the North Sea, but did not underwent a peer-review process. Byrkjedal and Åkervik (2009) simulated the wind resource and wind power potential at the Norwegian economic zone. The WRF model produced the simulated data for 2000-2008 used in their wind power calculations. However, the simulated data set was not validated against observations and the report was not peer-reviewed. Byrkjedal et al. (2010) used the WRF model to simulate the offshore wind power potential in the North Sea from 2000 to 2009. Based on their 10-year WRF-simulation they estimated wind power and identified areas with the greatest wind power potential, in addition to the dependency between separation distance and the correlation between two wind power production sites. The model performance was not compared to observations and the report was not evaluated in a peer-review process. The more recent data set, the New European Wind Atlas (NEWA), has been created as a cooperation between several European wind energy-related companies. The NEWA aims to provide a high-resolution, freely available data set on wind energy resources in Europe (NEWA, 2020). NEWA uses met masts on land to validate the onshore model data, while the offshore data is validated at 10 meters above sea level (m.a.s.l.) using satellite data. In addition, a validation at 100 m.a.s.l. is conducted by extrapolating the 10-m data. NEWA underwent a peer-review process. A peer-reviewed validation of wind model simulations before using the data for offshore wind power purposes is very important. The degree of data set validation and peer-review process of the results in the preceding studies are either limited or nonexistent.

50   In this study we perform an in detail validation of a new and freely available high-resolution data set called *the NORwegian hindcast Archive* (NORA3) to be used for offshore wind resource assessment and wind power estimates. NORA3 is a high-resolution atmospheric dynamic downscaling of the state-of-the-art reanalysis data from ECMWF, called ERA-5. The downscaling of ERA-5 is performed by the NWP model HARMONIE-AROME (H-A). H-A is a high-resolution NWP model developed and used by many European weather forecast and research institutions (Seity et al., 2011; Bengtsson et al., 2017).

55 The creation of NORA3 will contribute to the growing ensemble of wind resource data sets. Since all currently existing wind resource data sets are generated by the WRF model, the creation of NORA3 by a different NWP model will contribute to a diversity in the available wind resource data sets. When the ensemble of these data sets are considered in wind power planning

the overall uncertainty in power production can be better quantified. The usefulness of multi model ensembles has become increasingly clear over the last few decades in research field such as weather prediction and climate change. By this extensive validation of the NORA3 data set and documenting the quality of the simulated wind resource and related wind power estimates from a new model we wish to contribute to the growing literature on offshore wind resources.

The novelty of this study is the in depth validation of the model data using a new NWP model. Through advanced statistical measures we perform a near-hub-height validation of the NORA3 estimated wind resource and the related wind power production. Besides validation measures like arithmetic mean, standard deviation, relative difference between the data sets, temporal correlations, and seasonality of the variables, we also include comparison and validation of data distributions, hourly ramp rates, spatial correlation, analysis on the zero wind power events including extreme value analysis. Since this is the first paper to evaluate the wind resource estimates from NORA3 the focus is put on a detailed validation against observations. A comparison of NORA3 against the host data set (ERA5) is also conducted to document the improvement of the downscaling process. To our knowledge this is the first peer-review paper focusing on evaluation of simulated wind resource and wind power estimates against offshore observations in the North seas and adjacent ocean regions increasing the relevance of the present study.

We validate the NORA3 data set for wind power purposes using observational wind data from six offshore sites. Details regarding the model and observational data, in addition to the data processing routines are found in Sections 2.1 and 2.2. Chapter 2.3, 2.4, 2.5, and 2.6 describe the methods used. The result of the downscaling process of ERA-5 is quantified through a comparison between output data from NORA3 and ERA-5 in Section 3. In Chapter 4 we investigate how well NORA3 captures the statistical wind speed measures and the related distributions. We also study the model performance in terms of the wind speed ramp-rates (Section 4.1), spatial wind speed gradient (Section 4.2), and wind direction (Section 4.3). In addition, uncertainties related to observations sampled at large structures are discussed in (Section 4.4). After converting wind speed data to hourly wind power data, we examine the performance of NORA3 related to wind power climatology (Chapter 5), including wind power variables such as median production and capacity factor (CF). Revealing the wind power potential in an area also requires mapping the wind power intermittency and variability at different spatial and temporal scales. The ability of the model to capture wind power variability and intermittency are investigated using hourly wind power ramp-rates (Section 5.1) and long-term variability in CF (Ch. 5.2). In addition to temporal variability, we also consider the ability of NORA3 to capture the spatial co-variability between production sites (Section 5.3). It is crucial for a data set to reveal the length, duration, and total number of hours of zero wind power production, and NORA3's performance against these measures is discussed in Section 5.4. Moreover, we calculate and validate the maximum expected length of a zero-event occurring during the turbine lifetime (Section 5.5). In the last chapter (Chapter 6) we summarize the validation results.

Figure 1. The domain (red rectangle) covered by the HARMONIE-AROME simulation and the locations of the six sites used in verifying the NORA3 data set (red dots), in addition to the met mast located at Frøya. A zoom-plot of the positions and the names of the stations is also shown. Details for the sites are given in Table 1.

2 Data and method

2.1 Model data

NORA3 is obtained by high-resolution atmospheric dynamic downscaling of the state-of-the-art ERA-5 reanalysis data set from the ECMWF (Hersbach et al., 2020). ERA-5 covers the Earth in an approximately 31 x 31 km horizontal grid, providing hourly information in 137 vertical layers. The model used in the downscaling process is the nonhydrostatic, convection-permitting NWP model HARMONIE-AROME (H-A) (Cycle 40h1.2). Boundary values from ERA-5 are provided to the model every 6 h. Hourly[1] NORA3 output data is stored (some outputs are stored every third hour). The model domain in NORA3 encloses almost the entire northern part of the Atlantic Ocean (see Fig. 1), and the model runs with a horizontal resolution of 3 x 3 km, with the atmosphere divided into 65 vertical layers.

H-A is a high-resolution NWP model solving the fully compressible Euler equations using forward time integration on a non-staggered horizontal grid. H-A is used in short-range operational forecasting and research by many European weather services and research institutes (Seity et al., 2011; Bengtsson et al., 2017). The NORA3 data set is a hybrid between a hindcast and a reanalysis data set because of to the way the observations are treated in the model. The H-A model performs data assimilation of 2-m temperature and 2-m relative humidity. NORA3 is continiously being generated. When the model integration is finalized (in autumn 2021) the NORA3 data will cover the time period from 1979 to present, and will be regularly updated in the coming years when ERA-5 data becomes available. We will focus on the period 2004-2016 in this study due to the time coverage of the observational data. For further details on the model set-up and the NORA3 generation process see Haakenstad et al. (2021).

2.2 The observational data

The observations used in the verification of NORA3 are hourly wind observations [2] from five oil- and gas platforms (Ekofisk, Sleipner, Gullfaks C, Draugen, and Heidrun) and one met-mast (Fino1) (see Fig. 1 for the location of the sites and Table 1 for further site information), retrieved from the Norwegian Meteorological Institute. The observational data was quality checked prior to the validation of NORA3. For a detailed description of this quality check process see Solbrekke et al. (2020). In addition to the routine described in Solbrekke et al. (2020) we also exclude all records of zero-wind conditions (u=0) that are likely to be erroneous according to the following:

$$u_{obs}(i) = 0 \quad \wedge \quad u_{n3}(i) \geq \frac{5}{m}\frac{1}{n}\sum_{j=1}^{m}\sum_{i=1}^{n}|u_{obs}(i,j) - u_{n3}(i,j)| = 5\overline{MAD}, \tag{1}$$

where $u_{obs}(i,j)$ and $u_{n3}(i,j)$ are the observed and modeled wind speeds, respectively, at hour i for site j. n is the total number of hours; m is the total number of sites; and $\overline{MAD}$ is the mean absolute deviation between the observed and modeled wind speeds averaged over all sites. In other words, whenever the observed wind speed at hour i and site j is zero and the corresponding modeled wind speed exceeds $5\overline{MAD} = 7.2 \text{ ms}^{-1}$, the observed value at hour i is excluded from the time series for site j. This

---

[1]instantaneous values

[2]10 min average values provided at every hour

additional quality control leads to the exclusion of up to 5 h of observations per site, except at Heidrun which excludes 58 h of observations. For Heidrun, the removal of these erroneous records of zero-wind conditions (u=0) corresponds to an exclusion of approximately 0.035 % of the total data.

2.3 Wind speed interpolation

To avoid introducing additional uncertainties into the observational data set, we verify the wind variables from NORA3 at the wind sensor heights, ranging from 68-140 m.a.s.l., for each site (see "WSH" in Table 1 for the sensor heights). By contrast, the wind power verification is performed at a typical hub-height, at 100 m.a.s.l., to ensure the production estimate are comparable between sites. The interpolation of wind speed data to another height is usually done by either the *logarithmic law*, the *power law*, or a combination of the two methods (e.g. the *Deaves & Harris*). Gualtieri (2019) reviewed the three aforementioned methods for 96 different locations worldwide. He concluded that the power law was the most reliable and also the most frequently used extrapolation method. In addition, according Sill (1988) the usage of logarithmic law (log-law) is most suited near the surface. Despite the aforementioned results from Gualtieri and Sill we have compared the performance of the log law and the power law (with time varying power exponent) for the offshore sites. The result of the comparison show that the model bias using log-law is larger than using the power law method. Therefor, the interpolation of wind speed data to sensor height or hub-height is done using the power law relation (Emeis (2018)).

When using the power the interpolated wind speed is sensitive to the choice of the power law exponent α. Usually, α is assigned based on assumptions about atmospheric stability and surface roughness, both of which can introduce erroneous results. However, the data from NORA3 allows us to calculate α for each time step (i). Rearranging the power law relation, we get the following expression for the power law exponent α:

$$\alpha(i) = \frac{ln\frac{u(i)_{z2}}{u(i)_{z1}}}{ln\frac{z_2}{z_1}}, \tag{2}$$

where the height-subscripts $_1$ and $_2$ corresponds to the two layers within which the wind shear is calculated. The heights used to calculate α depend on the wind-sensor height (WSH) at the site in question: if WSH < 100 m.a.s.l. then α is calculated using NORA3 wind-shear between the two model layers $z_1 = 50$ m.a.s.l. and $z_2 = 100$ m.a.s.l.. If WSH >100 m.a.s.l. then alpha is calculated using the wind shear between $z_1 = 100$ m.a.s.l. and $z_2 = 250$ m.a.s.l.. The mean α for the whole time period for the six stations ranges from 0.05-0.08 between 50 and 100 m.a.s.l., and 0.03-0.06 between 100 and 250 m.s.a.l..

2.4 Normalized wind power

To ensure our validation results are as general as possible, and since the wind farm at each site is only imaginary and of unknown capacity, we use normalized power calculations $P_w(i) = \frac{P_w^T(i)}{P_w^{max}}$ to validate the wind power potential at each site

(Solbrekke et al., 2020). $P_w^T(i)$ is the produced wind power at each time step (i) for a given site, and P_w^{max} is the nameplate capacity. Hence, the normalized wind power $P_w(i)$ is defined as follows:

$$P_w(i) = \begin{cases} 0 & u(i) < u_{ci}, \\ \frac{u(i)^3 - u_{ci}^3}{u_r^3 - u_{ci}^3}, & u_{ci} \le u(i) < u_r, \\ 1, & u_r \le u(i) < u_{co}, \\ 0, & u_{co} \le u(i), \end{cases} \tag{3}$$

where $u(i)$ is the wind speed at hour i, $u_{ci} = 4$ ms^{-1} is the cut-in wind speed, $u_r = 13$ ms^{-1} is the rated wind speed, and $u_{co} = 25$ ms^{-1} is the cut-out wind speed. These numbers were retrieved from the SWT-6.0-154 turbines used in Hywind, Scotland - the first floating wind farm in the world (AG, 2011).

Table 1. Relevant information for the sites used in the validation of NORA3. "Abb" lists the site-name abbreviations. "Lat" and "Lon" are the latitude and longitude for the site locations, respectively. "WSH" (in meters above sea level) corresponds to the wind sensor height at each site. The sensor type is listed under "Sensor", and the data period for the available observations for each site is listed under "Data period". In addition, the percentage of valid observations is also shown under "Valid obs (%)".

					Site information		
Site	**Abb**	**Lat**	**Lon**	**WSH**	**Sensor**	**Data period**	**Valid obs (%)**
Fino1	f1	54.02	06.59	100	A100LK Cup-anemometer	01.01.2004-31.07.2009	95.8
Ekofisk	ek	56.52	03.22	68/102	Vaisala WMT703	01.01.2000-31.09.2016	85.3
Sleipner	sl	58.37	01.91	136	Gill Ultrasonic	01.01.2000-31.09.2016	83.4
Gullfaks C	gf	61.22	02.27	140	Gill Ultrasonic	01.01.2000-31.09.2016	80.2
Draugen	dr	64.35	07.78	78	Gill Ultrasonic	01.01.2002-31.09.2016	66.6
Heidrun	he	65.33	07.78	131	Gill Ultrasonic	01.01.2000-31.09.2016	84.6

2.5 Ramp rates

To validate the ability of NORA3 to capture the wind speed and wind power variability we calculate the ramp rates (R), defined as how much the wind speed (u) or wind power (P_w) changes during a time-increment τ (Milan et al., 2014):

$$R_{P_w}(i) = P_w(i) - P_w(i + \tau), \tag{4a}$$

$$R_u(i) = u(i) - u(i + \tau), \tag{4b}$$

setting τ to be 1, we validate the model performance on hourly ramp rates. To gain a general picture of the model performance in terms of how much the wind speed or wind power changes from one hour to the next we calculate the mean absolute ramp rate (MAR) for each site, for both the observational data and the modeled data. MAR is defined as follows:

$$MAR = \frac{1}{n}\sum_{i=1}^{n}|R(i)|, \tag{5}$$

160   where $R(i)$ is the ramp rate at hour i, and n is the total number of hours.

2.6   Zero-event duration using extreme value theory

A wind turbine has an expected life time of approximately 20 years. If the right steps are taken, the lifetime can be extended 15% -25 % depending on whether the structure is bottom-fixed or floating (Wiser et al., 2016). This means that the lifetime is expected to increase to 23-25 years. Therefore, determining the duration of long-lasting shutdowns expected to happen during
165   the lifetime of a turbine is important for estimating the levelized cost of energy (LCOE). The 25-year return value of the duration of a zero-event (a period of zero wind power production), the corresponding confidence interval, and the p-values are calculated from the observations and the model data using two statistical methods, "block maxima" (BM) in which the data is fitted to a generalized extreme value (GEV) distribution using yearly values of maximum zero-event duration, and "peak over threshold" (POT) in which the data is fitted to a generalized Pareto distribution (for more information see Smith
170   (2002)) using the 99th-percentile of zero-event duration (the highest 1 % of zero-event in terms of duration) as the selected threshold. We calculate the Kolmogorov-Smirnov p-value (KS_p) to test the null hypothesis. The null-hypothesis states: the empirical data is **not** drawn from the chosen data distribution (GEV or Pareto). Testing the null-hypothesis is done by the Kolmogorov–Smirnov statistic calculating the distance between the empirical and theoretical cumulative distributions. Hence, the cumulative distribution function from the BM data (POT data) is compared to the cumulative distribution function from
175   the GEV (Pareto) distribution. Thus, given a significance level of p=0.025, if the KS_p value is small ($KS_p < p$), the distance between the cumulative distributions is too large, and we can conclude that the empirical data (BM or POT) was sampled from a different population than the theoretical GEV or Pareto distribution with a probability of 1-p. If the result from the Kolmogorov-Smirnov test tell us that we cannot exclude the possibility that the data is drawn from either of the two data distributions (Gev or Praeto) we fit the observational based and model based maximum zero-event durations to GEV and
180   Pareto and find the corresponding 25-year return values for the five sites (Fino1 is excluded from the extreme value analysis due to the shorter time series: 2004-2009).

3   Comparison of NORA3 and ERA-5

The NORA3 wind estimates in 10 m.a.s.l. are extensively validated against observations and compared to the ERA-5 reanalysis in Haakenstad et al. (2021). Nevertheless, we compare the performance of NORA3 and ERA5 towards the observed wind speed
185   climatology to see the result of the downscaling process in the six wind sensor heights (68-140 m.a.s.l.).

Table 2. Seasonal average of the observed wind speed and the model deviation in percentage (%) for both NORA3 (n3) and ERA5 (e5). "DJF" corresponds to December-January-February, "MAM" is March-April-May, "JJA" is June-July-August, and "SON" is September-October-November.

| | Seasonal mean wind speed | | | | | | | | | | | |
| | DJF | | | MAM | | | JJA | | | SON | | |
Site	obs	n3 (%)	e5 (%)	obs	n3 (%)	e5 (%)	obs	n3 (%)	e5 (%)	obs	n3 (%)	e5 (%)
FINO1	11.27	**-1.3**	-2.7	9.75	**-2.5**	-4.5	8.41	**-2.6**	-3.1	10.63	**-2.3**	-2.5
Ekofisk	12.74	**-5.4**	-7.5	10.08	**-2.8**	-5.9	8.60	**-5.0**	(-7.5	11.40	**-5.2**	-6.9
Sleipner	13.85	**-9.2**	-11.3	10.75	**-8.1**	-11.3	8.98	**-9.0**	-11.6	12.33	**-8.9**	-10.9
Gullfaks C	13.39	**-6.4**	-9.9	10.53	**-6.7**	-10.3	9.09	**-7.9**	-11.3	11.92	**-5.9**	-8.8
Draugen	11.81	**-4.9**	-8.3	9.52	**-3.5**	-6.7	8.06	**-5.6**	-9.6	10.76	**-4.4**	-7.2
Heidrun	12.37	**-6.6**	-8.7	10.09	**-6.5**	-8.6	8.28	**-7.6**	-10.0	11.24	**-7.1**	-8.7
Average	12.57	**-5.6**	-8.1	10.12	**-5.0**	-7.9	8.57	**-6.3**	-8.9	11.38	**-5.6**	7.5

Table 3. Seasonal standard deviation of the observed wind speed and the model deviation in percentage (%) for both NORA3 (n3) and ERA5 (e5). "DJF" corresponds to December-January-February, "MAM" is March-April-May, "JJA" is June-July-August, and "SON" is September-October-November.

| | Seasonal wind speed standard deviation | | | | | | | | | | | |
| | DJF | | | MAM | | | JJA | | | SON | | |
Site	obs	n3 (%)	e5 (%)	obs	n3 (%)	e5 (%)	obs	n3 (%)	e5 (%)	obs	n3 (%)	e5 (%)
FINO1	5.32	**-3.6**	-6.6	4.37	**-3.7**	-9.2	3.84	**-2.6**	-5.7	4.60	**-2.2**	-5.9
Ekofisk	5.85	**-4.4**	-6.8	4.47	**-1.8**	-6.3	4.05	**-3.2**	-6.7	5.12	**-5.5**	-8.0
Sleipner	6.41	**-7.5**	-9.5	4.96	**-6.7**	-9.7	4.38	**-6.4**	-10.1	5.49	**-6.9**	-8.9
Gullfaks C	6.41	**-4.2**	-4.8	5.18	**-4.6**	-6.8	4.59	**-4.1**	-7.8	5.51	**-3.8**	-5.8
Draugen	5.88	**-6.1**	-7.3	5.48	**-8.2**	-11.3	4.50	**-9.3**	-13.8	5.70	**-6.7**	-9.7
Heidrun	5.94	**-8.4**	-11.4	5.33	**-7.5**	-10.1	4.27	**-8.0**	-11.2	5.65	**-8.3**	-11.7
Average	5.97	**-5.7**	-7.7	4.97	**-5.5**	-8.9	4.27	**-5.6**	-9.2	5.35	**-5.6**	8.3

The observed seasonal average and standard deviation of the wind speed are shown in Table 2 and Table 3, respectively. In addition, the tables also contain the relative difference (in percentage) between the observations and NORA3 (n3 (%)) and the observations and ERA5 (e5 (%)). Table 2 illustrate that the modeled average seasonal wind speeds from NORA3 are consistently closer to the observed values for all the seasons and for all the sites. The standard deviation (std) is here a measure of the variability in the wind speed. The seasonal variability is shown in Table 4. Compared to ERA5, NORA3 is consistently closer to the observed seasonal std for all the six sites.

Figure 2 shows a quantile-quantile plot (qq-plot) between the observed wind speed and modeled wind speed by NORA3 and ERA5. The qq-plot determines if the modeled and observed data sets are drawn from the same sample distribution. If the circles lie on the reference line the data sets comes from the same data distribution. For all the six sites the models perform best for the lowest wind speeds ($u \leq 10\ ms^{-1}$). For both models the deviation from the reference line ("ref line") increases with increasing wind speed percentile. Nevertheless, NORA3 is consistently closer to the reference line compared to ERA5, especially for wind speed exceeding a typical cut-off wind speed ($u \geq u_{co}$). A technical feature called "high wind ride through" enables the turbine to exploit more of the very strong wind speeds ($u \geq u_{co}$). In offshore areas higher winds are occurring more frequently. Therefor, the importance for a NWP model to correctly estimate these strong wind events increases. NORA3 outperforms ERA5 for these high wind speeds ($u \geq u_{co}$).

As illustrated in Fig. 2 the largest difference between the observations, NORA3 and ERA-5 is found for wind speeds exceeding a typical cut-out limit of $25\ ms^{-1}$ ($u \geq u_{co}$). Since the power production is terminated or at least reduced when $u \geq u_{co}$ we calculate the wind power capacity factors (CF) for the three data sets. This is done to see how the the models perform in terms of power production, where the strongest wind speeds are not influencing the result due the power production cut-out limit. Table 4 contains the CF for the observed data, NORA3, and ERA-5 for the six sites. NORA3 performs consistently better than ERA5, where NORA3 is on average 1.8 percentage point closer to the average observed CF-value compared to ERA-5.

The required rate of return when planning offshore wind projects is typically 5-10 %. A deficiency of 3 percentage point (approximately 6 % difference in the average power output) in the CF is a sizable error, and might be too large in terms of profitability. Nevertheless, this highlights the need for building up archives of different NWP simulations to be able to conduct informed uncertainty calculations for the power production in regions where observational data are limited. However, the comparison of CF between NORA3 and ERA5 shows that the ERA5-based CFs are on average 5 percentage point (approximately 10 % difference in the power output) lower than the observational based CFs. Hence, the improvements using NORA3 over ERA5 gives more realistic wind power profitability measures.

The validation of wind climatology in NORA3 and ERA-5 show that the downscaling of ERA-5 in the process of creating NORA3 has resulted in an improved wind resource data set. The remainder of this study will focus on the validation of NORA3 towards observed wind climatology.

4 Validation of NORA3 wind speed

Prior to exploiting NORA3 as a wind resource data set in the planning-phase of future offshore wind power installations the data set has to be validated and verified against observational data. We start with the validation of mean quantities and wind speed distributions. The most relevant wind speed measures can be seen in Table 5. Arithmetic mean (μ) and standard deviation (σ) are used as measures of the average wind speed and the corresponding variability. Mean wind speeds (μ) for the six sites lie within the interval $10 - 12\ \mathrm{ms}^{-1}$. For all the sites the observed mean wind speeds are higher than the wind speeds from NORA3, indicating that the model underestimates the mean wind speed. The largest difference can be seen for Sleipner where the observed mean wind speed is 8.9 % higher than the simulated wind speed. The wind speed at each site is highly variable,

Figure 2. Quantile-quantile plot between the observed wind speed (obs) and the modeled (mod) wind speed; NORA3 (red) end ERA5 (black), for all the six offshore sites.
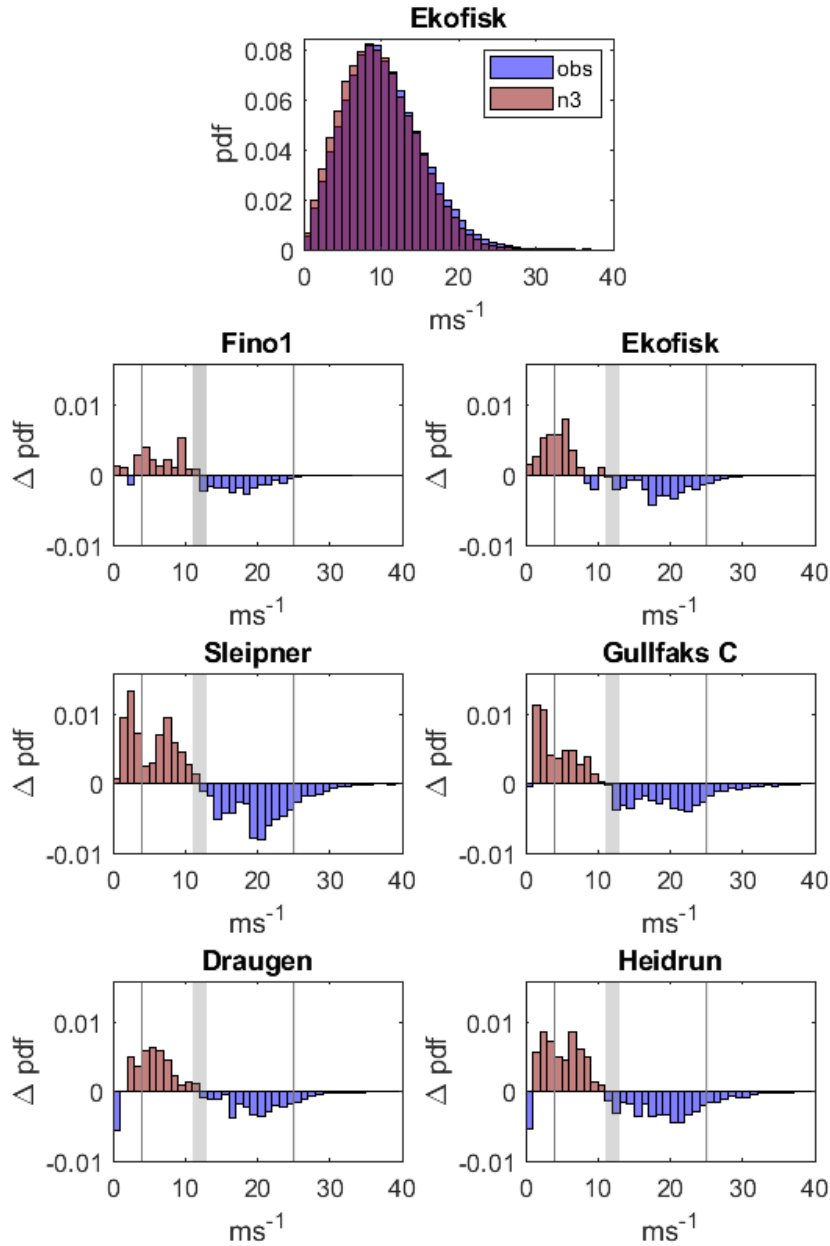
Figure 3. a) Example wind speed probability density function (pdf) (Ekofisk) for NORA3 (n3) in red and observations (obs) in blue. b)-h) Differences between NORA3 and observational wind speed probability density functions ($\Delta pdf = pdf_{mod} - pdf_{obs}$) for the six sites. When $\Delta pdf = 0.01$ the probability that the given wind speed will occur is 1 % higher in the model output. The large gray area corresponds to the range within which the rated wind speed usually falls. The gray vertical lines at the left and right mark the cut-in and cut-out wind speed limits used in this study, respectively.

Table 4. Capacity factor calculated from the observations (obs), NORA3 (n3) and ERA5 (e5) for the six sites. In addition, the difference (diff) between NORA3 and observations and between ERA5 and observations are also listed.

	Capacity factor		
Site	obs	n3 (diff)	e5 (diff)
FINO1	47.4	46.1 (**-1.3**)	45.2 (-2.2)
Ekofisk	51.2	48.6 (**-2.6**)	46.8 (-4.4)
Sleipner	54.7	49.6 (**-5.1**)	47.8 (-6.9)
Gullfaks C	53.4	49.5 (**-3.9**)	46.9 (-6.5)
Draugen	45.3	43.2 (**-2.1**)	40.8 (-4.5)
Heidrun	48.5	44.6 (**-3.9**)	43.1 (-5.4)
Average	**50.1**	**46.9 (-3.2)**	**45.1 (-5.0)**

225   with the std (σ) for the observations varying from $4.7 - 6.5 \mathrm{~ms}^{-1}$, where the model wind speed being slightly less variable (3-19 %). Hence, the observed wind speed is somewhat more intermittent and variable than the modeled wind speed, indicating that HARMONIE-AROME is missing some of the variability embedded in the wind field.

The Weibull scale parameter ("λ" in Table 5) indicates the height and width of the distribution. A larger scale parameter indicates a wider and lower probability distribution. All the observed scale parameters are slightly higher than the modeled;

230   the modeled scale parameters are on average 3.74 % lower than the observed. In other words, the observations contain more wind speed events at the tails of the Weibull distributions resulting in a larger scale parameter.

As all observed and modeled Weibull shape parameters ("k" in Table 5) are less than 2.6 the distributions are positively skewed, with a long tail to the right of the mean. The observed shape parameter is equal or smaller than the modeled (on average 7.3 % lower) indicating that the observed data is more positively skewed with a longer right tail, again emphasizing

235   that the observed data contain more high wind speed events than the NORA3 wind speed data.

According to Table 5 the model underestimates the wind speed at all sites. Since the wind power production is a function of the wind speed cubed the wind power is highly sensitive to systematic deviations between the observed and simulated wind speeds. However, the sensitivity varies with wind speed and is especially strong within the interval between cut-in and rated wind speeds. Fig. 3b-h shows the differences in the observed and modeled wind speed probability density functions (Δ pdf =

240   pdf mod - pdf obs) for the six sites, in addition to the wind speed distribution for Ekofisk (Fig. 3a). The main finding is that the model underestimates the number of events with high wind speed, and overestimates the number of events with low wind speed for all sites. The model is biased towards too few high-wind events and too many low-wind events than observed, and the transition occur near typical rated wind speeds ($11 - 13 \mathrm{~ms}^{-1}$) for state-of-the-art offshore wind turbines (the widest gray area in Fig. 3, panels b-h). This model-bias will have a large impact on the difference between the observed and modeled wind

245   power.

Table 5. Statistical measures of the wind speed for the observations (Obs) and the model (n3). μ and σ are the arithmetic mean and standard deviation, respectively. λ and k are the Weibull scale and shape parameters, respectively. The wind speed validation is performed at the sensor height to avoid uncertainties related to power-law extrapolation (see Table 1 for information on heights).

| | Wind speed | | | | | | | | | | |
| | μ | | | σ | | | λ | | k | |
Site	Obs	n3	%	Obs	n3	%	Obs	n3	Obs	n3
FINO1	9.96	9.72	**-2.41**	4.73	4.52	**-4.65**	11.24	10.97	2.27	2.27
Ekofisk	10.35	9.85	**-4.83**	4.97	4.75	**-4.42**	11.61	11.12	2.12	2.18
Sleipner	11.70	10.66	**-8.89**	5.83	5.42	**-7.03**	12.99	12.02	1.94	2.05
Gullfaks C	11.45	10.70	**-6.55**	5.90	5.66	**-4.07**	12.64	12.06	1.82	1.97
Draugen	9.87	9.44	**-4.36**	5.45	5.07	**-6.97**	10.95	10.65	1.75	1.94
Heidrun	10.56	9.87	**6.53**	5.67	5.22	**-7.94**	11.50	11.13	1.64	1.97
Average	**10.65**	**10.04**	**-5.60**	**5.43**	**5.11**	**-5.85**	**11.82**	**11.33**	**1.91**	**2.06**

4.1 Wind speed ramp rates

The wind speed ramp rate (ms^{-1}) is a measure of the hourly variability in the data set. In other words, the ramp rate quantifies how much the wind speed changes during 1 h. Figure 4 shows the distributions of observed and modeled hourly wind speed ramp rates for Ekofisk (the other sites have similar distributions). The distribution is wider for the observations than for the modeled data, illustrating that the observed wind speed change from one hour to the next is greater than that in the modeled wind speed data.

The mean absolute ramp rate (MAR) for the observed and modeled wind speed (u) is shown in Table 6. Typically observed MAR is around $1\ \mathrm{ms}^{-1}$ and the difference between modeled and observed ramp rates indicates that the model underestimates the variability in hourly wind speed by 30 %-36 %.

Table 6. Mean absolute ramp rate (MAR) for the observed and modeled wind speed (u). The difference between the modeled and observed MAR_u divided by the observed MAR_u, is given as a percentage (%).

| | MAR_u | | |
Site	Obs	n3	%
Fino1	0.98	0.67	**-31.27**
Ekofisk	1.04	0.67	**-35.58**
Sleipner	1.15	0.75	**-34.78**
Gullfaks C	1.15	0.81	**-29.57**
Draugen	1.31	0.85	**-35.11**
Heidrun	1.22	0.80	**-34.43**

Figure 4. The probability density distribution (pdf) of the modeled (n3) and observed (obs) hourly wind speed ramp rates.

4.2 Far offshore to coastal wind speed gradient

Table 7. The average change in wind speed at sensor height per 100 km kilometer for Heidrun-Draugen and Draugen-Frøya.

	Wind speed gradient		
Site	obs	n3	diff
Heidrun-Draugen	0.66	0.43	**-0.23**
Draugen-Frøya	1.52	1.68	**0.16**

An important feature of a model wind data set is the ability to properly estimate the horizontal wind speed gradient from far offshore to coastal areas. There are limited possibilities to investigate this using the available observational data. However, we made use of data from an observational met-mast situated on the coastal island of Frøya (see Fig. 1) to present some indicative results. Generally, using wind speed data at sensor height for the three sites Heidrun (far offshore), Draugen (near coastal) and Frøya (coastal) shows that there is no clear bias in the model (see Table 7). NORA3 underestimates the local far-offshore to near-coastal wind speed gradient, but slightly overestimates the near-coastal to coastal gradient.

Figure 5. Difference in the occurrence (%) of wind events categorized in different wind direction-intervals (30°-intervals) between NORA3 and observations (model-obs) for a) Ekofisk; b) Sleipner; c) Gullfaks C; d) Draugen; e) Heidrun. For each wind direction interval the wind events are divided into four different wind speed categories, the first one corresponds to u less than cut-in wind speed ($u < u_{ci}$), the second is the wind speed interval where the the wind power is a function of the wind speed cubed ($u_{ci} \leq u < u_r$), the third interval contain the wind speeds corresponding to rated wind power production ($u_r \leq u < u_{co}$), and the last interval is where the wind speeds are too strong resulting in a terminated wind power production ($u_{co} \leq u$).

4.3 Wind direction

Another important factor for planning a wind farm using simulated data is the quality of the modeled wind direction. State-of-the-art wind turbine technology allows the wind turbines to yaw to face the main wind direction. Mapping the wind direction climatology is important for the wind farm layout. Wind-rose plots (see Chapter A Fig. A1) demonstrate that the modeled and observed data in general show the same wind direction distributions, with only small differences, except for Fino1. Fino1 is excluded from the verification of wind direction because the wind rose for that site shows a clear directional disturbance, as the wind is affected by the observation mast (see Fig. A1, lower most row). Fig. 5 graphs the differences between the modeled and observed data (%) in the number of wind direction events (30°-intervals) for four wind speed categories ($u < u_{ci}, u_{ci} \leq u < u_r, u_r \leq u < u_{co}$, and $u_{co} \leq u$). There is no systematic bias in wind direction that can be seen across the sites, and the biases in frequency are less than 5% for all directional intervals and all sites. The wind speed interval with the greatest difference between the model and the observations features wind events corresponding to $u \geq u_{co}$. The wind speed interval with the smallest difference between the model and the observation are the "too low" wind events ($u < u_{ci}$). Hence, the model is better at capturing the wind direction when the wind speed is low.

Sleipner is the site with the greatest difference between model and observations for almost all wind direction intervals. The mismatch between the observed and modeled wind direction events for Sleipner is probably tied to the model performance. But, we cannot exclude the possibility that the platform design at Sleipner affects the flow field more than the design of the other platforms.

4.4 Uncertainties in observed wind speed

Working with observational data and numerical weather prediction models involves dealing with data that contains uncertainties and errors of known or unknown character. The majority of the observational sites used in this study (five of six sites) are oil- and gas platforms. The platforms are large structures that may influence the upcoming flow. On the other side, an observational mast may also influence the flow when the upcoming wind is guided to pass through the mast before being recorded by the sensor.

To what extent these large offshore structures influence the ambient flow field is unclear (Berge et al., 2009; Vasilyev et al., 2015; Furevik and Haakenstad, 2012). Nevertheless, using observations from oil- and gas platforms enable us to validate NORA3 over ocean areas where observational data is sparse. A potential alteration of the wind would be a function of the platform layout, the atmospheric stability, the upcoming wind direction, and the ambient wind speed. To investigate the distortion caused by these large structures, we compared wind speed data from the platforms with data from Fino1, and from the met mast at the Frøya field station.

The main results are illustrated and discussed in Appendix B (Figs. B1 and B2) and indicates that flow-disturbance by large oil- and gas platforms is to some extent visible in the wind speed data for some of the platforms. However, results illustrated in Fig. B2 and Fig. 5 show no clear signal indicating whether the large structures alter the flow field.

5 Comparison of estimated wind power from observed and modeled wind speed

295   Because the conversion from wind speed to wind power is non-linear (see Eq. 3), the wind power distribution differs greatly from the wind speed distribution. The statistical measures for the wind power are shown in Table 8. Median (q_{50}) and inter-quartile range (IQR) are independent of data distribution and are therefore good representations of the average wind power production and the related intermittency, respectively. All wind power estimates are calculated at a hub height of 100 m.a.s.l. using the wind interpolation method discussed in Section 2.3 and the normalized power curve described in Section 2.4.

300   Both the observation-based and model-based median wind power production estimates reveal very good wind power potential for the six sites (see Table 8). Nevertheless, since the model underestimates the wind speed events exceeding the rated wind speed, this partly counteracts the model's overestimation of the lower wind speed events ($u < u_r$), making the modeled average power production slightly underestimated. Therefore, the observation-based estimates of the median hourly power production q_{50} span from 0.3-0.5 (i.e. the median power production for a given hour would typically be 30 %-50 % of installed

305 capacity), compared to 0.3-0.4 for the model-based estimates. IQR, a measure of the variability, is the range between the first and third quartiles ($q_{75} - q_{25}$). Since the range of the normalized wind power is 0-1, IQR values close to 1 correspond to high variability, since almost the entire data range is present between the first and third quartiles. Hourly IQRs range from 0.86-0.95 for the observation-based estimates and from 0.80-0.94 for the model. There is no systematic difference between the IQRs of the model-based estimates and the observation-based estimates.

310   The capacity factor (CF) is another statistical measure quantifying the wind power potential. CF is here defined as the average wind power potential divided by the installed capacity. The observation-based estimates of CF vary between 46 % to 55 %, and the CF values from the model-based estimates are slightly smaller. The observation-based CF values exceed the modeled values by an average of 3 percentage point.

5.1 Wind power ramp rates

315 Figure 6 shows the distribution of observation-based and model-based hourly normalized wind power ramp rates for Ekofisk (the other sites have similar distributions). As for the distribution of hourly wind speed ramp rates, the distributions of hourly wind power ramp rates are wider for the observation-based ramp rates than for the model-based, illustrating that the hourly estimated wind power variability based on observations is greater than the estimated variability based on NORA3 data. The difference in MARs indicate an hour-to-hour variability typically of 7 %-9 % (Table 9) of the installed capacity based on

320 observations. In contrast, the variability for model-based estimates is 5 %-6 % and is underestimated at all sites.

5.2 Inter-annual and seasonal capacity factor

In addition to encompass short-term variations in wind speed and estimated power production it is essential for a model data set to contain the correct long-term variations. In this section we evaluate NORA3's ability to capture the longer-term climatic variability of the wind power potential for a given site. The inter-annual and seasonal variations in CF provide a good indication

325 of how NORA3 performs in terms of long-term wind power fluctuations.

Table 8. Statistical measures of the observation-based (Obs) and model-based (n3) wind power production. q_{50} is the hourly median production, IQR is the inter-quartile range of the hourly production, and CF is the wind power capacity factor. The wind power measures and estimates are performed at a typical hub-height of 100 m.a.s.l. using the interpolation of observed wind speeds as outlined in Section 2.3 and the power curve given in Section 2.4 for all the sites.

		Wind power							
	q_{50}			CF			IQR		
Site	Obs	n3	%	Obs	n3	n3-obs	Obs	n3	%
FINO1	0.37	0.36	**2.70**	47.55	45.85	**-1.70**	0.86	0.80	**-9.37**
Ekofisk	0.45	0.40	**-11.1**	51.02	48.39	**-2.63**	0.88	0.91	**3.41**
Sleipner	0.54	0.42	**-22.2**	54.82	49.60	**-5.22**	0.88	0.91	**3.41**
Gullfaks C	0.51	0.42	**-17.7**	53.30	49.54	**-3.77**	0.90	0.92	**2.22**
Draugen	0.33	0.30	**-9.10**	45.45	43.34	**-2.11**	0.95	0.91	**-4.21**
Heidrun	0.39	0.32	**-17.95**	48.33	44.67	**-3.66**	0.93	0.94	**1.08**
Average	**0.43**	**0.37**	**-13.95**	**50.08**	**46.89**	**-3.18**	**0.90**	**0.91**	**1.11**

Table 9. Mean absolute ramp rate (MAR_{P_w}) for the observation-based and model-based estimates of the wind power output. The difference between the modeled and observed MAR_{P_w} divided by the observed MAR_{P_w} is given in percentage (%).

	MAR_{P_w}		
Site	Obs	n3	%
Fino1	0.075	0.050	**-32.54**
Ekofisk	0.079	0.049	**-37.97**
Sleipner	0.077	0.051	**-33.77**
Gullfaks C	0.079	0.054	**-31.65**
Draugen	0.092	0.060	**-34.78**
Heidrun	0.084	0.055	**-34.52**

Fig. 7a and b illustrate the inter-annual and seasonal CF, respectively, from the observation-based estimates. In addition, the CF deviations (Δ CF) between the model-based estimates and the observation-based estimates are illustrated in Fig. 7c and d. The observed year-to-year variation in CF is substantial, varying up to 0.12 (12% of installed capacity) from one year to the next. Fig. 7c shows that the yearly CF values from the model are systematically lower than the observed CF values. This result is most pronounced for Sleipner, where the difference in Δ CF $\approx$ -5, meaning that the model-based CF is on average 5 percentage point lower than the observation-based CF.

The model's underestimation of CF can also be seen in the seasonal CF-values. Fig. 7d shows that Δ CF < 0 for all the sites. The underestimation of the seasonal CF values is largest during the summer months (May-September), meaning that

Figure 6. The probability density functions (pdfs) of the ramp rates for observation-based (obs) and model-based (n3) hourly normalized wind power.

the relative importance of the summer months in wind power production will be slightly underestimated in the model-based estimates.

5.3 Spatial wind power co-variability

Many studies have shown that interconnection of wind power production sites mitigates wind power intermittency (Kempton et al., 2010; Reichenberg et al., 2014; St. Martin et al., 2015; Reichenberg et al., 2017; Solbrekke et al., 2020). Therefore, simulated data sets for use in decision making about future wind power installations should be able to represent spatial and temporal co-variability between wind power sites.

Figure 8 illustrates the ability of NORA3 to capture the spatial co-variance in estimated hourly wind power production between the six sites. The figure demonstrates how the correlation between two sites changes as a function of the separation distance, both for the observation-based estimates (blue) and the model-based estimates (red). For almost all separation distances the model overestimates the correlation between two connected sites. The overestimation is generally small, but is greatest for small separation distances. This result indicates that NORA3 is better at capturing the large-scale spatial variability than variance on smaller scales.
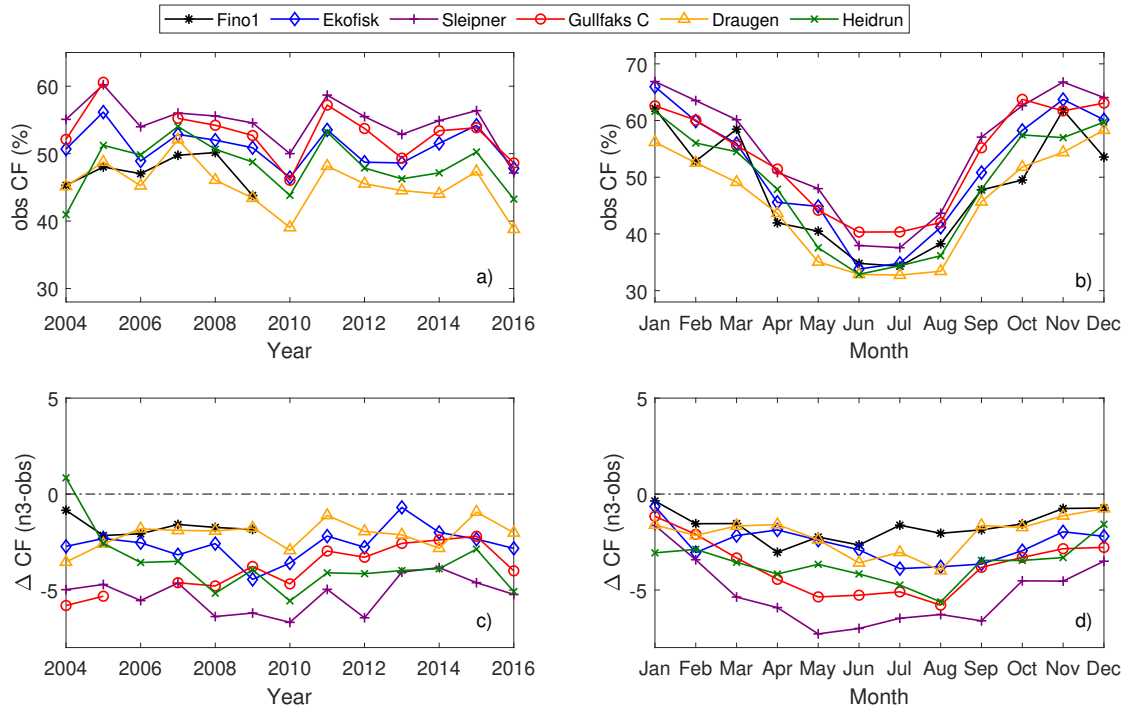
Figure 7. a) The observed inter-annual variation in the capacity factor (obs CF). b) Seasonal variation in CF for the observations. c) The difference in the inter-annual CF (Δ CF) between the model and the observations (n3-obs). d) The difference in the seasonal CF between the model and the observations (n3-obs). A specific year was excluded from the plot if more than one-half of the data for that year was missing.

A general description of the dependency between correlation and separation distance can give us information on the de-correlation length for the sites used in this study. Using the station-pair correlations we identify a best fitting exponential curve and a de-correlation length L (in kilometers). Connecting sites separated by a distance greater than the de-correlation length
350 ensures that the collective wind power intermittency from the two sites is substantially reduced compared to the intermittency from one of the sites. We use the e-folding distance[3] as a measure of the offshore de-correlation length L. The exponential curves and the corresponding de-correlation lengths for both the observations and NORA3 are presented in Fig. 8. The observation-based L is 426 km compared to a 508-km L based on NORA3. The model-based estimates indicate that to ensure relatively independent hourly power production, a greater interconnection distance is needed than that indicated by the observation-based
355 estimates.

Figure 8. Correlation of wind power time series as a function of the distance between the connected site-pairs for the observations (obs, blue) and NORA3 data (n3, red). An exponential fit is also shown (e^{bx^a}) for both the data sets with the corresponding de-correlation lengths, L.

5.4 Zero wind power events

Knowing about the risk, duration, and frequency of zero-events (periodes of zero wind power production) is important for decision-making and also in turbine maintenance planning, as these measures influence the levelized cost of energy and hence the decision-making process (Cory and Schwabe, 2009). A zero-event is caused by a wind speed that is too low ($u < u_{ci}$) or too high ($u \geq u_{co}$), and these events depend to some extent on the technical specifications of a wind turbine, but also, and more significantly, on the ambient wind climate in the area of interest. Table 10 shows the percentages of all hourly wind speed values that fall into each wind power category ($u < u_{ci}$, $u_{ci} \leq u < u_r$, $u_r \leq u < u_{co}$, and $u_{co} \leq u$) for each site. In addition, the table lists the total risk of having zero wind power production ($P_w = 0$). The percentage of hours when the wind is too weak to produce wind energy ($u < u_{ci}$) ranges from 8 % to 14% in the observation-based estimates and is overestimated by the model by an average of 1.6 percentage point for all sites. On the other hand, the observation-based estimates indicate that the fraction of hours in which the wind speed is too high ($u \geq u_{co}$) is about 0.2 %- 2 %, and the model underestimates this by approximately 0.6 percentage points. The model's overestimation of the number of hours with winds that are too weak to produce wind power and its underestimation of the number of hours with winds that are too strong results in a well-captured total numbers of hours of zero wind power production, which differs from the observed value by 1 percentage point.

---

[3]The distance where the correlation has dropped to $\frac{1}{e} = 0.37$.

Figure 9. a) Yearly occurrence and corresponding duration of observation-based zero-events caused by to wind speeds lower than cut-in wind speed ($u < u_{ci}$). b) The differences between model-based (n3) and observation-based (obs) zero-event occurrences divided by the total number of observed occurrences of "too low" wind speeds. Values over or under each bar correspond to the differences (n3-obs) in the number of yearly occurrences between the model and observations. Abbreviations: f1: Fino1; ek: Ekofisk; sl: Sleipner; gf: Gullfaks C; dr: Draugen; he: Heidrun.

Figure 10. a) Yearly occurrence and corresponding duration of observation-based zero-events caused by to wind speeds higher than cut-out wind speed ($u \geq u_{co}$). b) The differences between model-based (n3) and observation-based (obs) zero-event occurrences divided by the total number of observed occurrences of "too high" wind speeds. Values over or under each bar correspond to the differences (n3-obs) in the number of yearly occurrences between the model and observations. Abbreviations: f1: Fino1; ek: Ekofisk; sl: Sleipner; gf: Gullfaks C; dr: Draugen; he: Heidrun.

370 The atmospheric conditions causing winds that are too weak for wind power production are very different from those causing winds that are too strong. Therefore, we split the zero-events accordingly. Figs. 9 and 10 illustrate the ability of the NORA3 to capture the observation-based estimates of zero-events of different duration. Fig. 9a shows the observation-based numbers of zero-events of varying duration caused by too weak winds. As expected, the number of zero-events decreases as the duration of the events increases, ranging from around 90-130 yearly events lasting less than 3 h for most sites to close to zero such events

375 lasting longer than 2 days. Figure 9b graphs the relative differences (in percentage) between the NORA3 and observation-based estimates of the numbers of zero-events by duration. The model-based estimates typically have 40 %-50 % too few zero-events of short duration (1-3 h) compared to the observations. For longer zero-events the model is biased towards too many events.

The model's underestimation of short zero-events caused by too low wind speeds and its overestimation of longer zero-events occur as a result of the model having lower variability than the observations, as seen in the ramp-rate analysis (see Section

380 5.1). This lower variability means that when these zero-events occur in the model they tends to be of longer duration, but the frequency of such events is too low.

From Fig. 10a it is evident that the yearly average occurrence of zero-events caused by too strong winds is a factor of ten lower than the number of zero-events caused by winds that are too weak. Hence, one zero-event caused by too strong winds happens for approximately every 10 zero-events caused by too weak winds. The model underestimates the number of zero-

385 events caused by too strong winds for all sites (Fig. 10b); depending on the zero-event duration, NORA3 typically has 40 %-70 % too few zero-events caused by too strong winds.

Table 10. The percentages of observed wind speeds (Obs) and modeled wind speeds (n3) that fall into the following four categories: 1) the wind speed is less than the cut-in limit ($u < u_{ci}$), 2) the wind speed interval in which the wind power is a function of the cube of the wind speed ($u_{ci} \leq u < u_r$), 3) wind power production is rated ($u_r \leq u < u_{co}$), and 4) wind speed exceeds the cut-out limit ($u_{co} \leq u$). In addition, the total hours of zero wind power production ($P_w = 0$) divided by the total number of observations, are shows as a percentage.

Site	$u < u_{ci}$ Obs	n3	$u_{ci} \leq u < u_r$ Obs	n3	$u_r \leq u < u_{co}$ Obs	n3	$u \geq u_{co}$ Obs	n3	$P_w = 0$ (%) Obs	n3
Fino1	9.79	10.18	66.27	66.82	24.75	22.86	0.19	0.14	**9.98**	**10.32**
Ekofisk	8.88	10.42	63.45	64.87	27.07	24.42	0.60	0.29	**9.48**	**10.71**
Sleipner	7.52	10.62	54.08	57.67	36.48	30.85	1.92	0.86	**9.45**	**11.48**
Gullfaks C	8.96	11.52	53.98	55.72	35.07	31.53	2.00	1.23	**10.96**	**12.75**
Draugen	13.81	14.12	59.67	62.45	25.59	22.97	0.93	0.46	**14.75**	**14.58**
Heidrun	11.05	12.69	58.12	60.85	29.41	25.83	1.42	0.63	**12.52**	**13.32**
Average	**10.00**	**11.59**	**59.26**	**61.34**	**29.72**	**26.42**	**1.18**	**0.60**	**11.19**	**12.19**

5.5 Expected maximum zero-event duration over the turbine lifetime

In this section we attempt to validate the model's ability to provide reliable estimates of extremely long-lasting zero-events. This is done by estimating the 25-year return value for the duration of a zero-event (the typical length of a zero-event that statistically

Figure 11. The 25-year return value with the corresponding confidence interval of the maximum duration of a zero-event generated by fitting a generalized Pareto distribution to the POT-data (peak over threshold) using both observations (obs) and modeled data (n3) for each of the sites. Abbreviations: f1: Fino1; ek: Ekofisk; sl: Sleipner; gf: Gullfaks C; Draugen: Dr; he: Heidrun.

390 would occur at least once over a 25-year period) using the method outlined in Section 2.6. Using the Kolmogorov-Smirnov test, we cannot exclude the possibility that the BM-data and POT-data are drawn from a GEV distribution and a Pareto distribution, respectively. Thus, it is reasonable to fit the observation-based and model-based extreme zero-event duration estimates to these distributions and find the 25-year maximum expected zero-event duration.

   Figure 11 displays the results from fitting the Pareto distribution to the POT-data (the results fitting the BM-data to the GEV
395 distribution are similar). From the observed data the typical length of the longest zero-event expected to occur at least once during the lifetime of a turbine is on the order of 40-60 h, but a zero-event of more than five days cannot be ruled out. The uncertainty in the estimations make it difficult to judge which sites have the shortest and longest maximum zero-even duration. Using the model data, the estimates are typically longer than the observation-based estimates (not significant at the 2.5 % significance level for four of five sites); and are in line with the lower variability in the modeled hourly wind speed and wind
400 power as seen in the ramp-rate analysis (see Sections 4.1 and 5.1). In conclusion, using NORA3 to estimate extreme zero-event duration would lead to a conservative estimate of the return values, and the duration might be overestimated due to the lower variability in the model.

6 Summary

An in detail validation of NORA3 offshore wind resource and power production has been conducted for the time period 2004-2016. NORA3 is a new and freely available high-resolution (3 km) numerical mesoscale weather simulation data set from the Norwegian Meteorological Institute. The validation has been performed using observations from six offshore sites along the Norwegian continental shelf. In addition, the performance of NORA3 compared to the host reanalysis data set (ERA-5) has also been quantified. Through advanced statistical measures we have validated both the NORA3 wind resource and the related wind power production. Validation measures like arithmetic mean, standard deviation, relative difference between the data sets, temporal correlations, and seasonality of the variables have been conducted. In addition, we also include comparison and validation of hourly data distributions, hourly ramp rates, spatial correlation, and analysis on the zero wind power events including extreme value analysis. The general picture is that the NORA3 data is rather well suited for wind power estimates in the absence of in situ data. Nevertheless, there is a tendency towards the model generating slightly conservative estimates, and the results are summarized below.

The comparison between NORA3 and ERA-5 demonstrated that NORA3 outperforms ERA-5 in terms of mean and standard deviation of the wind speed climatology for all seasons and for all wind speed intervals, especially for the very strong winds ($u \geq u_{co}$). Since the very strong winds are not contributing to power production, the average power capacity factor (CF) was also compared. Again, NORA3 differs from the observation based CF by on average 3 percentage point compared to ERA-5's deficiency by 5 percentage point. The validation of wind climatology in NORA3 and ERA-5 show that the downscaling process resulted in an improved wind resource data set.

For all the six offshore sites NORA3 seems to be biased towards lower mean wind speeds ($u_{obs} = 10.59 \text{ ms}^{-1}$, $u_{n3} = 10.04 \text{ ms}^{-1}$). The differences in wind speed distribution between the observations and the model output reveal that the model underestimates the number of events with wind speed exceeding the rated wind speed and overestimates the number of events with wind speeds below the rated wind speed (see Fig. 3). The transition between over- and underestimation by the model occurs near a typical rated wind speed ($11 - 13 \text{ ms}^{-1}$). As the model underestimates the wind episodes above the rated wind speed, this partly counteracts the model's overestimation of low wind speeds, making the total modeled power production slightly underestimated.

NORA3 is also slightly biased towards less variable wind speeds on hourly timescales. Analysis of hourly wind speed ramp rates show that the hour-to-hour variability is typically slightly above 1 ms^{-1} while the model-based ramp rates are slightly below 1 ms^{-1}, resulting in an underestimation of wind speed ramp rates on the order of 30% (see Table 6).

Generally, estimates of wind power from NORA3 are biased towards too low median values ($P_{w,obs} = 0.43, p_{w,n3} = 0.37$) and wind power CF ($CF_{obs} = 50 \text{ %}$, $CF_{n3} = 47 \text{ %}$). The negative bias is a consistent feature seen in all years and for all months for all the six sites.

The wind power ramp-rate analysis shows that the hourly wind power variability of the NORA3-based estimates are too low. The observation-based wind speed variability leads to a corresponding wind power ramp rate that is typically 0.08 (8% of installed capacity), while the model-based ramp rate estimated is typically 0.05.

By interconnection of site-pairs we demonstrate that the spatial co-variability in estimated hourly wind power production between sites is slightly higher for the NORA3 data than for the observational data. Hence, the decorrelation length is estimated to be 16 % longer in the model-based estimates.

440   The estimation of the occurrence and duration of zero-events show a well-captured total risk of hourly zero-events (n3 = 12.19 %, obs = 11.19 % of the time). We split the zero-events into episodes of no wind power production caused by either too low ($u < u_{ci}$) or too high ($u \geq u_{co}$) wind speeds. For zero-events caused by winds that are too strong NORA3 underestimates the occurrence of zero-events for all durations. For winds that are too weak, NORA3 underestimates the number of short zero-events (1-3 h) but is biased towards an excess of zero-events with longer duration. As a result, when a zero-event occurs

445 in the NORA3 data, it tends to be of longer duration, but the frequency of such events is too low. This deviation from the observation-based zero-events is in line with the lower variability in hourly wind speeds seen in the ramp rate analysis (Section 4.1 and Section 5.1).

     In the extreme-value analysis we found that at least once during the lifetime of a turbine (25 years) a zero-power event is expected to last for 1 to 3 days, depending on the site in question (see Fig. 11). However, a zero-event lasting longer than 5 days

450 cannot be ruled out for some sites. Overall, the 25-year return values from NORA3 is somewhat conservative, with a tendency towards longer maximum zero-event duration than seen in the observation-based return values.

     To a large degree NORA3 resembles the climatological offshore wind resource and wind power characteristics seen in the observations. However, the model slightly underestimates the wind resource and power potential, and the hourly variability in the model output is lower than in the observations. These characteristics should be kept in mind when using the NORA3 data

455 set in the planning-phase of a future offshore wind farm.

Author contributions. IMS and AS conceptualized the overarching research goals of the study, in addition to conducting formal analysis regarding statistical and mathematical techniques and methods. HH is the creator of NORA3 data set and is responsible for the data resources. IMS was responsible for the data curation and creation of software code, validation, visualization of data, and preparation of the original draft, with contribution from AS. In addition, AS was responsible for the supervision. All the authors contributed to the review and editing

460 process of the paper.

Competing interests. No competing interests to declare.

Acknowledgements. Thanks to the Meteorological Institute, and a special thank to Magnar Reistad, Øyvind Breivik, and Ole Johan Aarnes for providing all the observational data used in this study and for interesting discussions.

Appendix A: Wind direction

465   See Fig. A1 for the observed and modelled wind rose plot.

Figure A1. Observed (left column) and modeled (right column) wind roses for the six sites. The colorbars show the wind speed intervals in meters per second. The upper row corresponds to the northern most site, Heidrun, followed by Draugen, and the last row is southern-most site Fino1.
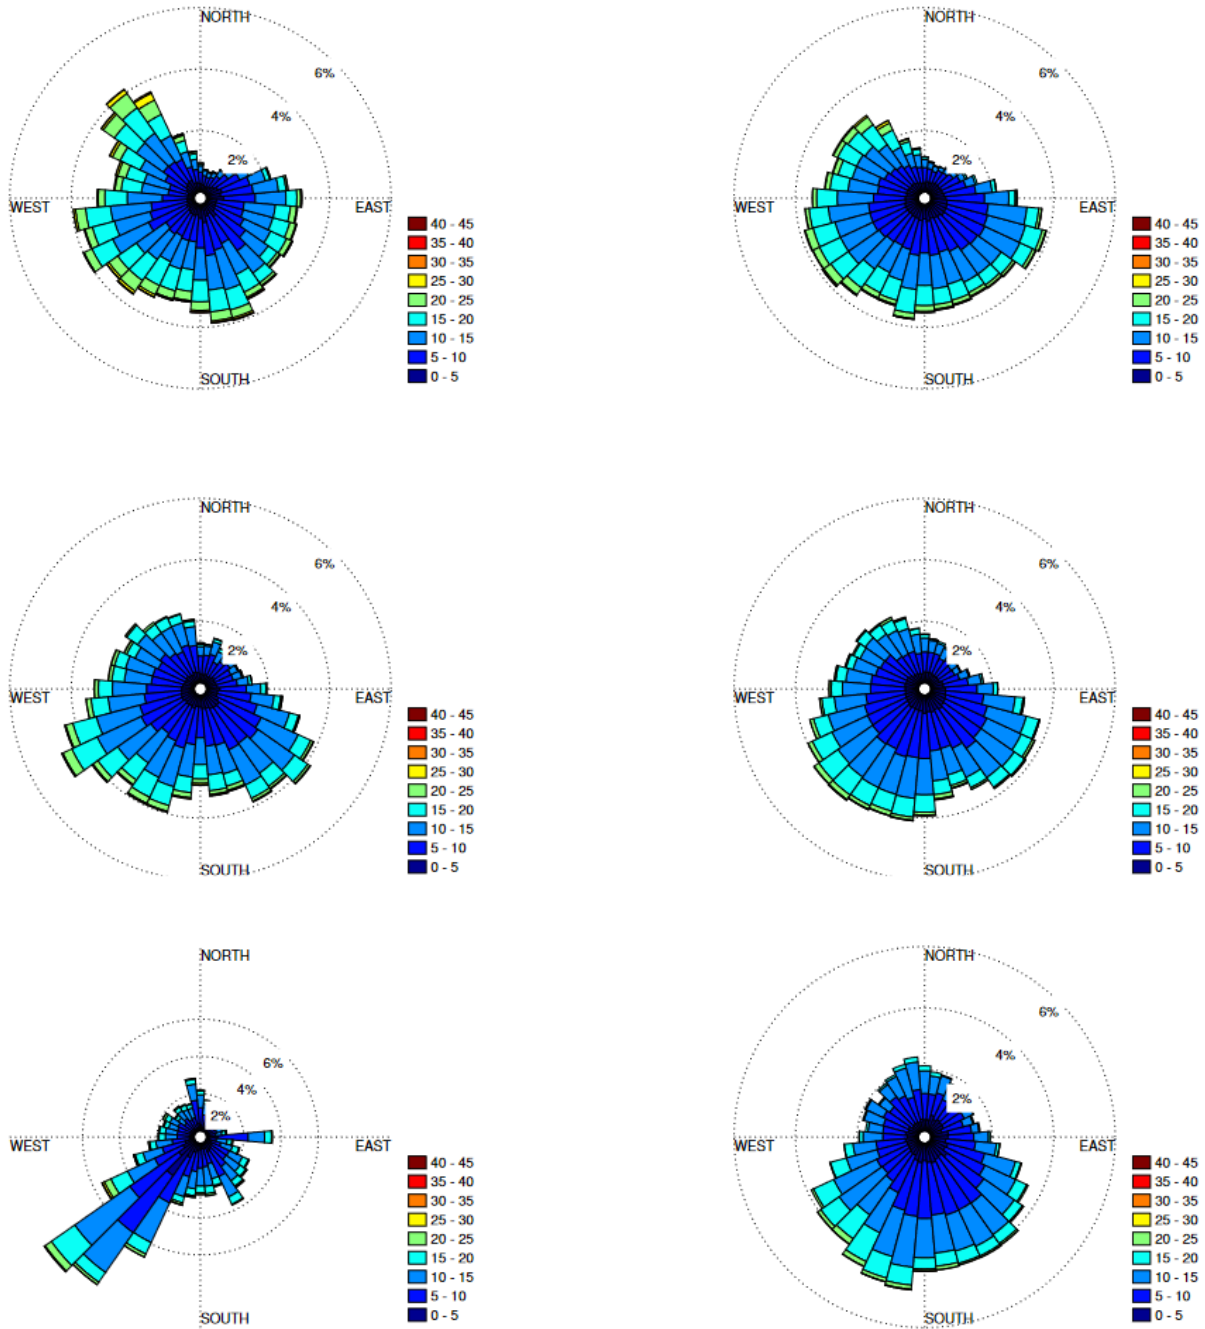
Figure A1.

Appendix B: Flow distortion by large structures

Figure B1. The wind speed difference (ms^{-1}) between the model and the observations (mod-obs) for the six sites. Dark green bars correspond to platform-sites (Ekofisk, Sleipner, Gullfaks C, Draugen and Heidrun) and light green to the met masts (Fino1 and Frøya).

To investigate a potential flow distortion effect in the observations caused by offshore platforms we sort the modeled wind speed data into categories according to the four wind power regimes: simulated wind speeds lower than cut-in wind speed $(u < u_{ci})$, simulated wind speeds falling between the cut-in and rated wind speed $(u_{ci} \leq u < u_r)$, simulated wind speeds corresponding to nameplate capacity $(u_r \leq u < u_{co})$, and simulated wind speeds exceeding the cut-out limit $(u_{co} \leq u)$. We then extract the observations for each hour that falls into each of the four categories, and determine the mean wind speeds for each category and for each site. The results are shown in Fig. B1. For the platform-sites, the observed wind speeds are greater than the modeled wind speeds in all four categories, with the largest difference seen in the rated wind regime $(u_r \leq u < u_{co})$. The finding that the observed wind speeds are greater than the modeled wind speeds for all four categories suggests a speed-up effect may be present for both weak and strong winds. For the two met masts, at Fino1 and Frøya, the model wind speed is weaker than the observed wind speed for the three first categories, but for the "too strong" wind regime $(u_{co} \leq u)$ the modeled wind speed is greater than the observed wind speed.

The five platforms have different layouts and therefore influence the upcoming flow field differently. Therefore, we calculate the difference in mean wind speed (n3-obs) for each wind direction interval (30°-interval) (see Fig B2) to investigate a potential flow distortion caused by the platforms. For almost all wind direction intervals the observed mean wind speeds are higher than

Figure B2. Differences in modeled (n3) and observed (obs) mean wind speeds (ms^{-1}) for each wind direction interval ($30°$-intervals). Abbreviations: ek: Ekofisk; sl: Sleipner; gf: Gullfaks C; dr: Draugen; he: Heidrun. Wind direction abbreviations intervals corresponds to the following: NNE = 0:29, NE = 30:59, NWW = 60:89, ESE = 90:119, SE = 120:149, SSE = 150:179, SSW = 180:209, SW = 210:239, WSW = 240:269, WNW = 270:299, NW = 300:329, NNW = 330:359.

the modeled speeds. For the Sleipner site, wind directions between 150:240 and 330:30 result in differences between the modeled and observed mean wind speeds of -2.5 ms^{-1} to -1 ms^{-1}, indicating some structural disturbance in the flow field.

References

AG, S.: Siemens 6.0 MW Offshore Wind Turbine, Tech. rep., 2011.

485 Bengtsson, L., Andrae, U., Aspelien, T., Batrak, Y., Calvo, J., Rooy, W. d., Gleeson, E., Hansen-Sass, B., Homleid, M., Hortal, M., Ivarsson, K. I., Lenderink, G., Niemelä, S., Nielsen, K. P., Onvlee, J., Rontu, L., Samuelsson, P., Muñoz, D. S., Subias, A., Tijm, S., Toll, V., Yang, X., and Køltzow, M. : The HARMONIE-AROME model configuration in the ALADIN-HIRLAM NWP system, Monthly Weather Review, 145, 1919–1935, https://doi.org/10.1175/MWR-D-16-0417.1, 2017.

Berge, E., Byrkjedal, , Ydersbond, Y., and Kindler, D.: Modelling of offshore wind resources. Comparison of a meso-scale model and 490 measurements from FINO 1 and North Sea oil rigs, in: European Wind Energy Conference and Exhibition 2009, EWEC 2009, vol. 4, pp. 2327–2334, 2009.

Byrkjedal, Ø. and Åkervik, E.: Vindkart for Norge, Tech. rep., Kjeller Vindteknikk, 2009.

Byrkjedal, Ø., Harstveit, K., Kravik, R., Løvholm, A. L., and Berge, E.: Analyser av offshore modellsimuleringer av vind, Tech. rep., Kjeller Vindteknikk, 2010.

495 Cory, K. and Schwabe, P.: Wind Levelized Cost of Energy: A comparison of Technical and Financing Input Variables, Tech. rep., National Renewable Energy Laboratory, 2009.

Emeis, S.: Wind Energy Meteorology, Springer, 2 edn., 2018.

Furevik, B. R. and Haakenstad, H.: Near-surface marine wind profiles from rawinsonde and NORA10 hindcast, Journal of Geophysical Research Atmospheres, 117, 1–14, https://doi.org/10.1029/2012JD018523, 2012.

500 Gualtieri, G.: A comprehensive review on wind resource extrapolation models applied in wind energy, Renewable and Sustainable Energy Reviews, 102, 215–233, https://doi.org/https://doi.org/10.1016/j.rser.2018.12.015, 2019.

Haakenstad, H., Breivik, Ø., Furevik, B. R., Reistad, M., Bohlinger, P., and Aarsnes, O. J.: NORA3: A non-hydrostatic high-resolution hindcast for the North Sea, the Norwegian Sea and the Barents Sea, Submitted to Journal of Applied Meteorology and Climatology, https://folk.uib.no/iso056/Paper_NORA3_Haakenstad_etal.pdf, 2021.

505 Hasager, C. B., Hahmann, A. N., Ahsbahs, T., Karagali, I., Sile, T., Badger, M., and Mann, J.: Europe's offshore winds assessed with synthetic aperture radar, ASCAT and WRF, Wind Energy Science, 5, 375–390, https://doi.org/10.5194/wes-5-375-2020, 2020.

Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J., 510 Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thépaut, J. N.: The ERA5 global reanalysis, Quarterly Journal of the Royal Meteorological Society, 146, 1999–2049, https://doi.org/10.1002/qj.3803, 2020.

Kempton, W., Pimenta, F. M., Veron, D. E., and Colle, B. A.: Electric power from offshore wind via synoptic-scale interconnection, Proceedings of the National Academy of Sciences of the United States of America, 107, 7240–7245, https://doi.org/10.1073/pnas.0909075107, 515 2010.

Milan, P., Morales, A., Wächter, M., and Peinke, J.: Wind Energy: A Turbulent, Intermittent Resource, pp. 73–78, https://doi.org/10.1007/978-3-642-54696-9_11, 2014.

NEWA: The New European Wind Atlas, https://www.neweuropeanwindatlas.eu/mission, 2020.

Regjeringen: Opening sites at the Norwegian continental shelf for wind power production, https://www.regjeringen.no/no/aktuelt/opner-omrader/id2705986/, 2020.

Reichenberg, L., Johnsson, F., and Odenberger, M.: Dampening variations in wind power generation-The effect of optimizing geographic location of generating sites, Wind Energy, 17, 1631–1643, https://doi.org/10.1002/we.1657, 2014.

Reichenberg, L., Wojciechowski, A., Hedenus, F., and Johnsson, F.: Geographic aggregation of wind power—an optimization methodology for avoiding low outputs, Wind Energy, 20, 19–32, https://doi.org/10.1002/we.1987, 2017.

Seity, Y., Malardel, S., Hello, G., Bénard, P., Bouttier, F., Lac, C., and Masson, V.: The AROME-France convective-scale operational model, Monthly Weather Review, 139, 976–991, https://doi.org/10.1175/2010MWR3425.1, 2011.

Sill, B. L.: Turbulent boundary layer profiles over uniform rough surfaces, Journal of wind engineering and industrial aerodynamics, 31, 147–163, 1988.

Skeie, P., Steinskog, D. J., and Näs, J.: Kraftproduksjon og vindforhold, Tech. rep., StormGeo AS, 2012.

Smith, E. P.: An Introduction to Statistical Modeling of Extreme Values, https://doi.org/10.1198/tech.2002.s73, 2002.

Solbrekke, I. M., Kvamstø, N. G., and Sorteberg, A.: Mitigation of offshore wind power intermittency by interconnection of production sites, Wind Energy Science, 5, 1663–1678, https://doi.org/10.5194/wes-5-1663-2020, 2020.

St. Martin, C. M., Lundquist, J. K., and Handschy, M. A.: Variability of interconnected wind plants: Correlation length and its dependence on variability time scale, Environmental Research Letters, 10, 44 004, https://doi.org/10.1088/1748-9326/10/4/044004, http://dx.doi.org/10.1088/1748-9326/10/4/044004, 2015.

Vasilyev, L., Christakos, K., and Hannafious, B.: Treating wind measurements influenced by offshore structures with CFD methods, in: Energy Procedia, https://doi.org/10.1016/j.egypro.2015.11.425, 2015.

Wiser, R., Jemmi, K., Seel, J., Baker, E., Hand, M., Lantz, E., and Smith, A.: Forecasting Wind Energy Costs and Cost Drivers: the Views of the World' S Leading Experts, Tech. rep., IEA Wind, NREL, Berkeley Lab, file:///C:/Users/iso056/AppData/Local/Temp/Ade26H1RQxa0ee2avOyS_lbnl-1005717.pdf, 2016.

Zheng, C. W., Li, C. Y., Pan, J., Liu, M. Y., and Xia, L. L.: An overview of global ocean wind energy resource evaluations, Renewable and Sustainable Energy Reviews, 53, 1240–1251, https://doi.org/10.1016/j.rser.2015.09.063, http://dx.doi.org/10.1016/j.rser.2015.09.063, 2016.