

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

Stijn Ally¹, Pieter-Jan Daems¹, Timothy Verstraeten¹, Ann Nowé¹, and Jan Helsen Vub¹

¹Vrije Universiteit Brussel

November 9, 2023

Abstract

Power production of offshore wind farms depends on many parameters and is significantly affected by wake losses. Due to the intermittency of wind power and its rapidly increasing share in the total energy mix, accurate forecasting of wind farm power production becomes increasingly important. This paper presents a data-driven methodology for forecasting power production and wake losses of wind farms, taking the dynamics of weather conditions into account. A modular approach is adopted by integrating multiple deep neural networks, resulting in a digital twin of the wind farm that can be interfaced with weather forecasts of different meteorological service providers. Another key advantage of the employed data-driven approach is its high prediction speed compared to physics-based methods, such that it can be employed for applications where real-time power forecasting is required. The methodology has been applied to two large offshore wind farms located within the Belgian-Dutch wind farm cluster in the North Sea.

RESEARCH ARTICLE

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

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

¹Artificial Intelligence Lab Brussels,
Vrije Universiteit Brussel, Belgium
²Faculty of Engineering, Vrije
Universiteit Brussel, Belgium

Correspondence

*Stijn Ally, Faculty of Engineering, Vrije
Universiteit Brussel, Pleinlaan 2, 1050
Elsene. Email: stijn.ally@vub.be

Abstract

Power production of offshore wind farms depends on many parameters and is significantly affected by wake losses. Due to the intermittency of wind power and its rapidly increasing share in the total energy mix, accurate forecasting of wind farm power production becomes increasingly important. This paper presents a data-driven methodology for forecasting power production and wake losses of wind farms, taking the dynamics of weather conditions into account. A modular approach is adopted by integrating multiple deep neural networks, resulting in a digital twin of the wind farm that can be interfaced with weather forecasts of different meteorological service providers. Another key advantage of the employed data-driven approach is its high prediction speed compared to physics-based methods, such that it can be employed for applications where real-time power forecasting is required. The methodology has been applied to two large offshore wind farms located within the Belgian-Dutch wind farm cluster in the North Sea.

KEYWORDS:

wind farm power forecasting, wake loss, deep learning, digital twin, virtual sensor, offshore wind farm

1 | INTRODUCTION

Over the last decades, renewables have been expanding quickly, but the global energy crisis has kicked them into an 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 growth¹.

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. In addition, as wind turbines and 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 an individual *wind turbine* depends on many parameters. Obviously, 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 an important influence on the power production, as well as wind shear and veer. Furthermore, a turbine can be derated due to technical reasons or be limited by a wind farm power controller.

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 deficit². 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 direction³, but also on other parameters, such as the wind turbulence intensity, atmospheric stability and surface roughness. The power loss of a downstream turbine can easily reach 40% in full-wake conditions³. 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^{4,5}. However, validation of the possible power gain under real-case conditions in a wind farm is still an active field of research^{6,7,8}.

The characteristics of the *wind inflow* into a wind farm (wind speed, 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^{9,10}. Wind blockage by neighbouring farms and the wind farm itself may also reduce the inflow speed and deflect the inflow direction^{9,11,12}. 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 land¹³.

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 turbines¹⁴. 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 (some examples: Jensen Park model¹⁵, FLOW Redirection and Induction in Steady State (FLORIS)¹⁶, curled wake model^{17,18} and TurbOPark¹⁹). In recent years, some further developments have made it possible to model also some heterogeneous and dynamic environmental conditions (e.g. FLORIDyn^{20,21}, UFloris²²). 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 models². 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 SOWFA²³, PALM²⁴ and SP-Wind²⁵. Finally, medium fidelity models (such as DWM²⁶, FAST.Farm²⁷, WakeFarm²⁸) are based on simplifications of the Navier-Stokes equations. Despite the simplifications, the computing time for medium fidelity models remains significant (order of minutes on PC)⁴.

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. Recent developments related to deep learning resulted in a significant leap forward in the modelling of large and complex data sets²⁹. In literature, different neural network architectures are described that model an individual wind turbine or a wind farm. These models differ in the selection of input features (often limited to wind speed), the data source used for the training of the model (often data generated by a physics-based model), limitations on the wind conditions (such as only steady-state conditions or only a few specific wind directions), length of the time steps, forecast horizon and complexity of the machine learning model (some examples: ^{30,31,32,33}). In some recent studies, some physics-based knowledge is integrated in the data-driven model, depicted as a physics-guided data-driven (PGDD) model, in an attempt to increase the interpretability and the generalization of the model to farm layouts or operating conditions other than those included in the training phase^{34,35,36,37,38}.

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 techniques³⁹. 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).

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 parameters that are required as input 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 model.

The methodology proposed in this paper is a data-driven technique that is not based on simplified physics-based equations, nor limiting assumptions about the temporal or spatial behaviour of the wind. It is based on deep learning models, trained solely with physical data measured at the wind farm itself. In contrast to most traditional models, this model is built for the complete wind farm as a whole, rather than for individual wind turbines. The resulting model is capable of predicting the total farm power and wake loss in a few milliseconds on PC, making it significantly faster than even physics-based low fidelity models. It takes into account a large number of different wind and farm input features, without simulating first the spatial wake profile throughout the farm. While this methodology is primarily intended for existing wind farms, it can also be trained with simulation data from physics-based high fidelity models. This may be useful, as the resulting trained neural network can predict the farm power under other weather conditions much faster than the high fidelity model itself.

Generally, when using neural networks, compared to physics-based models, the interpretability and explainability of the results is more challenging. Not only the accuracy of the model is essential to ensure reliable operation of wind farms being safety-critical systems, but insight in the uncertainty of the power predictions is crucial as well⁴⁰.

In order to maximize the interpretability and explainability of the proposed farm power model, only well-known physical features are used as input parameters for the model. Furthermore, farm-internal and farm-external influences on the farm power are modelled in separate sub-models, each having a simple neural network structure. In addition, the farm power model quantifies also the prediction uncertainty. Finally, for the validation of the model, not only traditional global performance metrics are used, but also the sensitivity for each input feature is analysed and represented in power curves overlaid onto corresponding measurement data. These show physically meaningful results which can be interpreted easily by wind energy professionals.

This work differs from other publications about farm power forecasting, as the methodology incorporates at the same time numerous wind and farm parameters, captures their temporal variations and is applied to multiple real-case large offshore wind farms. It enables wind farm owners to build their own dedicated farm model which can be interfaced with weather forecasts of different commercial weather forecast providers, without the need of additional expensive instrumentation or computing hardware. In addition, using an ensemble of multiple weather forecasts can possibly result in even more reliable farm power forecasts.

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 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 (ML) models are described in section 2.2.

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, 1-hour and 1-hour time resolution. Measurement data from local measurements is typically much 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. 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 1 minute is sufficient. Moreover, a 1-minute time resolution is also adequate for power forecasting applications related to trading on energy markets. Indeed, energy blocks are typically traded in 15-minutes or 1-hour blocks, with a lead time down to 5 minutes. Furthermore, a 1-minute power forecasting model may also prove useful for energy balancing applications.

In sections §2.1.1 to §2.1.5, each of the data sources used in the proposed methodology is described in more detail.

2.1.1 | Turbine SCADA data

The raw turbine SCADA data of the wind farms that are used as showcases in this paper, has a sampling time of 1 second. For each of the farms, the turbine data comprises the following measurements (among many others):

- 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 a turbine could technically actually produce (in case of sufficient wind):

- 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 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 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 produced by the farm can be calculated easily 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 information 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 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}
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 (almost) orthogonal to a side of the wind farm, all turbines in the 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 (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 with the highest wind speed are retained. The farm inflow wind parameters (inflow wind speed, inflow wind direction, inflow turbulence intensity and inflow wind direction variance) can then be 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 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. Due to the fast growing share of wind energy in the energy mix, the use of farm power control will 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 fleet 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

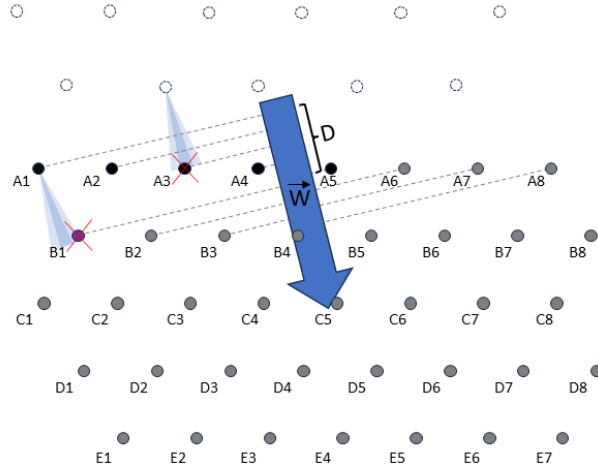


Figure 1 Determination of the upstream turbines of the wind farm with turbines A1 to E7, for the case of an average wind \vec{W} from north-northwest direction. Turbines A1 to A5 are located in the most upstream zone of the farm with length D along \vec{W} . As in this example the minimal quantity of 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).

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 multiple weather forecast providers have been used, each 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) and three-days-ahead forecasts.

2.1.5 | Historical weather data

In addition, ECMWF reanalysis v5 (ERA5) wind speed and direction data for the geographic location of the farms (center point of the farm layout) at a height of 100m has been used. ERA5 data has a time resolution of 1 hour.

2.1.6 | Train, test and validation data

After integrating the different data sources and selecting 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,




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.

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 precisely 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 ML model

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 may 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.

Time sequences from t_i to t_{i-T} of the following K input parameters are passed to the K convolution branches:

- 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),
- 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:

- 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*⁴¹. 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)⁴². For each time step t_i of the test data set, the following ratio is calculated:

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

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 Monte Carlo dropout method $\sigma_{\hat{P}_i}^2$ is re-calibrated as:

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

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:

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

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:

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

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

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.

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 consequently 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 (as far as the available data from the specific wind farm allows to identify such particular conditions).

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 the yaw control system will orientate the turbine perpendicular to the incoming wind direction (because for the wind farms used as examples in this paper no wake steering is done by applying yaw control).

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 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 an 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 direction. These auxiliary models are specific for a particular wind farm.

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. 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 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. As weather forecasts typically do not have a 1-minute time resolution, the mapping has been done with a 10-minute time resolution. For the farms used as examples in this paper, 1-hour weather forecasts have been interpolated to 10-minute forecasts, and the measured inflow wind speed and direction have been averaged to 10-minute data. For the training of the models, a training data set was generated containing for each time step only the most recent forecast for that time step (the forecast with shortest lead time), as that should be the most accurate forecast. In order to capture possible structural effects related to wind variations (such as delays), a sequence of five 10-minute forecast wind speed and direction data blocks centered around the particular time step is mapped to the corresponding measured ground-truth value.

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 a weather forecast provider 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 integrated 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.

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.

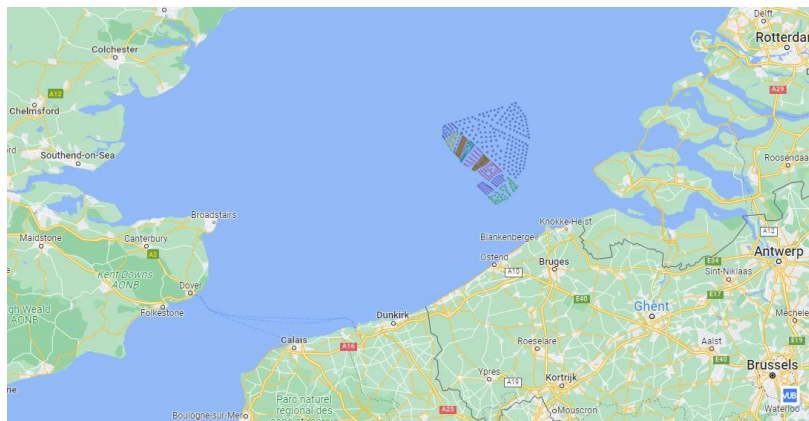


Figure 4 Position of the Belgian-Dutch offshore wind farm cluster in the North Sea

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

The results are grouped in the following three sections. In Section 3.1, results are presented related to the wind farm power model of both 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 with one of the two farm power models.

3.1 | Wind farm power model

3.1.1 | Data

For both Wind Farm 1 and Wind Farm 2, weather and turbine data was available for a period of about 2.5 years. Table 2 shows the quantity of 1-minute data points resulting from the preprocessing of the raw 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 consecutive days with strong wind fluctuations.

Table 2 Quantity of data points used for the two wind farm models

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

The data plots in figures 5a to 5c 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.

The polar data plots in figures 5d to 5f 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 wind turbulence 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. It can be noticed also 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, 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 air has a lower air density than dry air.

When comparing figures 5a and 5c 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 installed turbine capacity is never reached. This is due to the farm power controller, which limits the maximum producible power by the farm.

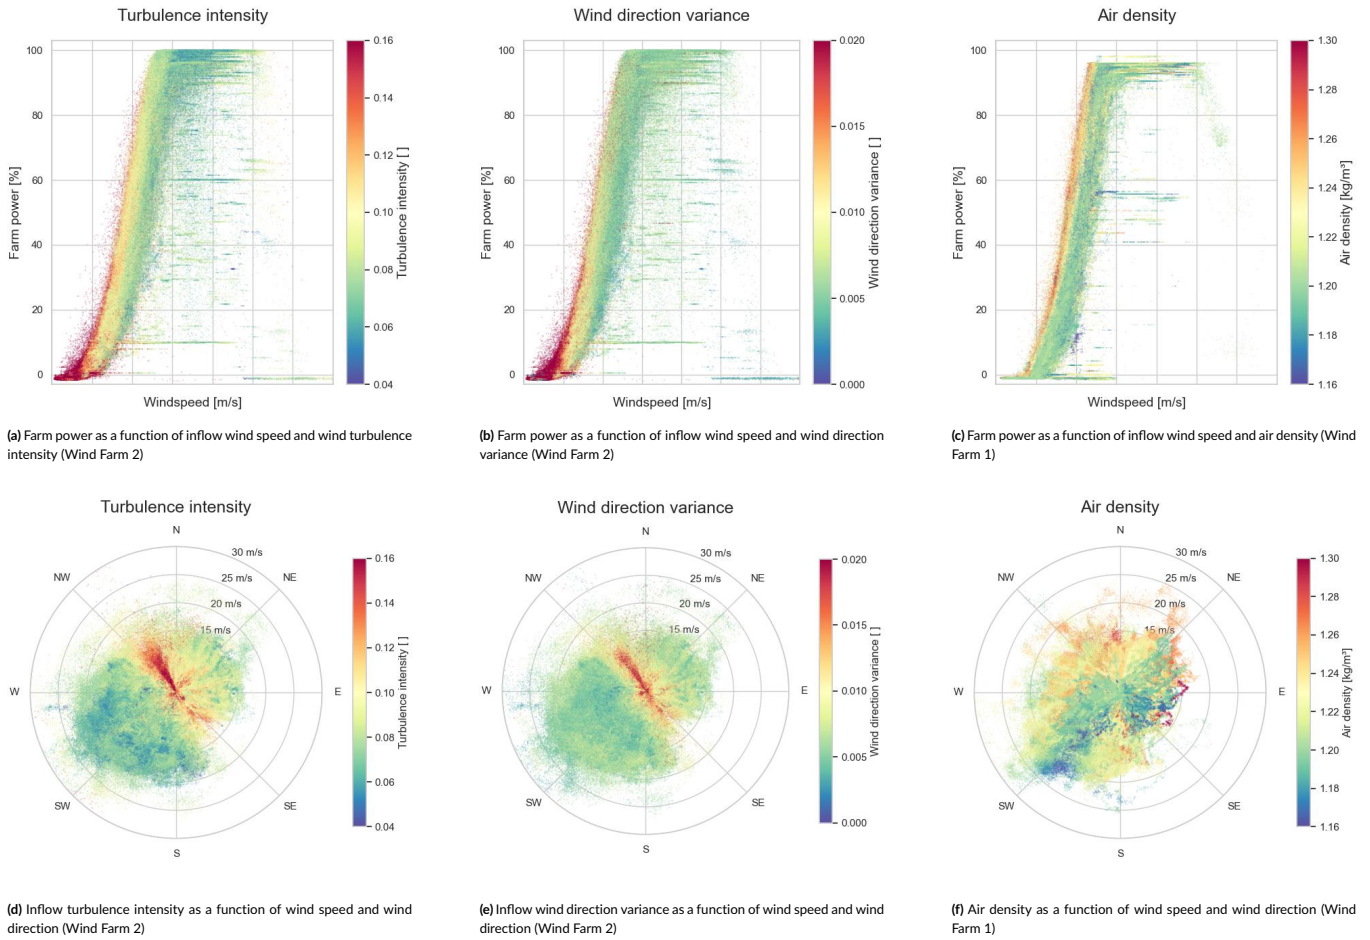


Figure 5 Inflow wind characteristics for Wind Farm 1 and Wind Farm 2 and their correlation with the farm power production

3.1.2 | Performance metrics

In Table 3, some performance metrics of the farm power models for Wind Farm 1 and Wind Farm 2 are listed: the root mean square error (RMSE), mean absolute error (MAE), mean error (ME) and the R2-score, for both the training and test data set. The mean error for each of the farms is near to 0% and the MAE is 2.42% and 2.14% of the installed capacity for the test data sets of Wind Farm 1 and Wind Farm 2 respectively. The difference between the performance metrics of the training and test data sets are very small, which is a sign for a satisfactory fit.

Table 3 Performance metrics of the models for Wind Farm 1 and Wind Farm 2.

	Wind Farm 1		Wind Farm 2	
	training	test	training	test
MAE	2.22 %	2.42 %	2.15 %	2.14 %
ME	-0.05 %	-0.17 %	-0.11 %	-0.10 %
RMSE	4.02 %	4.21 %	3.69 %	3.66 %
R2-score	0.99	0.99	0.99	0.99

Figure 6 gives 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

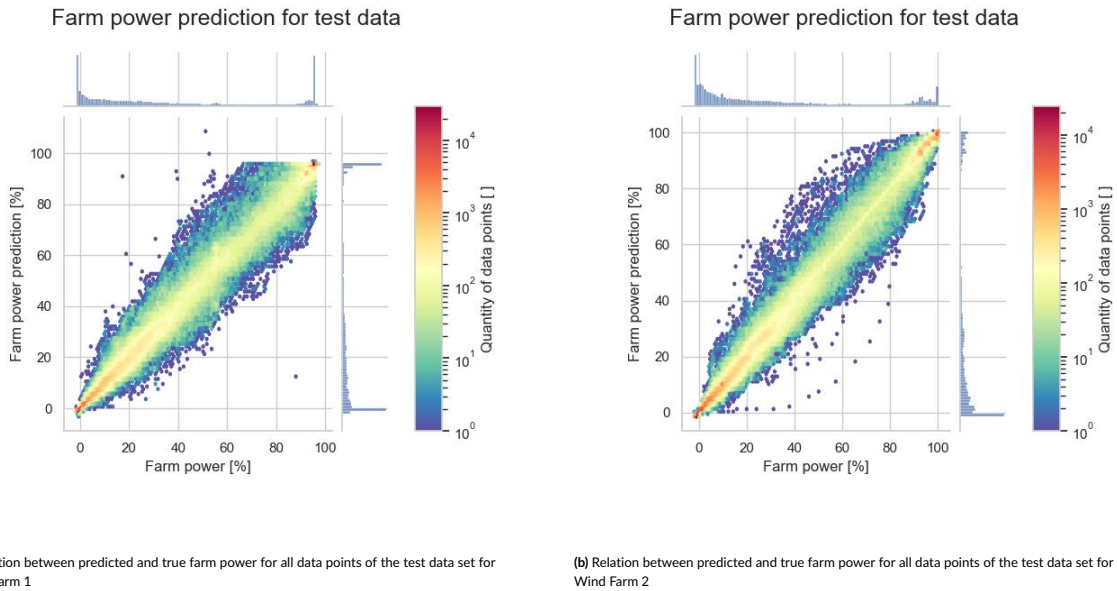


Figure 6 Relation between predicted and true farm power for all data points of the test data sets for Wind Farm 1 and Wind Farm 2




Figure 7 Predicted farm power and uncertainty intervals for a 60-hour time sequence of validation data (Wind Farm 1)

with wind speeds below the turbine cut-in speed and wind speeds above the rated turbine wind speed, where the power curve of the turbines is relatively flat.

3.1.3 | Validation time sequence

Figure 7 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 power curve 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 8 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 (when 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.

3.1.4 | Sensitivity analysis

Figure 9 illustrates the sensitivity of the farm power curve 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 based on the wind speed and a




Figure 8 Predicted farm power and uncertainty intervals for a 16-hour time sequence of test data with active power control (Wind Farm 2)

specific wind direction (see §2.2.4). 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.

Plots 9a to 9c 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 the data plots 5a to 5c. Plot 9d 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. Plot 9e shows the predicted power reduction for Wind Farm 2 when respectively one or two turbines are stopped. As mentioned in §3.1.1, the maximum power of this wind farm is limited by its farm controller. At high wind speeds, this maximum power limit can also almost be reached when one turbine is in standstill. Plot 9f 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°. The farm power for wind direction 260° is lower than for 325°, because of the higher farm-internal wake.

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, that cannot be modelled by traditional steady-state models. Two types of wind speed variations are tested: linear wind speed ramps and sinusoidal varying wind speeds with different frequencies.

Figure 10a shows power predictions for Wind Farm 1 for the same wind directions as shown in plot 9d. 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 10b shows the frequency response of the 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




Figure 9 Sensitivity analysis of the wind farm power models for different input parameters. Each plot shows predicted farm power curves and ground truth data points of which the features are similar to the input parameters of the power curve with the same color.

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.

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 higher than the oscillation period of the sinusoidal wind speed component. For higher oscillation periods, the time offset converges to a constant which seems to depend on the average wind speed, length of the cross-section of the farm and the magnitude of the wake. Logically, the time delay will never be higher than the time needed by the wind to cross the complete farm.

For small 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

As already shown in Figure 9d, 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.

Figure 11a shows the measured farm power for Wind Farm 1, as a function of the wind direction and wind speed. Plot 11b shows the subset of these data points for which the set point of the farm power controller is equal to 100% and not any or only one turbine is in standstill. Furthermore,




Figure 10 Dynamic behaviour of the wind farm models in case of linear wind speed ramps and sinusoidal oscillating wind speeds.

the scope of the plot has been limited to the region where wake is most predominant. Figure 11c 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 lower than for other wind directions. This can be seen more clearly after subtraction from the predicted turbine power. Figure 11d 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, and consequently each turbine in the farm can produce sufficient power so that the maximum farm power is reached (as mentioned in section §3.1.1, the maximum farm power for Wind Farm 1 is limited by a farm power controller).

3.3 | Power forecasting based on weather forecasts

The models presented in above sections, predict the farm power based on historical values of a set of wind characteristics and other parameters of the farm. Obviously, in order to forecast the farm power in the future with these farm models, forecasts of the wind speed and direction (and if possible also the other wind characteristics) are required.

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 12a shows the correction factor to be applied to the wind speed forecasts of a specific weather forecast provider for 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 factor is higher than one. This is due to the fact that the wind measurements on the turbines are not well calibrated and over-estimating these low wind speeds. For such wind speeds below the cut-in wind speed, turbines are shut off anyway.

In Figure 12b, 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, 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

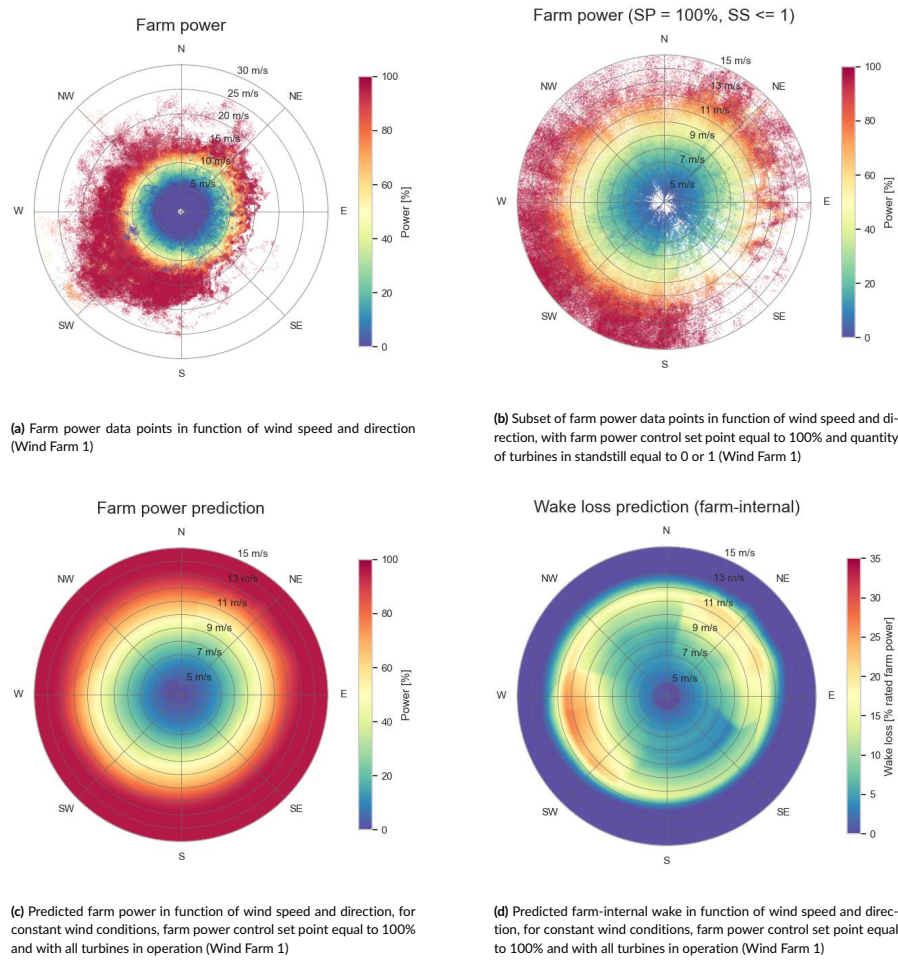


Figure 11 True and forecast wind power and farm-internal wake in function of wind speed and direction

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.

In addition, in Figure 12a 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 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 indication of the acceleration of the air flow at the corner of the cluster occurring jointly with the deflection due to the blockage effect. (Remind that the farm inflow is calculated by averaging the wind parameters measured at the most upstream turbines, thus the turbines located in the south-southwest corner for that wind direction (see 2.1.2)). 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 other weather forecasts providers and for historical ERA5 data, similar discrepancies in wind speed and wind directions are observed, however with different biases and/or variances. Anyhow, in order to forecast the farm power accurately, it is not crucial to identify exactly the origin(s) of these discrepancies.

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 13 shows the three-days-ahead, day-ahead and intra-day power forecasts for an eight-days sequence for Wind Farm 1. As weather forecasts with a shorter lead-time are usually more accurate than with longer lead-times, the resulting power forecasts are getting more accurate for shorter lead times as well. However, when comparing the accuracy of the forecasts in Figure 13 with the prediction accuracy of the farm power model shown in Figure 7, it is clear that the inaccuracy introduced by the weather forecasts is higher than the uncertainty inherent to the farm power model. This should be no surprise, as the weather forecasts used in this example have a time resolution of only a 1-hour.

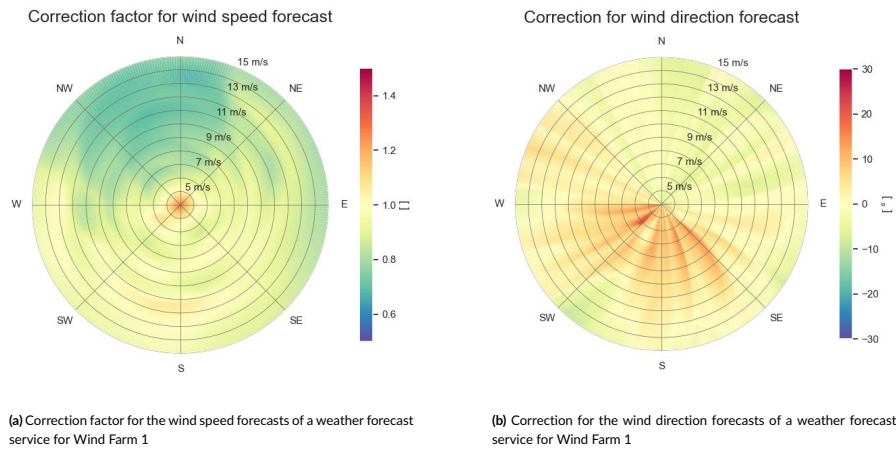


Figure 12 Corrections for the wind speed and direction forecasts of a specific weather forecast service for Wind Farm 1




Figure 13 Power forecasts for Wind Farm 1 based on wind forecasts with different forecast horizons

Furthermore, the used weather forecasts are updated only twice each day and, consequently, even the intra-day forecasts can have a lead time up to half a day. Obviously, using more accurate weather forecasts with shorter forecast horizons will result in better farm power forecasts.

4 | CONCLUSION

In the present paper, a novel methodology is proposed for forecasting the power production of a wind farm and predicting the wake loss. With a modular deep learning approach, a digital twin of the wind farm is created, which can be interfaced with weather forecasts of different meteorological service providers.

In contrast to many other traditional methods, this data-driven technique is not limited to homogeneous and steady-state wind inflows and does not depend on simplifications of physics-based equations, nor on technical data sheets of the turbines that are only valid for specific operation conditions.

The proposed farm power model relies solely on SCADA data from the wind farm itself, that is usually available to any wind farm operator. It captures the dynamics of several weather parameters, including wind speed, wind direction, turbulence intensity, wind direction variance and air density. Additionally, it models the behaviour of the farm power controller and takes the number of turbines in standstill 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 model is capable of predicting the total farm power and wake loss in only a few milliseconds on PC, making it significantly faster than even physics-based low-fidelity models.

A separate deep learning model is employed to generate corrections to the wind speed and direction forecasts from weather forecasts 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.

The proposed methodology has been applied to two large real-world offshore wind farms. Multiple performance metrics affirm a high prediction accuracy for both wind farms. 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. Additionally, the dynamic behaviour of the farm models has been tested by simulating wind speed ramps and sinusoidal fluctuating wind speeds. These dynamic scenarios result in additional insights into the dynamic behaviour of the wind farms.

In further research, we will integrate the wind farm models as digital twins into applications where their high predictive speed is crucial.


ACKNOWLEDGEMENTS

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

PEER REVIEW

The peer review history for this article is available at ...

ORCID

Stijn Ally  <https://orcid.org/0009-0009-0921-4914>

References

1. IEA . Renewables 2022 - Analysis and forecast to 2027. 2022.
2. Sanderse B, Pijl V. dS, Koren B. Review of computational fluid dynamics for wind turbine wake aerodynamics. *Wind energy* 2011; 14(7): 799–819.
3. Barthelmie RJ, Hansen K, Frandsen ST, et al. 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* 2009; 12(5): 431–444.
4. Boersma S, Doekemeijer BM, Gebraad PM, et al. A tutorial on control-oriented modeling and control of wind farms. In: IEEE. ; 2017: 1–18.
5. Verstraeten T, Daems PJ, Bargiacchi E, Roijers DM, Libin PJ, Helsen J. Scalable optimization for wind farm control using coordination graphs. *arXiv preprint arXiv:2101.07844* 2021.
6. Bossanyi E, Ruisi R. Axial induction controller field test at Sedini wind farm. *Wind Energy Science* 2021; 6(2): 389–408. doi: 10.5194/wes-6-389-2021
7. Fleming PA, Ning A, Gebraad PM, Dykes K. Wind plant system engineering through optimization of layout and yaw control. *Wind Energy* 2016; 19(2): 329–344.
8. Fleming P, King J, Simley E, et al. Continued results from a field campaign of wake steering applied at a commercial wind farm – Part 2. *Wind Energy Science* 2020; 5(3): 945–958. doi: 10.5194/wes-5-945-2020
9. Porté-Agel F, Bastankhah M, Shamsoddin S. Wind-turbine and wind-farm flows: A review. *Boundary-layer meteorology* 2020; 174(1): 1–59.
10. Pettas V, Kretschmer M, Clifton A, Cheng PW. On the effects of inter-farm interactions at the offshore wind farm Alpha Ventus. *Wind Energy Science* 2021; 6(6): 1455–1472. doi: 10.5194/wes-6-1455-2021
11. Bleeg J, Purcell M, Ruisi R, Traiger E. Wind farm blockage and the consequences of neglecting its impact on energy production. *Energies* 2018; 11(6): 1609.

12. Strickland JM, Gadde SN, Stevens RJ. Wind farm blockage in a stable atmospheric boundary layer. *Renewable Energy* 2022; 197: 50-58. doi: <https://doi.org/10.1016/j.renene.2022.07.108>
13. Van Der Laan M, Pena A, Volker P, et al. Challenges in simulating coastal effects on an offshore wind farm. In: . 854. IOP Publishing. ; 2017: 012046.
14. Nejad AR, Keller J, Guo Y, et al. Wind turbine drivetrains: state-of-the-art technologies and future development trends. *Wind Energy Science* 2022; 7(1): 387–411.
15. Jensen NO. *A note on wind generator interaction*. 2411. Citeseer . 1983.
16. NREL . FLORIS Wake Modeling and Wind Farm Controls Software, <https://github.com/NREL/floris> (FLORIS v3.4.1).
17. Martínez-Tossas LA, Annoni J, Fleming PA, Churchfield MJ. The aerodynamics of the curled wake: a simplified model in view of flow control. *Wind Energy Science* 2019; 4(1): 127–138. doi: 10.5194/wes-4-127-2019
18. Martínez-Tossas LA, King J, Quon E, et al. The curled wake model: a three-dimensional and extremely fast steady-state wake solver for wind plant flows. *Wind Energy Science* 2021; 6(2): 555–570. doi: 10.5194/wes-6-555-2021
19. Nygaard NG, Steen ST, Poulsen L, Pedersen JG. Modelling cluster wakes and wind farm blockage. In: . 1618. IOP Publishing. ; 2020: 062072.
20. Gebraad PM, Van Wingerden J. A control-oriented dynamic model for wakes in wind plants. In: . 524. IOP Publishing. ; 2014: 012186.
21. Becker M, Allaerts D, Wingerden vJW. FLORIDyn - A dynamic and flexible framework for real-time wind farm control. *Journal of Physics: Conference Series* 2022; 2265(3): 032103. doi: 10.1088/1742-6596/2265/3/032103
22. Foloppe B, Munters W, Buckingham S, Vandeveld L, Beeck vJ. Coupling of a dynamic wake model with WRF: a case study of the Belgian wind farms. In: ; 2022: 30–31.
23. NREL . Simulator fOr Wind Farm Applications (SOWFA), <https://github.com/NREL/SOWFA>.
24. Hannover oU. Parallelized Large-Eddy Simulation Model (PALM), <https://gitlab.palm-model.org>.
25. Sood I, Simon E, Vitsas A, Blockmans B, Larsen GC, Meyers J. Comparison of large eddy simulations against measurements from the Lillgrund offshore wind farm. *Wind Energy Science* 2022; 7(6): 2469–2489.
26. Larsen AMH, Bingöl F. Dynamic wake meandering modeling. 2007.
27. NREL . OpenFAST, <https://github.com/OpenFAST/openfast>.
28. Schepers J. *WakeFarm: nabij zog model en ongestoord wind snelheidsveld*. Energieonderzoek Centrum Nederland . 1998.
29. LeCun Y, Bengio Y, Hinton G. Deep learning. *nature* 2015; 521(7553): 436–444.
30. Yin X, Zhao X. Big data driven multi-objective predictions for offshore wind farm based on machine learning algorithms. *Energy* 2019; 186: 115704. doi: <https://doi.org/10.1016/j.energy.2019.07.034>
31. Perez-Sanjines F, Verstraeten T, Nowé A, Helsen J. Deep ensemble with Neural Networks to model power curve uncertainty. *Journal of Physics: Conference Series* 2022; 2362(1): 012029. doi: 10.1088/1742-6596/2362/1/012029
32. Liu T, Huang Z, Tian L, Zhu Y, Wang H, Feng S. Enhancing wind turbine power forecast via convolutional neural network. *Electronics* 2021; 10(3): 261.
33. Pombo DV, Göçmen T, Das K, Sørensen P. Multi-horizon data-driven wind power forecast: From nowcast to 2 days-ahead. In: IEEE. ; 2021: 1–6.
34. Sun H, Qiu C, Lu L, Gao X, Chen J, Yang H. Wind turbine power modelling and optimization using artificial neural network with wind field experimental data. *Applied Energy* 2020; 280: 115880.
35. Zehtabiyani-Rezaie N, Iosifidis A, Abkar M. Physics-guided machine learning for wind-farm power prediction: Toward interpretability and generalizability. *PRX Energy* 2023; 2(1): 013009.

36. Ti Z, Deng XW, Zhang M. Artificial Neural Networks based wake model for power prediction of wind farm. *Renewable Energy* 2021; 172: 618–631.
37. Howland MF, Dabiri JO. Wind farm modeling with interpretable physics-informed machine learning. *Energies* 2019; 12(14): 2716.
38. Park J, Park J. Physics-induced graph neural network: An application to wind-farm power estimation. *Energy* 2019; 187: 115883. doi: <https://doi.org/10.1016/j.energy.2019.115883>
39. Verstraeten T, Nowé A, Keller J, Guo Y, Sheng S, Helsen J. Fleetwide data-enabled reliability improvement of wind turbines. *Renewable and Sustainable Energy Reviews* 2019; 109: 428–437.
40. Meyers J, Bottasso C, Dykes K, et al. Wind farm flow control: prospects and challenges. *Wind Energy Science* 2022; 7(6): 2271–2306. doi: 10.5194/wes-7-2271-2022
41. Gal Y, Ghahramani Z. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In: Balcan MF, Weinberger KQ., eds. *Proceedings of The 33rd International Conference on Machine Learning*. 48 of *Proceedings of Machine Learning Research*. PMLR; 2016; New York, New York, USA: 1050–1059.
42. Laves MH, Ihler S, Fast JF, Kahrs LA, Ortmaier T. Well-calibrated regression uncertainty in medical imaging with deep learning. In: PMLR. ; 2020: 393–412.

How to cite this article: Ally S, Daems PJ, Verstraeten T, Nowé A and Helsen J (2023), Modular deep learning approach for wind farm power forecasting and wake loss prediction, *Wind Energy*, !!2023;00:1–6.