



Modular deep learning approach for wind farm power forecasting and wake loss prediction

Stijn Ally^{1,2}, Timothy Verstraeten^{1,2}, Pieter-Jan Daems², Ann Nowé¹, and Jan Helsen²

¹Artificial Intelligence Lab Brussels, Vrije Universiteit Brussel, Pleinlaan 9, 1050 Elsene, Belgium

²Faculty of Engineering, Vrije Universiteit Brussel, Pleinlaan 2, 1050 Elsene, Belgium

Correspondence: Stijn Ally (stijn.ally@vub.be), Timothy Verstraeten (Timothy.Verstraeten@vub.be), Pieter-Jan Daems (Pieter-Jan.Daems@vub.be), Ann Nowé (ann.nowe@ai.vub.ac.be), and Jan Helsen (Jan.Helsen@vub.be)

Abstract. Power production of offshore wind farms depends on many parameters and is significantly affected by wake losses. Due to the variability of wind power and its rapidly increasing share in the total energy mix, accurate forecasting of the power production of a wind farm becomes increasingly important. This paper presents a novel data-driven methodology to construct a fast and accurate wind farm power model. The deep learning model is not limited to steady-state situations, but captures also the influence of temporal wind dynamics and the farm power controller on the power production of the wind farm. With a multi-component pipeline, multiple weather forecasts of meteorological forecast providers are incorporated to generate farm power forecasts over multiple time horizons. Furthermore, in conjunction with a data-driven turbine power model, the wind farm model can be used also to predict the wake losses. The proposed methodology includes a quantification of the prediction uncertainty, which is important for trading and power control applications. A key advantage of the data-driven approach is the high prediction speed compared to physics-based methods, such that it can be employed for applications where faster than real-time power forecasting is required. It is shown that accuracy of the proposed power prediction model is better than for some baseline machine learning models. The methodology is demonstrated for two large real-world offshore wind farms located within the Belgian-Dutch wind farm cluster in the North Sea.

1 Introduction

Over the last decades, renewables have been expanding quickly, but the global energy crisis has kicked them into a new phase of even faster growth. Global wind capacity is expected to almost double the upcoming five years, with offshore projects accounting for one-fifth of the growth (IEA, 2022).

Due to the rapidly increasing share of wind power in the total energy mix and its variable and intermittent nature, accurate forecasting of the power production of wind farms becomes increasingly important for wind farm operators, balancing responsible parties (BRP) and transmission system operators (TSO). In addition, as wind turbines and wind farms are getting larger in scale, their influence on the wind flow and the resulting wake effects are becoming more pronounced.

The energy production of a wind farm is influenced by many factors (Lee and Fields, 2020). First of all, the energy production of an individual *wind turbine* depends on many parameters. The type of the turbine (with its specific mechanical and electrical design and control systems) and the wind speed are the most important parameters. Also wind turbulence and air density have



25 an important influence on the power production. Furthermore, a turbine can be derated due to technical reasons or be limited in power by a wind farm power controller.

Secondly, the energy production of a *wind farm* can be significantly reduced by wake losses. Wind turbines modify the air flow for downstream turbines, resulting in velocity deficits and increased turbulence. Reduced velocity of the air results in power losses, whereas increased turbulence results in a faster recovery of the velocity deficit (Sanderse et al., 2011). The magnitude of wake losses depends mainly on the layout of the wind farm (distance between the turbines, orientation of turbine rows and length of the cross-section of the farm) and the wind direction (Barthelmie et al., 2009), but also on other parameters, such as the wind turbulence intensity, atmospheric stability and surface roughness. The power loss of a downstream turbine can reach 40% in full-wake conditions (Barthelmie et al., 2009). When averaged over all different wind directions and wind speeds, annual power losses of a complete farm due to wake can range between 5% to 20%. Studies indicate that active wake control by applying yaw and/or induction control on the turbines may reduce wake losses and hence increase the power production of the farm (Boersma et al., 2017; Verstraeten et al., 2021). However, validation of the possible power gain under real-case conditions in a wind farm is still an active field of research (Bossanyi and Ruisi, 2021; Fleming et al., 2016, 2020).

Thirdly, the characteristics of the *wind inflow* into a wind farm (wind speed, wind direction, turbulence intensity, air density) are influenced by the environment surrounding the wind farm. Upstream wind farms may cause a reduction of the wind speed and an increase of the turbulence intensity (Porté-Agel et al., 2020; Pettas et al., 2021). Wind blockage by neighbouring farms and the wind farm itself may also reduce the inflow speed and deflect the inflow direction (Porté-Agel et al., 2020; Bleeg et al., 2018; Strickland et al., 2022). Also the proximity of coast lines can exert a significant influence on the wind characteristics, primarily attributable to differences in roughness and heat capacity between sea and land (Van Der Laan et al., 2017).

Modeling wind and wake flows within wind farms is challenging and since decades an active field of research. In addition to affecting power production, wake turbulence also impacts the loading of the turbines (Nejad et al., 2022). A wide range of different models have been developed. Based on the amount of detail they capture, they can be classified in three classes: low fidelity, medium fidelity and high fidelity models. Each of these types of models have some advantages and drawbacks. Low fidelity models describe only the dominant wake characteristics and are mostly limited to steady state simulations and homogeneous 2D wind inflows. These models are usually relatively fast (order of seconds on PC for a steady state simulation), but need tuning of some hyper-parameters and have a lower accuracy (Jensen Park model (Jensen, 1983), FLOW Redirection and Induction in Steady State (FLORIS) (NREL, b), curled wake model (Martínez-Tossas et al., 2019, 2021) and TurbOPark (Nygaard et al., 2020)). In recent years, some further developments have made it possible to model also some heterogeneous and dynamic environmental conditions (e.g. FLORIDyn (Gebraad and Van Wingerden, 2014; Becker et al., 2022), UFloris (Foloppe et al., 2022)). At the other side of the spectrum, high fidelity models, based on the 3D Navier-Stokes equations, describe flows in high detail. Large-eddy simulations (LES) resolve these equations on a coarse mesh and approximate smaller scale eddies with subgrid models (Sanderse et al., 2011). The main drawback of this type of models is the high computing load (order of days on a computing cluster). Some examples of high fidelity simulators are SOWFA (NREL, c), PALM (University of Hannover) and SP-Wind (Sood et al., 2022). Finally, medium fidelity models (such as DWM (Larsen and Binkgoel, 2007), FAST.Farm (NREL, a), WakeFarm (Schepers, 1998)) are based on simplifications of the Navier-Stokes equations. Despite, the



60 simplifications, the computing time for medium fidelity models remains significant (order of minutes on PC) (Boersma et al., 2017).

In contrast to physics-based models, data-driven techniques are not based on prior knowledge about the physical behaviour of the turbines or the air flow. Instead, they focus on fitting a general model on data. Recent developments related to deep learning resulted in a significant leap forward in the modelling of large and complex data sets (LeCun et al., 2015).

65 *Wind farm operators* have usually huge amounts of historical Supervisory Control and Data Acquisition (SCADA) data acquired by the instrumentation on their turbines. This data can be leveraged by machine learning techniques (Verstraeten et al., 2019). Taking advantage of this large amount of detailed information, accurate models of the wind farm and turbines can be built depending solely on this farm-specific data, and not on any other, often less specific or less accurate, data sources such as theoretical turbine power curves and weather forecast data.

70 *Commercial weather forecast services* provide nowadays weather forecasts for specific geographical locations. However, forecasts by distinct providers can differ among each other significantly, as they can be based on different weather models and data. In addition, although they provide forecasts for a specific geographical location such as the position of an offshore wind farm, they typically do not take into account the influence of the immediate surroundings of the wind farm, such as the presence of neighbouring offshore wind farms and the influence of coast lines. Moreover, they may not be able to provide all weather
75 parameters that are required as inputs for an accurate wind farm power model, and the provided variables may be calibrated differently than the instrumentation data used during the training phase of the data-driven model.

The *problem statement* that this paper addresses, is to construct a fast and accurate farm power forecasting model, which captures the temporal dynamics of the wind inflow as well as the behaviour of the farm power controller. A deep learning
80 now-casting power model is trained solely with SCADA data from the wind farm itself. Therefore, this farm power model is unique and independent from any weather forecasting data. In a multi-component pipeline, this single farm power model can be interfaced with multiple weather forecast services (from third-party providers) to forecast the wind farm power over multiple time horizons. The proposed model is capable of forecasting the farm power in a few milliseconds on PC, and this with an accuracy that surpasses the accuracy of other state-of-the-art data-driven models. This demonstrates that the proposed
85 farm power forecasting model can be used for applications that require both faster than real-time farm power forecasting and simulation of the farm power controller.

This work *distinguishes* itself from other publications about data-driven wind farm power forecasting on several aspects.

In literature, many models can be found for the power prediction of an individual turbine (Perez-Sanjines et al., 2022;
90 Zehtabiyani-Rezaie et al., 2023; Ti et al., 2021; Kisvari et al., 2021; Lin and Liu, 2020) (whether or not superposed afterwards into a farm). The model proposed in this paper predicts the power of a complete wind farm as a whole.

Whereas wind farm operators possess typically huge amounts of SCADA data, due to confidentiality reasons, such data is rarely accessible to the research community. Consequently, in many publications about wind farm power and wake modeling, simulation data is generated with physics-based models (such as the low-fidelity static model Floris (Yin and Zhao, 2019; Park



95 and Park, 2019) or mid-fidelity models (Zehtabiyan-Rezaie et al., 2023; Ti et al., 2021)), often only for relatively small virtual wind farms and for a limited set of wind conditions. In contrast, the data used in this paper is SCADA data from two large real-world offshore wind farms.

Many data-driven models in literature have only a limited set of input features (e.g. restricted to wind speed (Perez-Sanjines et al., 2022) and direction or historic wind power series (Wang et al., 2017)), whereas the wind farm power model presented
100 in this paper includes many more input parameters, all having a direct physical influence on the farm power (i.a. air density, turbulence intensity, wind direction variance, wind farm power setpoint and power limitations due to technical reasons). It is shown that adding these additional input parameters, improves significantly the accuracy of the model.

Data-driven wind farm power models found in literature are based often on data with a relatively coarse time resolution (10 minutes or 1 hour) (Liu et al., 2021; Kisvari et al., 2021; Wang et al., 2021). However, in this work, the wind farm power
105 prediction model is trained with time-series of 1-minute SCADA data in order to be able to capture the temporal dynamics of the inflow wind and the effect of setpoint changes of the farm power controller (of which 1-minute SCADA data was available for this work). In particular, the influence of inflow wind speed variations on the farm power production is analysed in this work.

A multi-horizon data-driven wind power forecasting method based on time series forecasting is proposed by Pombo et al. (Pombo et al., 2021). This method results only in good predictions for short time horizon forecasting. In order to obtain
110 long-term multi-horizon forecasts, the multi-component pipeline proposed in this work incorporates weather forecast data for multiple time horizons as a separate component, without modification of the now-casting farm power model trained solely with SCADA data from the wind farm. Another advantage of the multi-component pipeline, is that the sensitivity to multiple physical input features can be analysed. This is not possible for a model that consists of one single black box. Sensitivity plots
115 can be interpreted easily by wind energy professionals and increase the interpretability and explainability of the model.

In contrast to most other publications about farm power forecasting, the methodology proposed in this paper, includes also a quantification of the prediction uncertainty. Indeed, not only the accuracy of the model is essential for trading applications and to ensure reliable operation of wind farms being safety-critical systems, but insight in the uncertainty of the power predictions is crucial as well (Meyers et al., 2022; Braun et al., 2024).

120 In addition to farm power forecasting, in this work it is demonstrated also how farm-internal and farm-external power losses can be identified based on machine learning (ML) models.

The paper is organized as follows. First, the modular data-driven methodology is described in Section 2. Then, the results of the methodology applied to two offshore wind farms are presented and discussed in Section 3. Finally, the main conclusions
125 of the paper are summarised in Section 4.



2 Methodology

The proposed modular data-driven methodology is based on data from multiple data sources and integrates several deep learning models with each other. The used data sources are described in section 2.1 and the machine learning models are described in section 2.2.

130 2.1 Data

The proposed data-driven approach is based on multiple data sources: SCADA data from the wind turbines, from the wind farm power controller and from a weather station located in the wind farm, as well as weather forecasts from multiple weather forecast providers and ECMWF reanalysis v5 (ERA5) data. Typically, such data sources have different time resolutions and accuracy levels. For the wind farms used as examples in this paper, this data has respectively a 1-second, 1-minute, 10-minute, 135 1-hour and 1-hour time resolution. Measurement data from local measurements is typically more accurate than forecast or reanalysis data. Indeed, the accuracy of the former depends solely on the accuracy of the measurement instruments, whereas forecasts and reanalysis data depend on large-scale models and observations with a wide spacial grid that may be wider than a complete wind farm.

The optimal time resolution for a wind farm power model depends, on the one hand, on the purpose of the model (e.g. 140 which effects are to be captured by the model) and, on the other hand, may be limited by the available computing hardware and required prediction speed. In order to capture the dynamics of a varying wind propagating through the wind farm, the time resolution of the model has to be sufficiently shorter than the duration needed by the wind to cross the entire wind farm. For example, for an offshore wind farm with a cross-section length of 10 km, with wind speeds between turbine cut-in and cut-out wind speeds of respectively 4 m/s and 30 m/s, this duration is between 6 to 42 minutes. So, for that purpose, a sampling time of 145 1 minute should be adequate. In addition, the farm power set point data available for this work has a 1-minute time resolution, and so setpoint changes of the farm power controller can be captured as well.

In sections §2.1.1 to §2.1.4, each of the data sources used in the proposed methodology is described in more detail.

2.1.1 Turbine SCADA data

The turbine SCADA data of the wind farms that are used as showcases in this paper, has a sampling time of 1 second. For both 150 farms, the data of the turbines comprise the following measurements:

- wind speed (measured by an anemometer located on top of the turbine nacelle),
- wind direction (measured by a wind vane located on top of the turbine nacelle), and
- turbine active power (measured at the power terminals of the turbine).

For one of the two farms, an additional data field is available that expresses the maximum power that the turbine could technically produce at the moment in case of sufficient wind: 155



- turbine active power capability.

The maximum power can indeed be limited below the rated turbine power due to a technical problem of the turbine or by a curtailment imposed from externally (e.g. by the farm power controller).

Based on the wind speed and direction measurement data, two additional data features are built that can be used as measures
 160 for the wind turbulence: wind turbulence intensity ti and wind direction variance wd_{var} , as:

$$ti = \frac{\sigma(ws)}{\mu(ws)} \quad (1)$$

$$wd_{var} = \sigma^2(wd_{lat}) + \sigma^2(wd_{lon}) \quad (2)$$

with $\sigma(ws)$ the standard deviation of the wind speed, $\mu(ws)$ the average wind speed, and $\sigma^2(wd_{lat})$ and $\sigma^2(wd_{lon})$ the variance
 of respectively the lateral and longitudinal component of the wind direction, each during the 10-minute time interval centered
 165 around the 1-second data point.

As the wind farm and the turbine power models proposed in this paper have a temporal resolution of 1 minute, each of the
 above 1-second data feature sequences is averaged to 1-minute data blocks.

2.1.2 Wind farm data

Most of the farm data can be calculated by aggregating the SCADA data from the individual turbines. The active power
 170 produced by the farm is calculated as the sum of the individual turbine active powers. Similarly, the farm active power capability
 is calculated by summing up the power capability of each of the turbines.

Simply adding or averaging the wind measurement data of all individual turbines, however, would lead to a loss of infor-
 mation about the spatial variation of these features throughout the farm. Instead, the characteristics of the wind flowing into
 the wind farm are used, measured by a subset of turbines located in the upstream part of the farm. This allows also to isolate
 175 farm-internal wake and farm-external wake in separate ML models (the latter depending also on the operational status of the
 neighbouring farms). The set of upstream turbines is determined by the following algorithm (see Figure 1):

1. calculate wind vector \vec{W} as average of the wind speeds and directions measured by all the turbines of the farm
2. based on the farm layout, determine the projection of the position of each turbine on \vec{W}
3. select all turbines that are located in the most upstream zone of the farm with length D along \vec{W}
- 180 4. if the number of selected turbines is smaller than minimum quantity N , complete the set of turbines by adding additional
 turbines in order of their projected position on \vec{W}
5. from the set of N turbines, select the subset of M turbines with the highest wind speeds

An appropriate choice for parameter D , is a value slightly smaller than the distance between the outer turbine rows of the
 farm. As a consequence, in case that the wind vector \vec{W} is nearly perpendicular to a side of the wind farm, all turbines in the




Figure 1. Determination of the upstream turbines of the wind farm with turbines A1 to E7, for the case of an average wind \vec{W} from north-northwest direction. Turbines A1 to A5 are located in the most upstream zone of the farm with length D along \vec{W} . As in this example the minimal quantity of upstream turbines to be taken into account (N) equals 6, turbine B1 is added to the set of upstream turbines. Finally, as in this example the quantity of upstream turbines with the highest wind speed to be taken into account (M) equals 4, the two turbines with the lowest wind speed are not taken into account for determining the inflow wind characteristics. In this example, these are turbines A3 (which appears to be in a fully-waked position from a turbine of a neighbouring wind farm) and B1 (which is in a partially-waked position from turbine A1).

185 upstream row will be selected as upstream turbines, enabling to capture possible differences in wind characteristics over the complete width of the farm. In case that the wind vector \vec{W} is oriented towards a corner of the farm layout, however, only one single or very few turbines may be selected. Therefore, a minimum quantity of N of turbines is selected, with N a value larger than 1 and smaller than the quantity of turbines in a single row. Finally, in order to remove turbines that risk to be located in a narrow waked position caused by a turbine from the farm itself or from a neighboring farm, only the M turbines
 190 with the highest wind speeds are retained. The farm inflow wind parameters (inflow wind speed, inflow wind direction, inflow turbulence intensity and inflow wind direction variance) are then calculated as averages of the corresponding features from the individual upstream wind turbines. The overall methodology presented in this paper can also be applied with alternative algorithms for determining the upstream turbines. There is, however, always a trade-off between the responsiveness of the model on wind transients and the average accuracy of the model, due to the spatial variation of the wind characteristics across
 195 the wind farm.

An additional farm parameter, having a major influence on the farm power production, is the set-point of the farm power controller. Currently, for the majority of wind farms, active farm power control is still rarely used. Some farms, however, do nowadays already use actively farm power control in order to perform power balancing (Kölle et al., 2022). Due to the fast



growing share of wind energy in the energy mix, the use of farm power control may become more predominant in the near future. For the wind farms used as examples in this paper, the farm power set point data has a 1-minute time resolution.

Finally, one additional wind farm feature is composed from the turbine SCADA data: the quantity of stopped turbines. If one or more turbines are stopped (e.g. for maintenance reasons), the total power of the farm is reduced. On the other hand, a turbine in standstill will not cause wake for other downstream turbines. Depending on the available turbine SCADA data of the farms, the feature "quantity of stopped turbines" is built by counting the number of turbines that is producing zero power (or even consuming power) or as the number of turbines with power capability equal to zero. Notice that, based on this feature, the farm power model does not get the information which turbine(s) specifically are at standstill, which may lead to some loss of accuracy of the model. Indeed, for example, stopping a turbine in a waked position may lead to less power reduction than stopping a turbine in an upstream position with free wind inflow.

2.1.3 SCADA weather data

For the showcases in this paper, measurement data from a weather station located in one of the wind farms is used. The data set comprises the air temperature, humidity and pressure. Based on these three measurements, the relative air density is calculated. For the farms used as examples in this paper, the available measurements are data averaged over 10-minute intervals. In order to use the air density as input for the 1-minute farm and turbine power models, the data series is interpolated to a 1-minute data sequence.

2.1.4 Weather forecast data

For the showcases in this paper, wind speed and direction forecasts of a commercial weather forecast provider have been used, with a 1-hour time resolution. Based on the lead time of the forecasts, separate data sets have been composed with intra-day forecasts, day-ahead forecasts (before 11 a.m. the day before energy production) and three-days-ahead forecasts.

2.1.5 Train, test and validation data

After integrating the different data sources and generating the appropriate data features, the data is split in distinct data sets for model training, testing and validation. First, a long time-sequence is selected as validation data set. Thereafter, the test data set is established by selecting all data from some days-of-the-month from the remaining data. This, in order to guarantee that the test and training data sets are sufficiently independent, and, at the same time, are both representative for all seasons, hours-of-the-day and days-of-the-week. How the data split has been done for the wind farms used as examples in this paper, is explained more in detail in section §3.1.1.

2.2 Machine learning models

The modular deep learning approach presented in this paper combines multiple machine learning models. The different models are described in §2.2.1 to §2.2.5. How these models are finally integrated, is described in §2.2.6.



2.2.1 Wind farm power ML model

230 The core ML model proposed in this paper predicts the wind farm power based solely on measurement data from conventional turbine instrumentation. This type of data is usually available to any wind farm operator. Each of the input features of the model has a direct (physical) influence on the power production of the wind farm.

As it takes typically several minutes up to half an hour for the wind to cross a complete wind farm (depending on the wind speed and size of the farm), the wind and wake characteristics at time t_i across the farm do not only depend on the wind inflow at time t_i , but also on the evolution of the wind inflow between t_i and $t_i - T$ (with $T = 30$ minutes). Also control actions from the farm power controller and the quantity of stopped turbines in the past may have an influence on the spatial wind and wake profile across the farm. In contrast, some input parameters have an immediate discontinuous impact on the farm power production, independent of their historical values. For example, a farm power controller can reduce instantaneously the power of all turbines across the wind farm. Therefore, the proposed model is composed of two separate parts. For each input parameter that influences the wind and wake profile across the farm in a continuous way, a sequence of historic values is passed to a convolution branch. The outputs of these convolution branches are then passed to a feed-forward neural network, together with the input parameters that have (only) an immediate impact on the wind farm power production.

The structure of the proposed farm power model is shown in Figure 2. K separate convolution branches consist of two 1D convolutional layers. The outputs of these K branches are passed, together with L additional input features, through a dropout layer, to a feed-forward regression neural network composed of three dense layers, each followed by a dropout layer. Each convolution layer has 8 kernels with a kernel size of 3, a stride of 1 and no padding. The fully connected dense layers have respectively 128, 256 and 256 units with a rectified linear activation function.

Time sequences from t_i to $t_i - T$ of the following K input parameters (based on SCADA data) are passed to the K convolution branches:

- 250 – farm inflow wind speed,
- lateral component of the farm inflow wind direction,
- longitudinal component of the farm inflow wind direction,
- farm inflow turbulence intensity (Equation 1),
- farm inflow wind direction variance (Equation 2),
- 255 – air density,
- set point of the farm power controller (if available in the data set),
- quantity of stopped turbines, and
- active power capability of the farm (if available in the data set).

In parallel, the values at t_i of the following L input parameters are passed directly to the feed-forward component of the neural network:




Figure 2. Structure of the proposed deep learning farm power model. K separate convolution branches consist of two 1D convolutional layers. The outputs of these K branches are passed, together with L additional input features, through a dropout layer, to a three-layer regression model that is composed of three dense layers, each followed by a dropout layer.

- set point of the farm controller (if available in the data set),
- quantity of stopped turbines, and
- active power capability of the farm (if available in the data set).

Notice that the lateral and longitudinal components of the wind direction are used as inputs, instead of the wind direction itself (expressed in degrees). This is done to guarantee continuity in the data for wind direction $360^\circ/0^\circ$.

The dropout layers in the neural network model allow to predict not only the expected power production of the wind farm, but also to quantify the uncertainty of that prediction. The method applied in this paper is referred to as *Monte Carlo dropout* (Gal and Ghahramani, 2016). This is an epistemic method as it quantifies the uncertainty arising from the model architecture and the amount of data.

Instead of generating one single prediction of the farm power for time step t_i , N' different power predictions \hat{P}_i^n are generated by using the model with active dropout layers (in the same way as during the training phase of the model). The power prediction \hat{P}_i is then calculated as the average of these N' different power predictions. In addition, also the variance of these N' power forecasts $\sigma_{\hat{P}_i}^2$ can be determined, as well as an arbitrary set of percentiles and thus confidence intervals $[\hat{P}_i - \alpha_i, \hat{P}_i + \beta_i]$.

Unfortunately, the Monte Carlo dropout method is prone to miscalibration, i.e. the predictive uncertainty does not correspond well to the model error. Therefore, a method referred to as *sigma-scaling* is applied, calibrating jointly the epistemic uncertainty from the model and the aleatoric uncertainty from the data (e.g. due to sensor noise) (Laves et al., 2020). For each time step t_i



of the test data set, the following ratio is calculated:

$$q_i^2 = \frac{(\hat{P}_i - P_i)^2}{\sigma_{\hat{P}_i}^2} \quad (3)$$

280 with P_i the true farm power at time step t_i . According to Equation 3, q_i^2 is thus the ratio of the prediction squared error of the model for time step t_i and the variance calculated with the Monte Carlo dropout method for that time step. Analysis of the results for the wind farms used as example in this paper, shows that the prediction errors $\hat{P}_i - P_i$ and ratios q_i^2 depend mainly on the wind speed and the set point of the farm power controller. This could be expected, as for low wind speeds and for high wind speeds above the rated turbine wind speed, the power curve of the turbines is relatively flat. In contrast, for wind speeds slightly below the rated wind speed, the power curve of the turbines is the steepest. Therefore, a calibration function $q^2(ws, sp)$ is established, with ws the wind speed and sp the set point of the farm power controller. This is done by mapping a simple feed-forward neural network (with two hidden dense layers with 256 units) to the complete test data set. With this calibration function, the variance predicted with the Monto Carlo dropout method $\sigma_{\hat{P}_i}^2$ is re-calibrated as:

$$\hat{\sigma}_{\hat{P}_i}^2 = \sigma_{\hat{P}_i}^2 \times q^2(ws_i, sp_i) \quad (4)$$

290 Similarly, the confidence intervals generated with the Monte Carlo dropout method are re-scaled as:

$$[\hat{P}_i - \hat{\alpha}_i, \hat{P}_i + \hat{\beta}_i] = [\hat{P}_i - \alpha_i \times q(ws_i, sp_i), \hat{P}_i + \beta_i \times q(ws_i, sp_i)] \quad (5)$$

2.2.2 Farm internal wake loss

The power loss in a wind farm with J wind turbines due to internal wake, can be calculated as:

$$P_{wake} = \sum_{j=1}^J P_{WT}^j - P \quad (6)$$

295 with P_{WT}^j the power production of wind turbine j subjected to wind with the same characteristics as the inflow wind of the farm, and P the farm power. In case of J identical turbines, the farm power loss due to wake can be simplified as:

$$P_{wake} = J \times P_{WT} - P, \quad (7)$$

$$\text{where } P_{WT} = P_{WT}^j, \forall j. \quad (8)$$

300 Consequently, the wake loss can be modelled using the wind farm power model (as proposed in §2.2.1) and an equivalent model for a single turbine:

$$P_{wake}(ws, wd, ti, wd_{var}, \rho) = J \times P_{WT}(ws, ti, wd_{var}, \rho) - P(ws, wd, ti, wd_{var}, \rho) \quad (9)$$

with ws , wd , ti , wd_{var} and ρ respectively the farm inflow wind speed, wind direction, turbulence intensity, wind direction variance and air density. Notice that for the farm power model these wind parameters are time sequences. Consequently, the influence of temporal variations of these parameters are captured in the wake model as well.



305 2.2.3 Turbine ML model

Based on the same data set as used for the farm power model, also a turbine power model is built. This turbine model has thus also a 1-minute time resolution. However, in order to model a healthy turbine without any technical deration, reduced power mode or curtailment by its power controller, all data points with a reduced power capability and/or curtailment are removed from the training data set.

310 In contrast to a wind farm, the power production of a single turbine does not depend on the wind speed from multiple antecedent 1-minute time steps. Indeed, the response time of a single wind turbine, which is determined predominately by the inertia of its rotor, is significantly faster. Furthermore, the wind direction has no direct influence on the power of a turbine, as long as the yaw control system orientates the turbine perpendicular to the incoming wind direction. (For the wind farms used as examples in this paper no wake steering is done by applying yaw control).

315 As input data features, wind speed, turbulence intensity (Equation 1), wind direction variance (Equation 2) and air density are used. The turbine power is modeled by a simple feed-forward neural network, composed of three fully connected dense layers, each with 128 units with a rectified linear activation function.

2.2.4 Auxiliary ML models for missing input parameters

The farm power model (as specified in §2.2.1) has many input parameters related to the inflow wind. Sometimes the value of
320 some of these wind characteristics is not known, or at least not accurately. For example, weather forecast providers usually do not provide accurate information about the wind turbulence taking also the presence of neighbouring wind farms into account. In such case, one could decide to use a single average value for these parameters. However, the turbulence intensity, wind direction variance, air density and the number of stopped turbines depend all more or less on the wind speed and wind direction. Therefore, for each of these four parameters a model is built to predict its value based on the wind speed and wind
325 direction. These auxiliary models may be specific for a particular wind farm and are not the focus of this paper.

2.2.5 ML models for mapping weather forecasts

Commercial weather forecast services provide nowadays weather forecasts for specific geographical locations, such as the position of a wind farm. Forecasts of different providers can differ due to the use of different weather models or data. Usually, wind speed and wind direction forecasts do not take into account the presence of neighbouring wind farms or coast lines.
330 Furthermore, the forecasts may not have exactly the same calibration as the instrumentation that has been used to train the wind farm power model.

Therefore, for each weather forecast service a ML-model is trained to map the forecast wind speed and wind direction to the measured inflow wind speed and direction experienced by the wind farm.



2.2.6 Integration of all ML models

Figure 3 shows the complete modular structure, integrating all ML models presented in sections §2.2.1 to §2.2.5. Weather forecasts of weather forecast providers are converted to corrected forecasts of the wind inflow experienced by the farm. In doing so, phenomena such as external wake, wind farm blockage, coastal effects and other unknown systematic forecasting errors are accounted for. Based on these corrected inflow wind speeds and directions, possibly unknown wind parameters can be estimated with the respective auxiliary models. Finally, with all resulting input parameters, the wind farm power is forecasted, as well as the wind farm internal wake.

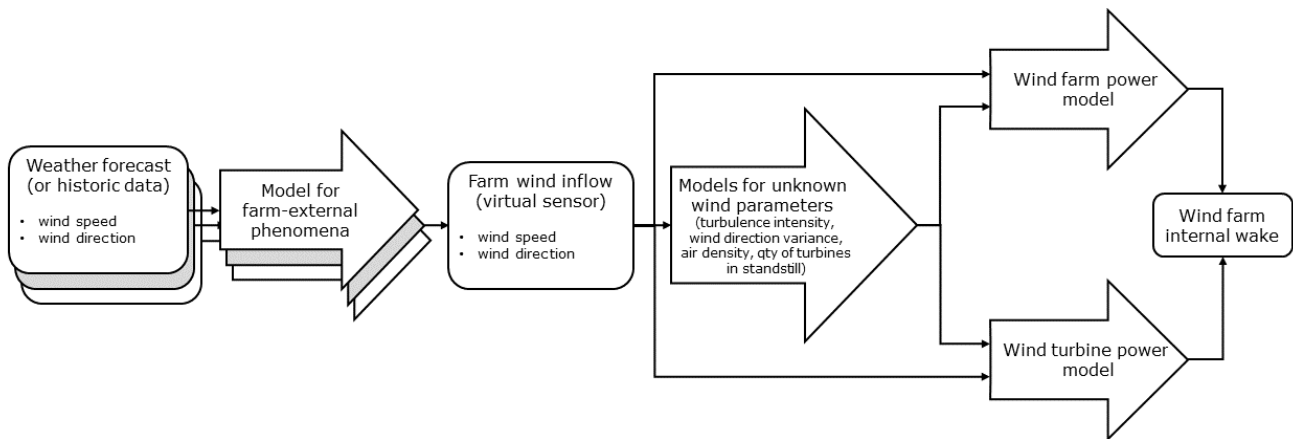


Figure 3. Overview of the modular structure of the multi-component pipeline with multiple ML models

2.3 Computing hardware & software

All computing performed for this paper, including the training of the ML-models, has been performed with a standard notebook (processor: 11th Gen Intel(R) Core(TM) i7-1165G7 @ 2.80GHz, internal RAM: 16.0 GB, 64-bit operating system). All code is written in Python, with Tensorflow as machine learning platform. This shows that the proposed model can be run on hardware that is readily accessible to wind farm operators.

3 Results

The methodology as described in section 2, has been applied to two offshore wind farms located in the Belgian-Dutch wind farm cluster in the North Sea (see Figure 4). This cluster comprises 13 wind farms. It is located on a distance of 20 km to 60 km from the Belgian and Dutch coastline. For confidentiality reasons, the two wind farms are denoted in this paper as *Wind Farm 1* and *Wind Farm 2*. The two farms are operated by different farm operators and have a different type of wind turbines. The main characteristics of the two farms and the complete wind farm cluster are listed in Table 1. All presented



results involving farm power and turbine power are normalised on respectively the installed capacity of the wind farm and the rated turbine power.

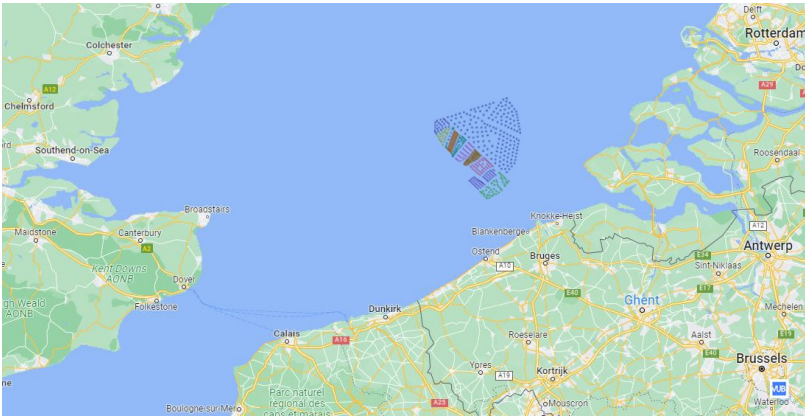


Figure 4. Position of the Belgian-Dutch offshore wind farm cluster in the North Sea (© Google Maps)

Table 1. Main characteristics of the Belgian-Dutch offshore wind farm cluster and the two wind farms used as examples in this paper.

| | Quantity of turbines | Installed capacity [MW] | Acreage [km ²] |
|---------------------------------|----------------------|-------------------------|----------------------------|
| Belgian-Dutch wind farm cluster | 572 | 3764 | 608 |
| Wind Farm 1 | ≥40 | ≥300 | ≥40 |
| Wind Farm 2 | ≥30 | ≥200 | ≥20 |

355 The results are grouped in the following three sections. In Section 3.1, results are presented related to the core wind farm
power prediction model of each of the two wind farms. Section 3.2 focuses on the estimation of farm-internal wake. Finally, in
Section 3.3, some results are shown related to the integration of weather forecasts.

3.1 Wind farm power model

3.1.1 Data

360 For both Wind Farm 1 and Wind Farm 2, weather and turbine data was available for a period of about 2.5 years. Table 3 shows
the quantity of 1-minute data points resulting from the preprocessing of the SCADA data, and split in a distinct training, test
and validation data set (as described in section 2.1). The test data set comprises the data from days 2, 3, 4, 16, 17, 18 and 28 of
each month. This, in order to guarantee that the test and training data sets are sufficiently independent, and, at the same time,
are both representative for all seasons, hours-of-the-day and days-of-the-week. The validation data is a time sequence of 11
365 consecutive days with strong wind fluctuations.



Table 2. Quantity of data points used for training, test and validation of the two wind farm power models

| Data set | Wind Farm 1 | Wind Farm 2 |
|------------|-------------|-------------|
| Training | 960695 | 978008 |
| Test | 292996 | 292078 |
| Validation | 15810 | 15810 |

The data plots in figures 5, 7 and 9 illustrate the dependency of the farm power on respectively the wind turbulence intensity, wind direction variance and air density. Higher wind turbulence results in higher farm power, as higher turbulence facilitates the wake recovery reducing so the possible power loss for downstream turbines. A higher air density results also in an increase of the turbine power, because the mass, and thus kinetic energy, of the moving air is higher.




Figure 5. Farm power as a function of inflow wind speed and wind turbulence intensity (Wind Farm 2)




Figure 6. Inflow turbulence intensity as a function of wind speed and wind direction (Wind Farm 2)

The polar data plots in figures 6, 8 and 10 give an indication of the correlation between respectively turbulence intensity, wind direction variance and air density with the inflow wind speed and wind direction. In the northwest and southeast wind directions, Wind Farm 2 has in its immediate vicinity neighboring wind farms with densely positioned turbines, resulting in a high turbulence intensity for these wind directions. Also in the northeast direction, Wind Farm 2 has a neighbouring wind farm. However, that wind farm is located further away and its turbines are positioned less densely. It can be seen also, that the wind turbulence is lower for high wind speeds than for slow speeds. Furthermore, it can be seen that the southwest direction is the wind direction with the highest wind speeds. In that direction, the wind is coming from over sea in parallel with the coastline,




Figure 7. Inflow wind direction variance as a function of wind speed and wind direction (Wind Farm 2)




Figure 8. Inflow wind direction variance as a function of wind speed and wind direction (Wind Farm 2)




Figure 9. Farm power as a function of inflow wind speed and air density (Wind Farm 1)




Figure 10. Air density as a function of wind speed and wind direction (Wind Farm 1)

through the narrow Strait of Dover, which has high cliffs. It can be noticed also that wind in parallel to the coast line typically has a lower air density compared to the orthogonal directions from and to the mainland. This corresponds to the fact that humid
 380 air has a lower air density than dry air.



When comparing figures 5 and 9 from respectively Wind Farm 2 and Wind Farm 1, it can be seen that Wind Farm 2 has markedly more data points with a reduced power. This is due to the power controller of the wind farm which, in some circumstances, regulates the farm power in a continuous manner. In contrast, for Wind Farm 1, there appear to be farm curtailments to a few specific discrete power set points (such as $\sim 60\%$ of the maximum farm power). Furthermore, it can be noticed also, that for Wind Farm 1, the total installed turbine capacity is never reached. This is due to the farm power controller, which limits the maximum producible power of the farm.

Table 3 shows the parameters used to determine the set of upstream turbines according to the algorithm described in §2.1.2. Wind farm 1 comprises more turbines than Wind Farm 2 (ref. parameter N) and these are positioned more widely apart (ref. parameter D).

Table 3. Parameters used for the selection of the upstream turbines

| Parameter | Wind Farm 1 | Wind Farm 2 |
|-----------|-------------|-------------|
| D | 1000m | 500m |
| N | 12 | 6 |
| M | 10 | 4 |

3.1.2 Performance metrics

In Table 4, some performance metrics (Piotrowski et al., 2022) of the farm power models for Wind Farm 1 and Wind Farm 2 are listed: the root mean square error normalised on the installed power capacity of the wind farm (nRMSE), the normalised mean absolute error (nMAE), normalised mean error (nME) and the R2-score, for both the training and test data set. The nME for each of the farms is near to 0% and the nMAE is 2.42% and 2.14% for the test data sets of Wind Farm 1 and Wind Farm 2 respectively. The differences between the performance metrics of the training and test data sets are very small, which is a sign for a satisfactory fit (little overfitting).

Table 4. Performance metrics of the power models for Wind Farm 1 and Wind Farm 2

| | Wind Farm 1 | | Wind Farm 2 | |
|----------|-------------|---------|-------------|---------|
| | training | test | training | test |
| nMAE | 2.22 % | 2.42 % | 2.15 % | 2.14 % |
| nME | -0.05 % | -0.17 % | -0.11 % | -0.10 % |
| nRMSE | 4.02 % | 4.21 % | 3.69 % | 3.66 % |
| R2-score | 0.99 | 0.99 | 0.99 | 0.99 |



In Table 5, as a comparison basis, the corresponding performance metrics are shown for two other ML models with a feed-
 400 forward neural network (FNN) architecture. Both FNNs have three dense layers, like the farm power model proposed in this
 work (§2.2.1), but have no convolution layers. The first FNN-model takes only the wind speed and wind direction at time-step
 t_0 as inputs, whereas the second FNN-model takes all input features from the proposed wind farm power model at time-step
 t_0 as inputs. Consequently, both FNN-models do not get input data from the past (preceding 30 minutes). In order to limit
 overfitting of these two FNN-models, the number of units per neural network layer has been reduced to 32. Comparing the
 405 performance metrics of the wind farm power model proposed in this work with those of the two FNN-models as baselines,
 one can see e.g. that for the test data the nMAE is reduced with about 47% and 18% respectively (for both wind farms). Thus,
 both the usage of additional input parameters (i.a. turbulence intensity, air density and setpoint of farm power controller) and
 information from the preceding time period (30 minutes) improve the performance of the farm power model significantly.

Table 5. Performance metrics for two baseline deep learning power models for both Wind Farm 1 and Wind Farm 2

| | FNN(ws_0, wd_0) | | | | FNN(all input features at t_0) | | | |
|----------|---------------------|-------|-------------|-------|-----------------------------------|--------|-------------|-------|
| | Wind Farm 1 | | Wind Farm 2 | | Wind Farm 1 | | Wind Farm 2 | |
| | training | test | training | test | training | test | training | test |
| nMAE | 4.16% | 4.58% | 4.48% | 4.06% | 2.38% | 2.94% | 2.54% | 2.59% |
| nME | 0.95% | 1.04% | 1.48% | 1.14% | 0.06% | -0.07% | -0.04% | 0.00% |
| nRMSE | 7.84% | 8.58% | 9.57% | 8.45% | 4.36% | 5.19% | 4.43% | 4.54% |
| R2-score | 0.95 | 0.94 | 0.92 | 0.94 | 0.99 | 0.98 | 0.98 | 0.98 |

410 Figures 11 and 12 give some insight in the prediction error for individual test data points. It can be seen that the prediction
 error is the smallest for conditions where the farm power is close to 0% or 100% of the installed capacity. Indeed, these regions
 comprise respectively many data points with wind speeds below the turbine cut-in wind speed and wind speeds above the rated
 turbine wind speed, where the power curve of the turbines is relatively flat.




Figure 11. Relation between predicted and true farm power for all data points of the test data set for Wind Farm 1




Figure 12. Relation between predicted and true farm power for all data points of the test data set for Wind Farm 2

3.1.3 Validation time sequence

Figure 13 shows the predicted and true farm power for a 60-hour time sequence of validation data for Wind Farm 1, as well as the predicted confidence intervals. The true farm power lays most of the time within the 68.3% confidence interval and rarely outside of the 95% confidence interval. The confidence intervals are indeed the smallest for high wind speeds resulting in maximum farm power and for low wind speeds with a farm power close to 0 MW. The uncertainty appears to be the highest when the farm power is fluctuating and peaking heavily. In contrast, for long continuous power increases or decreases, the uncertainty of the model appears to be relatively low.




Figure 13. Predicted farm power and uncertainty intervals for a 60-hour time sequence of validation data (Wind Farm 1)



Figure 14 shows the predicted and true farm power for a time sequence of test data during which the farm power controller of Wind Farm 2 is actively curtailing the farm power (i.e. the set point of the power controller is smaller than 100%). The uncertainty of the model appears to be very low when the wind speed is largely sufficient to attain the power set point.




Figure 14. Predicted farm power and uncertainty intervals for a 16-hour time sequence of test data with active power control (Wind Farm 2)

3.1.4 Sensitivity analysis

Figures 15 to 20 illustrate the sensitivity of the farm power curve generated by the farm power model for each of the input parameters. In each of the plots, the farm power is shown as a function of the wind speed. The other input parameters are predicted by the corresponding auxiliary models (see §2.2.4) based on the wind speed and a specific wind direction. One single input parameter is then adapted slightly in order to analyse the resulting impact on the predicted farm power. Each simulation is a steady-state simulation, i.e. the value of each of the input parameter sequences is constant in time.

Figures 15 to 17 show that the models predict a higher farm power for respectively an increased turbulence intensity, wind direction variance and air density. This confirms what can be seen in figures 5, 7 and 9. Figure 18 shows the predicted power for Wind Farm 1 for two different wind directions. The produced power for wind direction 240° is lower than for wind direction 180° . This difference is due to farm-internal wake (see also section 3.2). Indeed, wind direction 240° is parallel with a long cross-section of Wind Farm 1. In contrast, wind direction 180° is slightly off the short cross-section of the wind farm. For this wind direction, the internal wake is thus minimal. Figure 19 shows the predicted power reduction for Wind Farm 2 when respectively one and two turbines are stopped. As mentioned in §3.1.1, the maximum power of this wind farm is limited by its farm controller. As can be seen in Figure 19, at high wind speeds this maximum power limit can almost be reached with one turbine at standstill. Figure 20 shows the farm power for Wind Farm 2 in case the set point of the farm power controller equals 30%, for the wind directions 325° and 260° . For low wind speeds, the farm power for wind direction 260° is lower than for 325° , because of the higher farm-internal wake.




Figure 15. Predicted farm power as a function of inflow wind speed and wind turbulence intensity (Wind Farm 2)

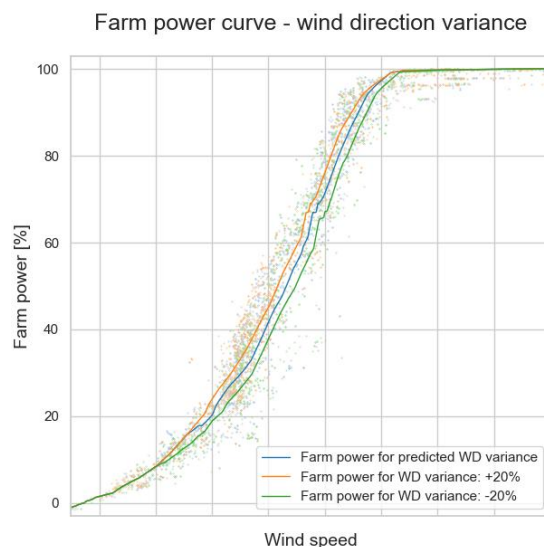


Figure 16. Predicted farm power as a function of inflow wind speed and wind direction variance (Wind Farm 2)




Figure 17. Predicted farm power as a function of inflow wind speed and air density (Wind Farm 1)




Figure 18. Predicted farm power in function of inflow wind speed and wind direction (Wind Farm 1)




Figure 19. Predicted farm power in function inflow of wind speed and the quantity of turbines in standstill (Wind Farm 1)




Figure 20. Predicted farm power in function of inflow wind speed and direction, with the farm power controller set point equal to 30% (Wind Farm 2)

3.1.5 Farm power dynamics

All power predictions presented in §3.1.4 are for steady-state conditions, i.e. all input parameters of the model are constant during the thirty-one 1-minute time steps. As explained in §2.2.1, the structure of the neural network of the farm model has been chosen specifically to be able to capture temporal variations of the inflow wind characteristics. In this section, the farm power is predicted for wind speeds that fluctuate over time. The objective is twofold: firstly, to assess the ability of the model to predict consistent and physically meaningful results under these dynamic conditions, and secondly, to gain insights into how wind speed variations impact farm power production, what cannot be modelled by traditional steady-state power models. Two types of wind speed variations are tested: linear wind speed ramps and sinusoidal varying wind speeds with different frequencies.

Figure 21 shows power predictions for Wind Farm 1 for the same wind directions as shown in figure 18. However, in addition to the two power curves for constant wind speeds, the power predictions are shown for the cases with a linear wind speed increase and decrease with a change rate of 0.05 m/s per minute. It can be seen that in case of increasing wind speed, the predicted farm power is lower. In contrast, for decreasing wind speed, the predicted power is higher. Indeed, if at the inflow side of the wind farm the wind speed is increasing, this means that downstream in the wind farm the wind speed is still lower than at the inflow side, which results in a lower total farm power. As can be seen on this plot as well, this effect is larger for wind direction 240° (blue arrows) than for wind direction 180° (orange arrows). This is because the cross-section of the wind farm in direction 240° is longer than in wind direction 180°. Consequently, changed wind speeds need more time to reach



the downstream turbines. In addition, on the plot it can be seen that there is a hysteresis for the start-up and shut-down of the turbines around the cut-in wind speed.




Figure 21. Predicted farm power in function of inflow wind speed with a linearly increasing and decreasing speed with a change rate of 0.05 m/s per minute, for wind directions 240° and 150° (Wind Farm 1).

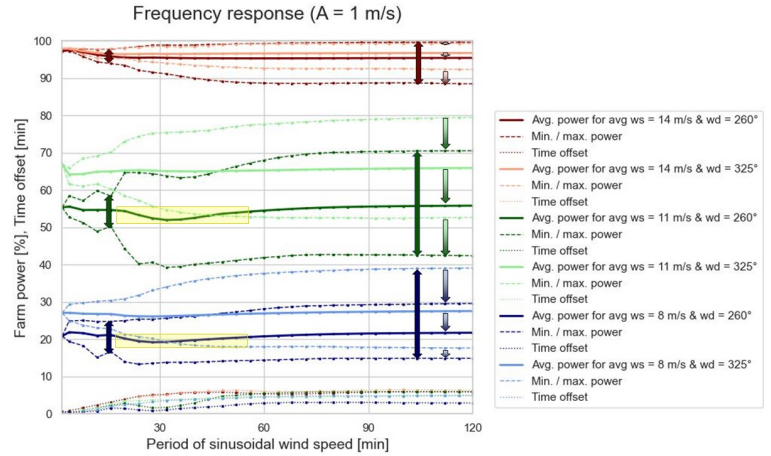


Figure 22. Frequency response of predicted farm power in case of a wind inflow with a sinusoidal oscillating wind speed component with an amplitude of 1 m/s. The plots show the average, maximum and minimum power of the predicted oscillating farm power in function of the period of the sinusoidal component of the wind speed. In addition, also the time offset between the oscillating farm power and wind speed is shown. The color of the curves indicate the average wind speed (resp. 14 m/s, 11 m/s and 8 m/s) and the wind direction (260° and 325°) (Wind Farm 2).

Figure 22 shows the frequency response of the power model for Wind Farm 2. The wind speed is simulated as a constant average speed superposed with a sinusoidal component with an amplitude of 1 m/s. Simulations were run for a period T of the sinusoidal component equal to 2 minutes up to 120 minutes ($T = 2, 4, 8, 12, 16, 20, \dots, 120$ minutes). This has been done for different average wind speeds (14 m/s, 11 m/s and 8 m/s) and wind directions (260° and 325°). The curves show the average, maximum and minimum values of the resulting oscillating farm power, as well as the time delay between the sinusoidal wind speed component and the oscillating farm power.

As a period of 120 minutes is much longer than the time needed for the wind to cross the complete wind farm, this is a quasi-static wind condition. If the period is smaller, the frequency of the wind oscillation is higher. However, for $T = 2$ minutes, taking into account the 1-minute time step, the wind speed is again constant (because $\sin(\frac{2\pi}{T} \times t_i) = \sin(\frac{2\pi}{2} \times n) = 0$).

As can be seen on the plot, for a specific wind speed and oscillation frequency, the average, maximum and minimum farm power for wind direction 260° (with high wake) is always lower than for wind direction 325° (with less wake) (see downward arrows). For faster fluctuating wind speeds (i.e. with shorter oscillation period), the amplitude of the farm power fluctuation decreases (see bidirectional arrows). In addition, in case of wind conditions with a high farm-internal wake, there appears to



be a decrease of the average farm power (see yellow markings). For small oscillation periods below 15 minutes, the average power is increasing again, converging back to the same value as for quasi-static wind conditions.

475 For short oscillation periods and consequently small wavelengths of the spatial wind speed distribution, the time delay of the farm power is converging to zero. Indeed, the time delay cannot be longer than the oscillation period of the sinusoidal wind speed component. For longer oscillation periods, the time offset converges to a constant value. Logically, the time delay will never be higher than the time needed by the wind to cross the complete farm.

480 For short oscillation periods below 15 minutes, the wavelength of the wind speed oscillation is getting smaller than the cross-section of the farm. This may cause the important decrease of the amplitude of the farm power oscillation. This might also explain the jigsaw shape of the maximum and minimum power in these conditions. To analyse these effects in more detail, a wind farm model could be used that does not only predict the total farm power, but also the power production of each individual turbine.

3.2 Farm-internal wake

485 As already shown in Figure 18, the farm power production depends on the wind direction due to the difference in wake loss. In order to calculate the farm-internal wake in absolute terms, an ML-model has been established for a single turbine, as described in §2.2.3. By subtracting the power predicted by the farm power model from the predicted turbine power under identical wind conditions (multiplied with the number of turbines in the farm), the farm-internal wake effect can be isolated from other influences on the farm power.

490 Figure 23 shows the measured farm power for Wind Farm 1, as a function of the wind direction and wind speed. Figure 24 shows the subset of these data points for which the set point of the farm power controller is equal to 100% and at maximum one turbine is at standstill. Furthermore, the scope of the plot has been limited to the wind speed range for which wake is most predominant. Figure 25 shows the corresponding predicted farm power by the farm model for steady-state wind inflow (during 30 minutes). It can be seen that for the directions west-southwest and east-northeast for a given wind speed, the farm power is
495 lower than for other wind directions. This can be seen more clearly after subtraction from the predicted turbine power. Figure 26 shows the power loss due to internal wake as a percentage of the installed capacity of the farm. The maximum internal wake loss is about 30%. The reason for the high wake in the west-southwest and east-northeast directions, is that these directions are parallel to the long axis of the wind farm, with multiple turbines positioned after each other. For wind speeds above 13 m/s, the power loss due to wake approaches 0 MW. This is because at such high wind speeds there is sufficient energy in the wind,
500 and consequently each turbine in the farm can produce sufficient power so that the maximum farm power is reached.

3.3 Power forecasting based on weather forecasts

The farm power prediction models presented in sections 3.1 and 3.2, predict the farm power based on measurement data of the wind inflow and some other parameters of the farm. In order to forecast the farm power in the future with these farm models, forecasts of the wind speed and direction (and if available also the other wind characteristics) are required.

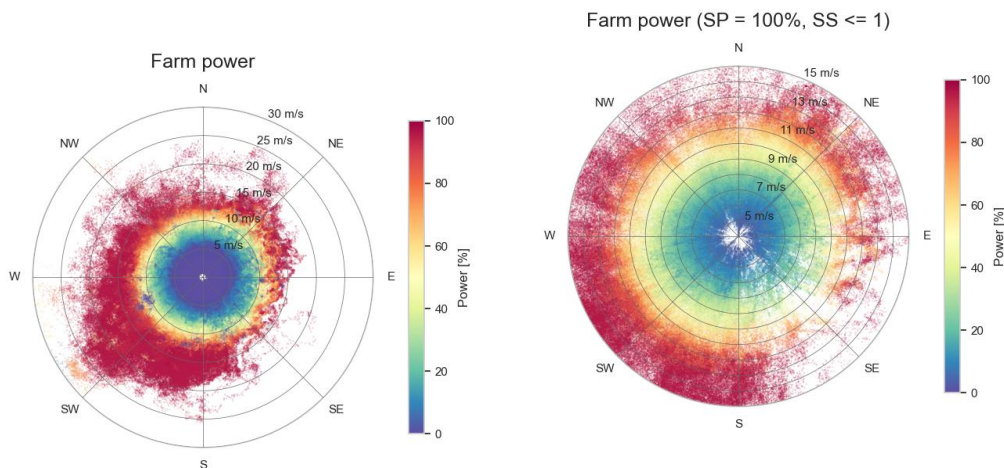


Figure 23. Farm power measurement data points in function of wind speed and direction (Wind Farm 1)

Figure 24. Subset of farm power measurement data points in function of wind speed and direction, with farm power control set point equal to 100% and quantity of turbines in standstill equal to 0 or 1 (Wind Farm 1)




Figure 25. Predicted farm power in function of wind speed and direction, for constant wind conditions, farm power control set point equal to 100% and with all turbines in operation (Wind Farm 1)

Figure 26. Predicted farm-internal wake in function of wind speed and direction, for constant wind conditions, farm power control set point equal to 100% and with all turbines in operation (Wind Farm 1)



505 Weather forecasts of different providers may differ among each other as they can be based on different weather models and data. Furthermore, they typically do not take (accurately) local effects into account, like for example neighbouring wind farms. These may cause a reduction of the wind speed, an increase of the wind turbulence intensity and a redirection of the wind due to wind blockage. Also coastal effects may have a significant impact on the wind speed and direction.

Figure 27 shows the correction factor to be applied to the wind speed forecasts of a specific weather forecast provider for
 510 Wind Farm 1, depending on the forecast wind speed and direction. This correction factor has been calculated by mapping historical wind speed forecasts (those with the shortest lead time) from that weather forecast provider to the corresponding measured inflow wind speed of that farm. As can be seen in the figure, for wind directions between north-northwest and north-northeast, the wind speeds predicted by the correction model are only about 80% of the forecast wind speeds. The reason is that in that upstream direction many wind farms are located in immediate proximity. For wind speeds below 5 m/s, the correction
 515 factor is higher than one. This is due to the fact that the wind measurements on the turbines are not well calibrated and are over-estimating these low wind speeds. For such wind speeds below the cut-in wind speed, turbines are shut off anyway.




Figure 27. Correction factor for the wind speed forecasts of a weather forecast service for Wind Farm 1




Figure 28. Correction for the wind direction forecasts of a weather forecast service for Wind Farm 1

In Figure 28, it can be seen that for forecast wind directions in the sector from northwest, over south to east, the wind directions are in reality about 10 to 20 degrees higher, meaning that air flows coming from these directions are deflected by about 10 to 20 degrees in clockwise direction (compared to the values forecasted by the weather forecast provider). In contrast,
 520 air flow coming from the sector east to northwest, are deflected in anti-clockwise direction. This corresponds to the fact that Wind Farm 1 is located in the south of the wind farm cluster and the cluster has a rectangular-like shape with the long axis in northwest-southeast direction. Due to blockage of the wind by the wind farm cluster, the air flow deviates through pressure build-up in front of the cluster slightly in the direction of the outside corners of the cluster, where it can flow next to the cluster. For lower wind speeds, thus with lower momentum, this deflection appears to be sharper than for higher wind speeds.



525 In addition, in Figure 27 it can be seen that on average the corrected inflow wind speeds are lower than the ones forecasted
 by the weather service provider, also for wind directions without upstream wind farms causing wake. This may be may be
 partly attributable to coastal effects (the coast is the nearest in east to south directions) and partly to the blockage effect of the
 wind farm cluster. For the wind directions in the sector south-southwest (directed towards the corner of the wind farm cluster)
 slightly increased wind speeds can be observed (especially for wind speeds between 11 m/s and 12 m/s). This might be an
 530 indication of the acceleration of the air flow at the corner of the cluster occurring jointly with the deflection due to the blockage
 effect. However, this may also be attributable to other reasons, such as under-estimation by the weather forecast provider of
 wind speeds in parallel with the coastline through the narrow Strait of Dover, as weather forecast models cannot model all local
 effects. For forecasts from another weather forecast provider and for historical ERA5 data, similar discrepancies in wind speed
 and wind directions are observed, however with different biases and/or variances.

535 Using the complete chain of models as shown in Figure 3, starting from the models correcting the wind speed and direction
 forecasts, the auxiliary models to predict the wind turbulence and air density, and finally the farm power model, the farm
 power can be forecasted for multiple time horizons. Figure 29 shows the three-days-ahead, day-ahead and intra-day power
 forecasts for an eight-days sequence for Wind Farm 1. Figure 30 shows the corresponding wind speed forecasts used as inputs
 for the power forecasts. As weather forecasts with a shorter lead-time are usually more accurate than with longer lead-times,
 540 the resulting power forecasts are getting more accurate for shorter lead times as well.

This can be seen also in table 6, which shows multiple error metrics for the different wind speed and power forecasts. It can
 be seen also that the mean error (ME) of the uncorrected wind speed forecasts is always larger than 0 m/s. This is probably
 mainly due to the fact that the weather forecasts ignore the wind speed reduction caused by the wake of the upstream wind
 farms. Remark that after applying the correction model to the wind speed forecasts (refer to figure 27), the ME is around 0 m/s
 545 and both the MAE and RMSE are reduced significantly. Remark that a wind speed error of 1 m/s represents about 12.5% of the
 range between the cut-in and rated wind speed of a turbine (~ 4 m/s to ~ 12 m/s). Consequently, a wind speed forecasting error
 of that magnitude can result into a large farm power forecasting error, larger than the inaccuracy inherent to the farm power
 model itself.

Table 6. Performance metrics for intra-day (ID), day-ahead (DA) and three-days-ahead (3DA) wind speed forecasts, corrected wind speed
 forecasts and power forecasts for test data of Wind Farm 1. In the last column, the performance metrics are shown for the corresponding
 power predictions by the wind farm power model based on the measured inflow data (refer to table 4, column "test" for Wind Farm 1).

| | Wind speed forecast | | | Corrected wind speed forecast | | | Power forecast | | | Power prediction |
|----------|---------------------|---------|---------|-------------------------------|----------|----------|----------------|--------|--------|-------------------|
| | ID | DA | 3DA | ID | DA | 3DA | ID | DA | 3DA | SCADA inflow data |
| MAE | 1.7 m/s | 1.8 m/s | 2.5 m/s | 1.2 m/s | 1.2 m/s | 2.0 m/s | 8.9 % | 9.2 % | 15.9 % | 2.4 % |
| ME | 0.6 m/s | 0.7 m/s | 0.7 m/s | -0.1 m/s | -0.1 m/s | -0.1 m/s | 0.6 % | 0.9 % | 0.6 % | -0.2 % |
| RMSE | 2.1 m/s | 2.2 m/s | 3.2 m/s | 1.6 m/s | 1.7 m/s | 2.7 m/s | 14.0 % | 14.4 % | 24.3 % | 4.2 % |
| R2 score | 0.74 | 0.72 | 0.41 | 0.86 | 0.85 | 0.58 | 0.84 | 0.84 | 0.52 | 0.99 |




Figure 29. Power forecasts for Wind Farm 1 based on wind forecasts with different forecast horizons for a time sequence of 8 days




Figure 30. Wind speed forecasts with different forecast horizons (from a third-party weather forecasting provider)

550 The relatively large errors in the weather forecasts should be no surprise, as the weather forecasts used in this example have a time resolution of only a 1-hour (to be compared to the 1-minute resolution of SCADA data and wind farm power model). Moreover, day-ahead forecasts have a forecasting lead time of up to 37 hours (from day-ahead 11 a.m. until 12 p.m.). In addition, the processing time for making these weather forecasts may take up to 6 hours, so that the day-ahead forecasts may be based on data from 43 hours ahead. The intra-day forecasts used in this example, can have a lead time of up to half a day,

555 because the weather forecast data available for this work is updated only twice per day. Using more accurate weather forecasts updated more frequently and with a shorter time resolution will result in better farm power forecasts. One could argue that for the specific application of long-term wind farm power forecasting, using an accurate wind farm power forecasting model with such short time resolution (1 minute) adds little value. It should be reminded, however, that if for example a 1-hour resolution time would be used for the training of the wind farm power model, the behaviour of the wind farm power controller and the

560 temporal dynamics of inflow wind transients would not be captured by the farm power model, which are, however, crucial for other targeted applications of the model.



4 Conclusion

In the present work, a novel methodology is proposed to forecast the power production of a wind farm. The methodology is based on a multi-component pipeline with as two main components a deep learning wind farm power model and a distinct machine learning model for integrating weather forecasts.

The proposed wind farm power model relies solely on SCADA data from the wind farm itself, which are usually available to any wind farm operator. It captures the influence of several weather parameters, including wind speed, wind direction, turbulence intensity, wind direction variance and air density. Additionally, it captures the temporal dynamics of the wind inflow and the behaviour of the farm power controller. Also the number of turbines that are in standstill is taken into account. Notable, the model does not only predict the farm power with a high accuracy, but generates also confidence intervals for these power predictions. Furthermore, the farm power model is capable of predicting the farm power in only a few milliseconds on PC, making it significantly faster than even low-fidelity physics-based models.

A separate deep learning model is employed to post-process the wind speed and direction forecasts from weather forecast providers. In doing so, it takes farm-external factors into account, such as wake generation by neighbouring wind farms, wind farm blockage, coastal effects and systematic forecasting errors from the respective weather forecast providers.

Furthermore, in conjunction with a data-driven turbine power model, also the farm wake losses can be predicted.

The proposed methodology has been applied to two large real-world offshore wind farms. Performance metrics affirm a significantly improved prediction accuracy compared to some baseline ML models. Validation sequences demonstrate in addition the reliability of the predicted confidence intervals. Sensitivity analyses, performed on each of the model's input features, yield interpretable and physically meaningful results. In addition, the prediction capacity of the farm power models is demonstrated for fluctuating inflow wind speeds.

It is shown that the accuracy of wind farm power forecasts with a long forecasting horizon are limited predominantly by the accuracy of the weather forecasts and not by the uncertainty inherent to the farm power model itself.

In further research, we will integrate the wind farm power prediction models as digital twins into applications where their high prediction speed and the simulation of the farm power controller are crucial. Furthermore, the wind farm power model will be extended for short-term power forecasts.

Author contributions. SA: conceptualization, data curation, formal analysis, investigation, methodology, software, validation, visualization and writing (original draft preparation, review and editing); TV and PJD: supervision and paper reviewing; AN: Funding acquisition; JH: Funding acquisition, supervision and paper reviewing.

Competing interests. The authors declare that they have no competing interests.

<https://doi.org/10.5194/wes-2024-94>
Preprint. Discussion started: 5 September 2024
© Author(s) 2024. CC BY 4.0 License.



Acknowledgements. This research was supported by the Flemish Government (AI Research Program) and financially supported by the Energy Transition Fund projects Poseidon and Beforecast.



References

- Barthelmie, R. J., Hansen, K., Frandsen, S. T., Rathmann, O., Schepers, J., Schlez, W., Phillips, J., Rados, K., Zervos, A., Politis, E., et al.:
 595 Modelling and measuring flow and wind turbine wakes in large wind farms offshore, *Wind Energy: An International Journal for Progress
 and Applications in Wind Power Conversion Technology*, 12, 431–444, 2009.
- Becker, M., Allaerts, D., and van Wingerden, J. W.: FLORIDyn - A dynamic and flexible framework for real-time wind farm control, *Journal
 of Physics: Conference Series*, 2265, 032 103, <https://doi.org/10.1088/1742-6596/2265/3/032103>, 2022.
- Bleeg, J., Purcell, M., Ruisi, R., and Traiger, E.: Wind farm blockage and the consequences of neglecting its impact on energy production,
 600 *Energies*, 11, 1609, 2018.
- Boersma, S., Doekemeijer, B. M., Gebraad, P. M., Fleming, P. A., Annoni, J., Scholbrock, A. K., Frederik, J. A., and van Wingerden, J.-W.:
 A tutorial on control-oriented modeling and control of wind farms, in: 2017 American control conference (ACC), pp. 1–18, IEEE, 2017.
- Bossanyi, E. and Ruisi, R.: Axial induction controller field test at Sedini wind farm, *Wind Energy Science*, 6, 389–408,
<https://doi.org/10.5194/wes-6-389-2021>, 2021.
- 605 Braun, M., Gruhl, C., Hans, C. A., Härtel, P., Scholz, C., Sick, B., Siefert, M., Steinke, F., Stursberg, O., and Berg, S. W.-v.: Predictions and
 Decision Making for Resilient Intelligent Sustainable Energy Systems, arXiv preprint arXiv:2407.03021, 2024.
- Fleming, P., King, J., Simley, E., Roadman, J., Scholbrock, A., Murphy, P., Lundquist, J. K., Moriarty, P., Fleming, K., van Dam, J., Bay, C.,
 Mudafort, R., Jager, D., Skopek, J., Scott, M., Ryan, B., Guernsey, C., and Brake, D.: Continued results from a field campaign of wake
 steering applied at a commercial wind farm – Part 2, *Wind Energy Science*, 5, 945–958, <https://doi.org/10.5194/wes-5-945-2020>, 2020.
- 610 Fleming, P. A., Ning, A., Gebraad, P. M., and Dykes, K.: Wind plant system engineering through optimization of layout and yaw control,
Wind Energy, 19, 329–344, 2016.
- Foloppe, B., Munters, W., Buckingham, S., Vandeveld, L., and van Beeck, J.: Coupling of a dynamic wake model with WRF: a case study
 of the Belgian wind farms, in: 18th EAWC PhD Seminar, pp. 30–31, 2022.
- Gal, Y. and Ghahramani, Z.: Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning, in: *Proceedings
 615 of The 33rd International Conference on Machine Learning*, edited by Balcan, M. F. and Weinberger, K. Q., vol. 48 of *Proceedings of
 Machine Learning Research*, pp. 1050–1059, PMLR, New York, New York, USA, <https://proceedings.mlr.press/v48/gal16.html>, 2016.
- Gebraad, P. M. and Van Wingerden, J.: A control-oriented dynamic model for wakes in wind plants, in: *Journal of Physics: Conference
 Series*, vol. 524, p. 012186, IOP Publishing, 2014.
- IEA: Renewables 2022 - Analysis and forecast to 2027, 2022.
- 620 Jensen, N. O.: A note on wind generator interaction, vol. 2411, Citeseer, 1983.
- Kisvari, A., Lin, Z., and Liu, X.: Wind power forecasting—A data-driven method along with gated recurrent neural network, *Renewable
 Energy*, 163, 1895–1909, 2021.
- Kölle, K., Göçmen, T., Garcia-Rosa, P. B., Petrović, V., Eguinoa, I., Vrana, T. K., Long, Q., Pettas, V., Anand, A., Barlas, T. K., et al.:
 Towards integrated wind farm control: Interfacing farm flow and power plant controls, *Advanced Control for Applications: Engineering
 625 and Industrial Systems*, 4, e105, 2022.
- Larsen, Gunner C, A. M. H. and Bingöel, F.: Dynamic wake meandering modeling, 2007.
- Laves, M.-H., Ihler, S., Fast, J. F., Kahrs, L. A., and Ortmaier, T.: Well-calibrated regression uncertainty in medical imaging with deep
 learning, in: *Medical Imaging with Deep Learning*, pp. 393–412, PMLR, 2020.
- LeCun, Y., Bengio, Y., and Hinton, G.: Deep learning, *nature*, 521, 436–444, 2015.



- 630 Lee, J. C. and Fields, M. J.: An overview of wind energy production prediction bias, losses, and uncertainties, *Wind Energy Science Discussions*, 2020, 1–82, 2020.
- Lin, Z. and Liu, X.: Wind power forecasting of an offshore wind turbine based on high-frequency SCADA data and deep learning neural network, *Energy*, 201, 117 693, 2020.
- Liu, T., Huang, Z., Tian, L., Zhu, Y., Wang, H., and Feng, S.: Enhancing wind turbine power forecast via convolutional neural network,
 635 *Electronics*, 10, 261, 2021.
- Martínez-Tossas, L. A., Annoni, J., Fleming, P. A., and Churchfield, M. J.: The aerodynamics of the curled wake: a simplified model in view of flow control, *Wind Energy Science*, 4, 127–138, <https://doi.org/10.5194/wes-4-127-2019>, 2019.
- Martínez-Tossas, L. A., King, J., Quon, E., Bay, C. J., Mudafort, R., Hamilton, N., Howland, M. F., and Fleming, P. A.: The curled wake model: a three-dimensional and extremely fast steady-state wake solver for wind plant flows, *Wind Energy Science*, 6, 555–570,
 640 <https://doi.org/10.5194/wes-6-555-2021>, 2021.
- Meyers, J., Bottasso, C., Dykes, K., Fleming, P., Gebraad, P., Giebel, G., Göçmen, T., and van Wingerden, J.-W.: Wind farm flow control: prospects and challenges, *Wind Energy Science*, 7, 2271–2306, <https://doi.org/10.5194/wes-7-2271-2022>, 2022.
- Nejad, A. R., Keller, J., Guo, Y., Sheng, S., Polinder, H., Watson, S., Dong, J., Qin, Z., Ebrahimi, A., Schelenz, R., et al.: Wind turbine drivetrains: state-of-the-art technologies and future development trends, *Wind Energy Science*, 7, 387–411, 2022.
- 645 NREL: OpenFAST, <https://github.com/OpenFAST/openfast>, <https://github.com/OpenFAST/openfast>, a.
- NREL: FLORIS Wake Modeling and Wind Farm Controls Software, <https://github.com/NREL/floris> (FLORIS v3.4.1), <https://github.com/NREL/floris>, b.
- NREL: Simulator fOr Wind Farm Applications (SOWFA), <https://github.com/NREL/SOWFA>, <https://github.com/NREL/SOWFA>, c.
- Nygaard, N. G., Steen, S. T., Poulsen, L., and Pedersen, J. G.: Modelling cluster wakes and wind farm blockage, in: *Journal of Physics: Conference Series*, vol. 1618, p. 062072, IOP Publishing, 2020.
- 650 Park, J. and Park, J.: Physics-induced graph neural network: An application to wind-farm power estimation, *Energy*, 187, 115 883, <https://doi.org/10.1016/j.energy.2019.115883>, 2019.
- Perez-Sanjines, F., Verstraeten, T., Nowé, A., and Helsen, J.: Deep ensemble with Neural Networks to model power curve uncertainty, *Journal of Physics: Conference Series*, 2362, 012 029, <https://doi.org/10.1088/1742-6596/2362/1/012029>, 2022.
- 655 Pettas, V., Kretschmer, M., Clifton, A., and Cheng, P. W.: On the effects of inter-farm interactions at the offshore wind farm Alpha Ventus, *Wind Energy Science*, 6, 1455–1472, <https://doi.org/10.5194/wes-6-1455-2021>, 2021.
- Piotrowski, P., Rutyna, I., Baczyński, D., and Kopyt, M.: Evaluation metrics for wind power forecasts: A comprehensive review and statistical analysis of errors, *Energies*, 15, 9657, 2022.
- Pombo, D. V., Göçmen, T., Das, K., and Sørensen, P.: Multi-horizon data-driven wind power forecast: From nowcast to 2 days-ahead, in: 2021 International Conference on Smart Energy Systems and Technologies (SEST), pp. 1–6, IEEE, 2021.
- 660 Porté-Agel, F., Bastankhah, M., and Shamsoddin, S.: Wind-turbine and wind-farm flows: A review, *Boundary-layer meteorology*, 174, 1–59, 2020.
- Sanderse, B., Van der Pijl, S., and Koren, B.: Review of computational fluid dynamics for wind turbine wake aerodynamics, *Wind energy*, 14, 799–819, 2011.
- 665 Schepers, J.: WakeFarm: nabij zog model en ongestoord wind snelheidsveld, Energieonderzoek Centrum Nederland, 1998.
- Sood, I., Simon, E., Vitsas, A., Blockmans, B., Larsen, G. C., and Meyers, J.: Comparison of large eddy simulations against measurements from the Lillgrund offshore wind farm, *Wind Energy Science*, 7, 2469–2489, 2022.



- Strickland, J. M., Gadde, S. N., and Stevens, R. J.: Wind farm blockage in a stable atmospheric boundary layer, *Renewable Energy*, 197, 50–58, <https://doi.org/https://doi.org/10.1016/j.renene.2022.07.108>, 2022.
- 670 Ti, Z., Deng, X. W., and Zhang, M.: Artificial Neural Networks based wake model for power prediction of wind farm, *Renewable Energy*, 172, 618–631, 2021.
- University of Hannover: Parallelized Large-Eddy Simulation Model (PALM), <https://gitlab.palm-model.org>, https://gitlab.palm-model.org/releases/palm_model_system.
- Van Der Laan, M., Pena, A., Volker, P., Hansen, K. S., Sørensen, N. N., Ott, S., and Hasager, C. B.: Challenges in simulating coastal effects
 675 on an offshore wind farm, in: *Journal of Physics: Conference Series*, vol. 854, p. 012046, IOP Publishing, 2017.
- Verstraeten, T., Nowé, A., Keller, J., Guo, Y., Sheng, S., and Helsen, J.: Fleetwide data-enabled reliability improvement of wind turbines, *Renewable and Sustainable Energy Reviews*, 109, 428–437, 2019.
- Verstraeten, T., Daems, P.-J., Bargiacchi, E., Roijers, D. M., Libin, P. J., and Helsen, J.: Scalable optimization for wind farm control using coordination graphs, *arXiv preprint arXiv:2101.07844*, 2021.
- 680 Wang, H.-z., Li, G.-q., Wang, G.-b., Peng, J.-c., Jiang, H., and Liu, Y.-t.: Deep learning based ensemble approach for probabilistic wind power forecasting, *Applied energy*, 188, 56–70, 2017.
- Wang, Y., Zou, R., Liu, F., Zhang, L., and Liu, Q.: A review of wind speed and wind power forecasting with deep neural networks, *Applied Energy*, 304, 117 766, 2021.
- Yin, X. and Zhao, X.: Big data driven multi-objective predictions for offshore wind farm based on machine learning algorithms, *Energy*, 186, 115 704, <https://doi.org/https://doi.org/10.1016/j.energy.2019.07.034>, 2019.
- 685 Zehtabiyani-Rezaie, N., Iosifidis, A., and Abkar, M.: Physics-guided machine learning for wind-farm power prediction: Toward interpretability and generalizability, *PRX Energy*, 2, 013 009, 2023.