RESEARCH ARTICLE

Statistical Post-processing of Turbulence-resolving Weather Forecasts for Offshore Wind Power Forecasting

Ciaran Gilbert¹ | Jakob W. Messner² | Pierre Pinson² | Pierre-Julien Trombe³ | Remco Verzijlbergh^{4,5} | Pim van Dorp^{4,6} | Harmen Jonker^{4,6}

¹University of Strathclyde, Glasgow, UK

²Technical University of Denmark, Denmark

³Vattenfall Vindkraft A/S, Esbjerg, Denmark

⁴Whiffle Weather Finecasting Ltd,
Molengraaffsingel 12, 2629 JD Delft, the Netherlands

⁵Delft University of Technology, Department of Engineering Systems & Services, Jaffalaan 5, 2628 BX Delft, the Netherlands

⁶Delft University of Technology, Department of Geosciences & Remote Sensing, Stevinweg 1, 2628 CN Delft, the Netherlands

Correspondence

C. Gilbert, Technology and Innovation Centre, University of Strathclyde, 99 George Street, Glasgow G1 1RD, Scotland. Email: ciarangilbert@strath.ac.uk

Abstract

Accurate short-term power forecasts are crucial for a reliable and efficient integration of wind energy in power systems and electricity markets. Typically, forecasts for hours to days ahead are based on the output of Numerical Weather Prediction models, and with the advance of computing power the spatial and temporal resolution of these models have increased substantially. The increasing ability to resolve smaller scale processes provides great potential for forecast improvements. However, high-resolution forecasts often exhibit spatial and/or temporal displacement errors and when regarding performance metrics such as root mean square error they often perform worse than smoother forecasts from lower-resolution models. Recent computational advances have enabled the use of Large Eddy Simulation (LES) in the context of operational weather forecasting, yielding turbulence resolving weather forecasts with a spatial resolution of 100 meters or finer and a temporal resolution of 30 seconds or less. This study investigates how these ultra-high resolution weather models can be used efficiently for forecasting at Horns Rev I in Denmark. It is shown that temporal smoothing of the forecasts clearly improves their skill, even for the benchmark resolution forecast, although potentially valuable high-frequency information is lost. Therefore, a statistical post-processing approach is explored based on smoothing and feature engineering from the high-frequency signal. The results indicate that for wind farm forecasting, using information content from both the standard and LES resolution models improves the forecast accuracy, especially with a feature selection stage, compared to using the information content solely from either source.

KEYWORDS:

forecasting, wind power, large eddy simulation, post-processing

1 | INTRODUCTION

The penetration of wind power in certain electric power networks around the globe has reached significant levels. At least in the near term future, the contribution of this stochastic weather-dependent power source in the overall energy mix is set to grow and efficiently integrating this power is critical to maintain an economic and secure source of supply¹. Within this substantial and challenging field, energy forecasting is essential for decision making when future generation is a factor². Applications include trading in balancing or day-ahead markets, balancing supply and demand, energy storage management, as well as other developing research themes such as wind farm operation and maintenance, and the role of predictability in investment decision making^{3,4}.

The time horizon of the forecast defines the suitability of methods and data for creating an informative model. Very short term forecasts (0– $\approx$ 6 hours ahead) are usually built from purely statistical models using information from individual or spatially distributed power time series. Looking further ahead in the future ($>\approx$ 6 hours ahead), the use of Numerical Weather Prediction (NWP) has shown to give higher forecasting skill⁵. Conventionally, power forecasts for these lead times are generated using statistical learning techniques to map the relationship between the measured power time series and concurrent explanatory variables engineered from the NWP. This general methodology implicitly accounts for site-specific effects such as turbine degradation, as well as any systematic NWP biases present and is therefore recommended best practice⁶. The pursuit of high resolution forecasts and appraising their value for wind power forecasting is not a new topic and has been discussed in literature with varying success^{5,7}. The situation has evolved substantially over the last few years, thanks to quantity of data being collected, increase in computational power (e.g. based on Graphics Processing Units – GPUs) and advances in data science. Consequently, this enables the high-resolution forecasting framework considered here, i.e., based on computationally demanding approaches to producing high-resolution weather forecasts at wind farm level, and post-processing stage via advanced statistical learning techniques.

In recent years, academic research has been focused on quantifying the uncertainty associated with energy forecasts due to the underlying chaotic weather systems driving the energy generation⁸. These uncertainty forecasts are useful for applications such operating reserve management⁹ and defining optimal bidding strategies¹⁰. However, currently practitioners prefer deterministic (point) forecasts of future generation which comprise of single valued best estimates. This is due to the ease of interpretation and incorporation into existing decision making systems, as well poor communication of the underlying information content of an uncertainty forecast⁶. Therefore improving point forecasts is still a relevant research pursuit. For comprehensive reviews of short-term wind power forecasting please refer to^{11,5}.

NWP models have been steadily improving over the last 40 years in terms of both forecast skill and temporal/spatial resolution due to advances in scientific research, computing capability, and the expansion of available observations used to initialise atmospheric models¹². Currently, the European Centre for Medium-Range Forecasting (ECMWF) offers global weather forecasts at a spatial resolution of $\approx$ 16km at 1–6 hour intervals within the high resolution 10-day forecast model. Wind speed forecasts retrieved directly from the ECMWF fields therefore have no resolved variance below these spatial and temporal scales. A common strategy to improve global weather forecasts is to downscale them using NWP models that cover only a limited area and are therefore also referred to as Limited Area Models (LAM) or mesoscale models. Mesoscale models have the advantage of representing the earth's surface in higher detail and resolve physical processes in higher temporal detail, but the formulation of physical processes like turbulence and (cloudy) convection is essentially the same as those of the global NWP models.

Large-eddy simulations (LES), on the other hand, have a typical resolution of 10m–100m or less and therefore directly resolve turbulence and boundary layer clouds. Furthermore, because a typical LES resolution allows for wind turbine resolving simulations, the use of LES to study flows through wind farms has received a lot of attention in the scientific literature lately, see for example^{13,14} and¹⁵. Until recently, though, the high spatial and temporal (roughly 5–10 seconds) resolution of LES prohibited their use in an operational forecast setting due to their high computational cost and excessively long run-times. Advances in computing like the use of GPUs to accelerate LES computations have, however, drastically shortened the run-times and have paved the way for operational weather forecasting using LES^{16,17}. A turbine resolving operational weather forecast on wind farm scale is therefore currently possible and raises the question how one can best improve these forecasts by statistical post-processing.

Increasing the resolution of the NWP model does not immediately translate to improvements in forecast skill, which is usually measured on the average performance, e.g., root mean square error (RMSE) and mean absolute error (MAE). A forecast which may represent better the inherent behaviour of the underlying processes, can then be heavily penalised due to phase and location errors where the exact timing or placement of any particular weather event is missed. As a result, high resolution forecasts can often perform worse in terms of RMSE or MAE when compared to low resolution forecasts. Therefore, the research aim and contribution of this paper is simple: to ascertain if it is possible to extract value from a high resolution NWP model in the context of wind power forecasting.

We therefore regard point forecasts of wind speed and power at the iconic Horns Rev I 160 MW wind farm off the coast of Denmark and compare the forecasting ability of the high resolution LES model with forecasts from ECMWF. In the post-processing of the raw NWP data into power and wind speed forecasts, smoothing techniques are used to account for the penalisation of temporal and spatial displacement errors. The disadvantage of smoothing is that high resolution information can be diminished; to remedy this, a simple feature engineering stage is proposed which can account for the variability in the atmospheric model. This is informed by¹⁸ where feature engineering is used to improve forecast skill for single solar and wind generating sites using information from neighbouring grid points in a typical NWP model. This work is intended as an initial feasibility study into the value of these forecasts in power/wind forecasting for a day ahead with typical statistical learning techniques used presently. Future studies should consider optimal frameworks for using the high spatial temporal resolution data such as hierarchical models or deep learning to extract useful features from the high frequency information.

This paper is organised as follows; Section 2 details the specifics of the NWP models used and Horns Rev I and Section 3 includes descriptions of statistical post-processing techniques. Section 4 presents and evaluates the results of the LES and the validation case study based on measurements from a individual turbine and subsequently the overall wind farm, and finally Section 5 details the conclusions and future work.

2 | CASE STUDY AT THE HORNS REV OFFSHORE WIND FARM

To evaluate the performance of the post-processed high resolution forecast generated from the LES model, results are presented for Horns Rev I and compared with forecasts post-processed from ECMWF meteorological forecasts. Horns Rev I was one of the first larger size offshore wind farms and therefore has a large bank of historical data. Both wind speed and wind power post-processing are considered as these are both extremely important for operational tasks such as trading and operation and maintenance (O&M).

2.1 | Horns Rev

Out of the 3 offshore wind farms located at Horns Rev (2 operational ones and one expected to start operation in 2019), we focus on that which was installed first, in operation since 2002) and commonly referred to as Horns Rev I. Horns Rev I consists of 80 2-MW turbines located approximately 18km off the coast of Denmark in the North Sea covering an area of around 21km². It was historically the first large offshore wind farm worldwide. The data used for the following study consists of average power generation and wind speed anemometer measurements from each constituent turbine from January 2015 to December 2016 at a 10 minute temporal resolution. Initially, to reduce the dimensions of this exploratory analysis, much of the post-processing results presented are for an individual turbine located at the north west corner of the park with a final analysis of the wind farm power as whole. It should be noted that a single turbine located at the south west corner of the wind farm is excluded from the analysis.

2.2 | Weather forecasts: from ECMWF to the Whiffle Large Eddy Simulation

The operational ECMWF deterministic forecasts covering the years 2015 and 2016 were retrieved at 3 hourly time resolution and approx 16km spatial resolution. Note, that forecasts are also available in 1 hour resolution but usually ECMWF offers only the 3 hour resolution data to commercial customers. To make the forecasts comparable to the turbine data, the forecast fields were interpolated in time using linear interpolation and model levels were interpolated to a hub height of 70 meters by linear interpolation of the closest model levels¹⁹.

Whiffle is a start-up that has focused on the computation of LES on Graphics Processing Units (GPUs), which enables their model GRASP (GPU-Resident Atmospheric Simulation Platform) to compute 48h hour ahead forecast within roughly an hour. The formulation of the model is provided in²⁰ and more recent features, including the GPU implementation and the method to drive an LES model with large-scale boundary conditions from a NWP model have been described in¹⁷ and²¹.

The simulation domain for this study is 8.2 km in the horizontal and 5 km in the vertical with 128³ grid points, yielding a horizontal resolution of 64m. In the vertical direction, the grid is stretched, with a resolution of roughly 16m near the surface and 80m near the top of the domain. Boundary conditions are taken from the ECMWF deterministic forecasts. Given the dissemination schedule of the ECMWF operational forecasts, this means that a day-ahead forecast can be delivered around 07:00 UTC, which corresponds to 08:00 CET and 09:00 CEST. The LES forecast is therefore well in time to be used for day-ahead trading.

Wind turbines have been modelled in the LES with a uniform actuator disk model as described by²². In addition, a concurrent precursor simulation has been run to prevent recirculation of the turbines wakes in the LES domain²³. The turbine parametrisation applies axial and radial forces to the LES wind fields that are based on the power and thrust curves of the turbine type. Therefore, the torque exerted on the rotor blades can readily be diagnosed from the simulation and the produced power can be exported as output variable with the same time resolution as the LES time-step, which is typically in the order of 5 seconds. All output variables that may be relevant for post-processing, such as power, rotor-disk averaged wind speed and direction and air density have been exported for each individual turbine location with a 30 seconds time resolution. Care was taken to produce a forecast dataset that was as representative as possible for day-ahead forecasting, so a maximum computation time of 1 hour per forecast was observed and the 00UTC ECMWF operational forecasts valid for the next 24-48h (local time) were used as boundary conditions.

Figure 1a shows a snapshot of the instantaneous wind speed field at hub height for a typical day with south-westerly winds. Using the pre-cursor simulation that has no wind turbines, it is also possible to calculate the difference between the free-stream wind speed and the wind speed in the wind farm. Figure 1b shows the wind speed deficit, i.e. the percent reduction in wind speed with respect to the free stream wind, in the simulation domain. The wind speed deficits for this particular day are roughly 20% to 40%, which is in agreement with values that have been reported in the literature²⁴.

Although the focus of this paper is on the prediction of wind speeds and wind farm power, the framework of a turbine resolving LES model driven by NWP boundary conditions opens more possibilities for future research and practical applications. Currently ongoing research efforts focus of wind resource assessments and annual energy prediction using LES with reanalysis fields as boundary conditions, air-sea interactions in and around wind farms, more advanced turbine parametrizations and in-depth analysis of the interactions of (far) wake effects and their environment.

FIGURE 1 Left: Instantaneous wind speed field for 2016-05-18 02:00 UTC. The turbine locations and their yaw angles are depicted with the grey bars. Right: Two hour average wind speed deficit (percent reduction with respect to free-stream wind speed) field for 2016-05-18 00UTC-02UTC

3 | POST-PROCESSING MODELS

The following describes the statistical post-processing models utilised and the general post-processing strategy for the case study. The approach can be summarised by using optimal statistical learning techniques for both wind speed and power forecasting with a focus on smoothing, feature engineering, and selection. Smoothing is an interesting approach in this high-resolution forecasting context, however it has been surprisingly overlooked in the applicable research literature for energy forecasting. Benefits have been reported with temporal and spatial feature engineering in the post-processing of coarse global atmospheric models^{18,25,26} with rolling average and leading/lagging wind speed variables among others. Therefore the aim here is to explore this relationship more explicitly with single input models and then to capture the value in the high-frequency content of the signal with other engineered features. This approach, compared to using the high resolution wind speed signal directly, aims to retain the value of the smoothed wind speed forecast whilst augmenting the model with selected summary statistics of the high frequency signal. The methodologies described are implemented in R using the packages `crch` and `h2o`^{27,28,29}.

3.1 | Wind speed forecasting

For wind speed forecasting the raw NWP signals are statistically post-processed to account for systematic biases between each signal and the measured time series at the turbine nacelle. A truncated linear regression model is used to account for the non-negativity of wind speed³⁰ which is defined as

$$y_t \sim \mathcal{N}_0(\mu_t, \sigma^2) \quad (1)$$

$$\mu_t = \beta_0 + \beta_1 x_t \quad (2)$$

where wind speed y_t at time t follows a zero truncated normal distribution with mean μ_t and standard deviation σ while μ_t is a linear function of the predicted wind speed x_t and β_0, β_1 , and σ are regression coefficients. The zero truncated normal distribution has a probability density function of

$$d(y_t) = \begin{cases} \frac{1}{\sigma} \phi\left(\frac{y_t - \mu_t}{\sigma}\right) & y_t > 0 \\ 0 & y_t \leq 0 \end{cases} \quad (3)$$

where $\phi(\cdot)$ is the probability density function of the standard normal distribution. The regression coefficients are estimated with maximum likelihood estimation as implemented in the R software package `crch`²⁸.

3.2 | Power forecasting

The Gradient Boosting Machine (GBM) is a supervised learning method whereby a series of individually weakly predictive regression base-learners (e.g. regression trees) are combined to make a powerfully predictive ensemble^{31,32}. Specifically for regression trees, the available input variable space is partitioned into disjoint regions which allows for the direct capture of non-linear relationships, such as the wind power curve, and makes this a powerful tool for energy forecasting applications^{18,25,26}. Additionally, the flexibility of the algorithm in terms of loss functions is a desirable attribute; in this study a squared loss function is used. For target variable y and a pool of explanatory variables $\mathbf{x}_t = (x_1, x_2, \dots)^\top$, the gradient boosting machine^{31,33} $F_N(\mathbf{x}_t)$ is defined as the sum of N base-learners $f_n(\mathbf{x}_t)$

$$y_t = F_N(\mathbf{x}_t) + \epsilon_t = \sum_{n=0}^N f_n(\mathbf{x}_t) + \epsilon_t \quad (4)$$

where $f_0(\mathbf{x}_t)$ is the initialisation guess, and the subsequent ensemble of base-learners is constructed sequentially by estimating the latest via

$$\operatorname{argmin}_{f_n} \sum_t L(y_t, F_{n-1}(\mathbf{x}_t) + f_n(\mathbf{x}_t)) \quad (5)$$

for some loss function (e.g., squared loss) $L(\cdot)$. To tackle this approach in practice, the base learner chosen here is the regression tree $f_n = h(\mathbf{x}; \theta_n)$, specified by a vector of tree parameters θ_n . Where $L(\cdot)$ is differentiable, the negative gradient $g_n(\mathbf{x})$ is defined as

$$g_n(\mathbf{x}_t) = - \left[\frac{\partial L(y_t, F_n(\mathbf{x}_t))}{\partial F_n(\mathbf{x}_t)} \right]_{F_n(\mathbf{x})=F_{n-1}(\mathbf{x})} \quad (6)$$

and the regression tree is efficiently fit to this negative gradient by least squares

$$\theta_n = \operatorname{argmin}_{\theta} \sum_t [g_n(\mathbf{x}_t) - h(\mathbf{x}_t; \theta)]^2 \quad (7)$$

The ensemble is then updated with

$$F_n(\mathbf{x}_t) = F_{n-1}(\mathbf{x}_t) + \lambda \rho_n h(\mathbf{x}_t; \theta_n) \quad (8)$$

where λ which is a user defined regularisation lever termed shrinkage, and ρ_n which is the best gradient step-size found via

$$\rho_n = \operatorname{argmin}_{\rho} \sum_t L(y_t, F_n(\mathbf{x}_t) + \rho h(\mathbf{x}_t; \theta_n)) \quad (9)$$

and specifically for regression trees a different ρ_n is computed for each terminal leaf. This model fitting optimisation strategy is then based on two stages: least squares fitting of the base learner followed by the parameter optimisation according to the general loss function³⁴. Regularisation is extremely important when deploying a GBM model. Generalisation performance is achieved via the tunable shrinkage parameter (λ) which allows the user to penalise the importance of each individual learner in the overall ensemble, the number of learners (N), and row/column²⁹ sub-sampling fractions where a random subset of rows or covariates respectively are used for each training round. Row and column sub-sampling are really useful for generalisation performance when a large set of input features are available. If the entire dataset is defined as $\{y_t, \mathbf{x}_t\}_1^T$ then a random sub sample of rows of size $\tilde{T} < T$ is given by $\{y_{r(i)}, \mathbf{x}_{r(i)}\}_1^{\tilde{T}}$, where $\{r(i)\}_1^{\tilde{T}}$ is random permutation of the integers $\{1, \dots, T\}$ ³⁴. A random subset of features/columns can be defined similarly. Other tree-specific tunable hyper-parameters include the depth (or number of splits) of each tree, and the minimum number of observations per terminal leaf of the tree. It is evident that the benefits of the GBM method come at the expense of the necessary fine-tuning stages compared to for instance a linear model. These regression models were created in R via the `h2o`²⁹ software package.

3.3 | Feature Engineering

Although high-resolution numerical models may be able to predict specific weather features better temporal and/or spatial misplacement errors can in fact reduce their skill compared to coarser global models that provide a smoother representation of the same events. Figure 2 shows an example time series of the raw wind speed forecasts versus the turbine anemometer measurements. Visually, it is clear that neither GRASP or ECMWF can perfectly predict the truth, although GRASP clearly shows a higher and thus more realistic variability. However, when regarding the Pearson correlation coefficient between the observations and NWP forecasts in Table 1 (top row) it can be seen that the high frequency information actually seems to disturb the forecast signal and that smoother forecasts are of advantage. This observation raises the idea of temporally smoothing the forecasts to obtain forecasts with the same underlying signal but with some of the (potentially disturbing) high frequency information removed. This is most simply done by calculating moving averages of the raw time series

$$\tilde{x}_t = \frac{1}{M+1} \sum_{r=t-M/2}^{t+M/2} x_r \quad (10)$$

where M is an even number that controls the degree of smoothing. For this study, we tested moving averages spanning time windows from 30 sec-

	GRASP	ECMWF
Raw	0.898	0.910
Smoothed	0.905	0.915

TABLE 1 Pearson correlation coefficient of raw and smoothed forecasts with wind speed observations. For smoothing the moving average spans 500 minutes.

FIGURE 2 Example time series of raw wind speed forecasts from both NWP sources versus the turbine anemometer measurement. The GRASP and ECMWF variables used are the disk averaged and 70 metre wind speed respectively.

onds to 16 hours. Although high frequency information is disregarded, Table 1 (bottom row) shows that this smoothing increases the correlation between the observations even in the ECMWF model. The justification of smoothing in post-processing, even for ECMWF forecasts, can be alternatively explained by the fact that the NWP global models resolve processes at resolutions that are less than that of the grid they are computed on. In addition, high-resolution forecasts expose themselves to the risk of being doubly penalized when predicting the right events, but slightly misplaced in time or in space. Smoothing somewhat dampens this risk. It should be emphasised that the GRASP data retains the high resolution spatial information and one example of spatial smoothing is retained in the analysis and is defined as the spatial average of all the disk averages wind speeds in the wind farm.

Although average scores may be improved by smoothing the numerical forecasts, the high frequency content still might contain some valuable information. Therefore, moving variances are derived as

$$z_t = \frac{1}{M+1} \sum_{r=t-M/2}^{t+M/2} (x_t - \tilde{x}_t)^2 \quad (11)$$

which are used as supplementary engineered features to summarize the variability of the forecast. Additionally, a separate strategy to exploit the higher frequency content of the NWP signals is proposed based on a rolling Fast Fourier Transform (FFT). A smoothed Power Spectral Density (PSD) estimate of the transformed signal is split into a number of bands; the average, sum, range, and variance of the power in these bands is used

to inform the models. The rolling discrete Fourier transform of the time-series centred on the window M is defined as

$$X_{t,k} = \sum_{r=t-M/2}^{t+M/2} x_r e^{-i2\pi rk/M}, \quad k = 0, 1, \dots, M-1 \quad (12)$$

at frequency domain point k . The corresponding rolling estimate of the PSD at this point is

$$P_{t,k} = \frac{1}{M} |X_{t,k}|^2 \quad k = 0, 1, \dots, M-1 \quad (13)$$

and since the Fourier transform of this real valued data is symmetric, the bands which are defined to engineer features within a frequency range are split equally within the $M/2$ range. The features retained in the final models are then based on summary statistics of $P_{t,k}$ within the highest frequency band which proved to be most informative. Other transformations could potentially better track the time-varying properties of the time-series such as the Hilbert-Huang Transform³⁵, however for this proof of concept study these aspects are reserved for future work.

Therefore, the modelling strategy can be summarised as follows: to use temporally smoothed forecasts as the driving signal for the post-processed models and supplement this with engineered features that inform the model with information on variability of the NWP. As discussed, feature engineering has proven extremely successful in both wind and solar energy forecasting^{18,25,26}. However, the temporal smoothing properties have not been explored explicitly, to the best of the authors knowledge, by single input models. It should be noted that all of the engineered rolling features are calculated per issue time of the forecast because of potential step changes in the variables across at this point.

4 | RESULTS

The results of this case study are presented as firstly an analysis of the raw LES output, following an exploratory analysis where a single turbine in the farm is selected for post-processing with single input models to explicitly extract the value in smoothing. Both wind speed and wind power post-processing are considered at this stage. Next, additional features are added to the turbine power models to capitalize on the high frequency content, before finally results for the wind farm level forecasting are presented.

For the post-processing analysis, the data is partitioned into tuning and testing data by 6 and 18 months blocks respectively, where the tuning data is used to improve the GBM models. Out-of-sample forecasts are generated using 5-fold cross-validation on the testing dataset with curtailment around the wind farm flagged and removed from the forecasting exercise. The definition of prediction error $e_t = y_t - \hat{y}_t$ is the difference between the measured y_t and forecast value $\hat{y}_t$. Evaluation of the post-processed forecasts is then based on the Root Mean Square Error (RMSE) over the testing dataset of length T

$$\text{RMSE} = \sqrt{\frac{1}{T} \sum_{t=1}^T e_t^2} \quad (14)$$

and the results are generally presented via boxplots showing the sampling distribution from bootstrap score averages. Bootstrap sampling can convey the distribution of the evaluation metric via sampling the errors with replacement, then calculating the average error metric, and repeating the process k (in this instance 100) times.

The power point forecasts generated by the GBM models are constructed to minimize the quadratic loss function in this study. Therefore the models created will be optimised for the mean squared error and this motivated the use of RMSE as the lead measure of accuracy. Though other scores such as mean absolute error (MAE) could also be similarly reported, the results will not be optimised for the measure unless the models were modified to carry out a median regression. Additionally, due to the wealth of models tested in this case study the MAE results are omitted for brevity.

4.1 | LES results

We first present a number of quantities from the unprocessed LES output to verify the implementation of the wind farm in the LES model set-up. Figure 3 shows the yearly average power production per turbine from the observations and as forecast by GRASP. For the turbines on the western and southern edge of the wind farm, the yearly average production shows a good agreement between the model and the observations. The deviation is roughly within a 20 kW range, which is 2% of the average yearly production. Deeper inside the wind farm, GRASP underestimates the production, or, equivalently, overestimates the wake effects. This is also confirmed by an initial analysis looking at the RMSE of the power as exported directly from the LES (results not shown here) and forms another justification for statistical post-processing.

4.2 | Exploratory Data Analysis

In order to reduce the dimensions of the problem, the following presents results of forecasting for a single turbine in the farm, located at the north-east perimeter of the farm in terms of both wind speed and power prediction performance. Here we have a few parameters which require

FIGURE 3 Yearly average power from observations (left) and forecasted by the GRASP LES model (right)

explanation: DA indicates the disk averaged wind speed which is a weighted average of the closest LES grid point wind speeds over the rotor disk plane, and DA-all is the spatial average of these disk averaged variables over all the turbine locations.

4.2.1 | Single Input Models

To explicitly investigate the influence of smoothing the wind speed forecast on both wind power and wind speed forecasting single-input models are first considered; a separate model for each NWP wind speed forecast is trained against the measured wind speed or power time-series.

Figure 4 shows the results for the wind speed post-processing case at 70m height where 0.5 and 10 minutes indicates no smoothing for GRASP and ECMWF data respectively. It should be noted that the ECMWF is already linearly interpolated to match the temporal resolution of the power measurements. The RMSE for the ECMWF is generally much lower for wind speed across the smoothing window choices, apart from the longest rolling window lengths of >900 minutes. For both models the optimal smoothing window is around 400-500 minutes which is longer than the original temporal resolution of the ECMWF data and implies that smoothing can improve the forecast accuracy of a post-processed model derived from both traditional NWP output as well as high-resolution LES output.

For wind power, shown in Figure 5, a similar profile is found with some very important distinctions; the GRASP based models give a lower error than ECMWF across the smoothing windows, and the characteristic dip of the wind speed plots is slightly shifted and not as pronounced. This suggests that from smoothing the error improvements at the optimal smoothing window are out-with the below rated region of the power curve because the non-linear effect of the power curve would tend to amplify improvements in this key region. However for both NWP sources, smoothing NWP forecasts for use in power forecasting does have a beneficial influence at a window of around 400 minutes.

The unusual difference between the wind speed and wind power performance between the NWP sources is explained in Figure 6 which shows the performance of four comparable models conditional on the wind speed and power measurements respectively. All models compared here are using a smoothing window of 400 minutes and at 70m height. At the key range for power prediction — the medium range wind speeds — the GRASP post-processed model is more accurate and the larger wind speed prediction errors of the GRASP model are negated. It should be noted that wind speed predictions could be improved at these regions by using (for instance) splines in the regression, however this is out-with the scope of this preliminary analysis which is mainly focused on wind power.

FIGURE 4 Wind speed RMSE results with different smoothing windows at 70m height. Forecasts were derived from a truncated linear regression model with different smoothed forecasts as input.

FIGURE 5 Wind power RMSE results with different smoothing windows. Forecasts were derived from GBM models with smoothed wind speed input variables at 70m height for ECMWF and the disk averaged wind speed (DA) for GRASP

Looking at the performance in terms of the model height is also informative. Figure 7 shows that for wind speed forecasting the ECMWF gives the highest accuracy at 50 metres, with broadly similar performance across the heights compared to GRASP; it is important to emphasise that these additional ECMWF height fields are linearly interpolated from the nearest available model heights, roughly 10m, 31m, 54m, 79m and 107m. For the LES model there is a pronounced improvement around the hub height of the turbine, which signals that there is more skill in the vertical profile of the LES compared to that of ECMWF. This is expected because the resolved wake effects are most prominent at hub height.

For power predictions over the different model heights, as shown in Figure 8, there is again a reversal of roles where GRASP gives much improved forecasts at four key model heights. Notably, the disk averaged (DA) wind speed variable is the best predictor of wind power. The wind farm spatial average of the disk averaged wind speed forecasts is not the best predictor here, which suggests that for this particular turbine the high spatial resolution is providing some benefit.

FIGURE 6 RMSE of Whiffle's GRASP model and ECMWF conditioned on 5 equally populated observation bins with a 400 minute smoothing window and at a height of 70m. The shaded areas show the interquartile range from bootstrap score averages.

FIGURE 7 Wind speed RMSE at different model heights with a 500 minute smoothing window

4.2.2 | Improving the Power Forecast

Thus far, the analysis has essentially excluded the high temporal resolution content. Including the proposed rolling variance and PSD band features aims to capitalize on this available information. Figure 9 is the culmination of the exploratory analysis; it shows the progression from the raw NWP forecast, to the smoothed signal, to the smoothed forecast with wind direction, then including PSD band features, or finally rolling variance variables. For GRASP, the forecast used is the disk averaged wind speed and wind direction at 70 meters, and for ECMWF the speed and direction forecasts are based at 70 meters. The smoothing and rolling window used for both is 400 minutes.

Although these single-input models have highlighted the influence and value of temporal smoothing, the errors can be clearly further reduced by incorporating wind direction and engineered features. As shown in Figure 9 incorporating wind direction into the model is valuable for the GRASP. Whereas for the ECMWF forecasts, a significant improvement is observed when including wind direction. It is suggested that could be due to some of the directional effects of the farm, such as wake deficits, are already resolved in the high resolution model, whereas for ECMWF these effects are obviously excluded. The plot also illustrates that incremental benefit is achievable for GRASP and even ECMWF using these PSD engineered features. However, using the more simple rolling variance degrades performance across both cases.

FIGURE 8 Wind power RMSE at different model heights with a 400 minute smoothing window

FIGURE 9 Wind power RMSE where ECMWF features are at 70m height and for GRASP the disk averaged wind speed and wind direction at 70m are used for GRASP. **A** is the raw NWP signal; **B** is the smoothed NWP signal over 400 minutes; **C** is the smoothed NWP and wind direction; **D** is the smoothed NWP, PSD features, and wind direction; **E** is the smoothed NWP, rolling variance, and wind direction

Overall at this turbine, Figure 9 shows that accuracy improvements are achievable using the high resolution LES data. Additionally, it is clear that to fully utilise the high resolution temporal content of the signal more advanced models are required.

4.3 | Wind Farm Power

Investigating the overall predictive power in terms of the wind farm power can shed some light on the average performance of both datasets. This section is based on a approach whereby the wind farm power is modelled with the entire set of explanatory variables in each case, and the inherent feature selection ability of the statistical learning technique is relied upon to select the most relevant features. The three cases considered are the separate GRASP and ECMWF datasets with smoothing and features engineered as shown previously, and finally a combined model with access to all available NWP inputs and features.

The dimensionality of the data in these three cases are an issue however because of the large pool of available information. Although the boosting algorithm provides an inherent feature selection capability, empirically it is found that a more rigorous selection and feature reduction stage can

FIGURE 10 Wind farm power results

improve model performance by making them more parsimonious²⁵. The feature selection stage involves fitting a regularised GBM model with all the available inputs and retaining only the features which have the highest influence. A second and final model is then trained using these selected variables. This process removes the influence of unnecessary predictor variables where no additional valuable information is available but may have been used sparingly in the primary model training. Regularisation levers in the GBM framework provide the necessary framework for this feature selection stage and have been compared to other more familiar algorithms such as penalised regression³³.

Figure 10 shows the results of the wind farm forecasting case study. From this it is clear that, even in this context, the high resolution data is competitive with ECMWF. However the real value comes from using the information from both GRASP and the standard ECMWF forecast, which gives significant improvements in RMSE over both single dataset models. This is understandable given that the model is selecting the best of both worlds in terms of information content. Lastly, a feature selection stage is clearly shown to improve the accuracy of the 3 models tested.

5 | CONCLUSIONS

The value in post-processing ultra high resolution weather forecasts for power prediction has been shown to improve the accuracy of point forecasts at Horns Rev I. The case study evaluates the performance of wind speed and power forecasting at a single turbine in the array to first explore the dataset and then investigates the performance of overall wind farm power prediction. These predictions are compared against forecasts generated purely from the standard resolution ECMWF model.

The exploratory investigation used single-input models to characterise the value in temporal smoothing of both NWP sources which is shown to improve performance. Generally, ECMWF proved to be better at prediction of wind speeds and Whiffle's LES model GRASP superior at power prediction at this single turbine. It is revealed that this is due to the high resolution model providing more accurate wind speed forecasts in key-regions of the turbine power curve, which is expected because the wake effects resolved by the LES model are most prominent in this regime. Improving the power forecast by including wind direction and engineered features to capitalize on the high frequency content of the GRASP signals is investigated. Whilst wind direction measurably improves both forecasts further work is necessary to fully exploit the temporal information in the LES model. For wind farm power forecasting using information from both ECMWF and GRASP proved to give significant increases in accuracy, especially with a feature selection stage, compared to using information content solely from either single source, which are comparable to each other. This can be explained by the statistical learning technique model selecting the most relevant information content from a diverse pool.

Clearly the high resolution LES is shown to approximate much more closely the underlying behaviour of the wind speed signal than the benchmark NWP. However the double penalisation of spatial and temporal errors mean that average error metrics are perhaps not the best framework for evaluating these forecasts; event-based metrics and applications in ramp forecasting of wind power could be very useful for optimally leveraging the GRASP information content. Future work should consider this as well as utilising the high spatial and temporal information in a modelling framework more suitable to the data. For instance, the engineered features such as the rolling variance which quantify the variability of the signal could be more valuable in probabilistic forecasting for modelling the upper and lower ends of the distribution via quantile regression^{18,36}. Understandably, a hierarchical model where each turbine is used to generate a consistent wind farm forecast could be an optimal way of using the high

spatial content of the data. Additionally, a more in depth study focused on extracting value from the temporal content of the wind forecast signal such as deep-learning³⁷, instantaneous frequency transforms³⁵, or wavelet decomposition³⁸ techniques should be explored. In the context of this initial feasibility study it is clear that it is possible to use the high resolution data with relatively simple models to improve on conventional wind farm forecasting by combining information content from GRASP and ECMWF.

ACKNOWLEDGEMENTS

Ciaran Gilbert is supported by the University of Strathclyde's EPSRC Centre for Doctoral Training in Wind and Marine Energy Systems, grant number EP/L016680/1.

DATA STATEMENT

Due to confidentiality agreements with research collaborators, access to wind power data and high resolution weather forecasts is restricted. ECMWF numerical weather predictions used in this study are available at www.ecmwf.int for qualifying research institutions or commercial customers.

References

1. IEA Renewables 2017: Analysis and Forecasts to 2022 - Executive Summary. tech. rep., International Energy Agency; 2017. [Date Accessed] 07/05/2018: www.iea.org/renewables/.
2. Morales JM, Conejo AJ, Madsen H, Pinson P, Zugno M. *Integrating renewables in electricity markets: operational problems*. 205. Springer Science & Business Media . 2013
3. Mazzi N, Pinson P. 10 - Wind power in electricity markets and the value of forecasting. In: Kariniotakis G., ed. *Renewable Energy Forecasting* Woodhead Publishing Series in Energy. Woodhead Publishing. 2017 (pp. 259 - 278)
4. Rodrigo JS, Paredes LF, Girard R, Laquaine K, Stoffels N, Von Bremen L. 14 - The role of predictability in the investment phase of wind farms. In: Kariniotakis G., ed. *Renewable Energy Forecasting* Woodhead Publishing Series in Energy. Woodhead Publishing. 2017 (pp. 341–357)
5. Giebel G, Kariniotakis G. 3 - Wind power forecasting — a review of the state of the art. In: Kariniotakis G., ed. *Renewable Energy Forecasting* Woodhead Publishing Series in Energy. Woodhead Publishing. 2017 (pp. 59 - 109)
6. Bessa R, Möhrle C, Fundel V, et al. Towards Improved Understanding of the Applicability of Uncertainty Forecasts in the Electric Power Industry. *Energies* 2017; 10(9): 1402.
7. Giebel G, Badger J, Marti I, Louka P, Kallos G. Short-term Forecasting Using Advanced Physical Modelling - The Results of the Anemos Project Results from mesoscale, microscale and CFD modelling. In: European Wind Energy Association (EWEA); 2006.
8. Zhang Y, Wang J, Wang X. Review on probabilistic forecasting of wind power generation. *Renewable and Sustainable Energy Reviews* 2014; 32(0): 255–270.
9. Bessa RJ, Matos MA, Costa IC, et al. Reserve Setting and Steady-State Security Assessment Using Wind Power Uncertainty Forecast: A Case Study. *IEEE Transactions on Sustainable Energy* 2012; 3(4): 827–837.
10. Pinson P, Chevallier C, Kariniotakis G. Trading wind generation from short-term probabilistic forecasts of wind power. *IEEE Transaction on Power Systems* 2007; 22(3): 1148–1156.
11. Monteiro C, Bessa R, Miranda V, et al. Wind power forecasting : state-of-the-art 2009. tech. rep., Argonne National Laboratory (ANL); Argonne, IL (United States); 2009. doi: 10.2172/968212.
12. Bauer P, Thorpe A, Brunet G. The quiet revolution of numerical weather prediction. *Nature* 2015; 525(7567): 47–55.

13. Calaf M, Meneveau C, Meyers J. Large eddy simulation study of fully developed wind-turbine array boundary layers. *Physics of Fluids* 2010; 22(1): 1–16.
14. Wu YT, Porté-Agel F. Modeling turbine wakes and power losses within a wind farm using LES: An application to the Horns Rev offshore wind farm. *Renewable Energy* 2015; 75: 945–955.
15. Mehta D, Zuijlen vAH, Koren B, Holierhoek JG, Bijl H. Large Eddy Simulation of wind farm aerodynamics: A review. *Journal of Wind Engineering and Industrial Aerodynamics* 2014; 133: 1–17.
16. Schalkwijk J, Griffith EJ, Post FH, Jonker HJJ. High-performance simulations of turbulent clouds on a desktop PC. *Bulletin of the American Meteorological Society* 2012; 93(3): 307–314.
17. Schalkwijk J, Jonker HJJ, Siebesma AP, Bosveld FC. A Year-Long Large-Eddy Simulation of the Weather over Cabauw: an Overview. *Monthly Weather Review* 2015; 143: 828–844.
18. Andrade JR, Bessa RJ. Improving Renewable Energy Forecasting with a Grid of Numerical Weather Predictions. *IEEE Transactions on Sustainable Energy* 2017; 8(4): 1571–1580.
19. Drechsel S, Mayr GJ, Messner JW, Stauffer R. Wind Speeds at Heights Crucial for Wind Energy: Measurements and Verification of Forecasts. *Journal of Applied Meteorology and Climatology* 2012; 51(9): 1602–1617.
20. Heus T, Van Heerwaarden CC, Jonker HJJ, et al. Formulation of the Dutch Atmospheric Large-Eddy Simulation (DALES) and overview of its applications. *Geoscientific Model Development* 2010; 3(2): 415–444.
21. Schalkwijk J, Jonker H, Siebesma A, Van Meijgaard E. Weather forecasting using GPU-based large-Eddy simulations. *Bulletin of the American Meteorological Society* 2015; 96(5).
22. Meyers J, Meneveau C. Large Eddy Simulations of Large Wind-Turbine Arrays in the Atmospheric Boundary Layer. *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition* 2010(January): 1–10.
23. Stevens RJ, Graham J, Meneveau C. A concurrent precursor inflow method for Large Eddy Simulations and applications to finite length wind farms. *Renewable Energy* 2014; 68: 46–50.
24. Aitken ML, Banta RM, Pichugina YL, Lundquist JK. Quantifying wind turbine wake characteristics from scanning remote sensor data. *Journal of Atmospheric and Oceanic Technology* 2014; 31(4): 765–787.
25. Landry M, Erlinger TP, Patschke D, Varrichio C. Probabilistic gradient boosting machines for GEFCom2014 wind forecasting. *International Journal of Forecasting* 2016; 32(3): 1061–1066.
26. Silva L. A feature engineering approach to wind power forecasting: GEFCom 2012. *International Journal of Forecasting* 2014; 30(2): 395–401.
27. R Core Team. R: A Language and Environment for Statistical Computing. 2016.
28. Messner JW, Mayr GJ, Zeileis A. Heteroscedastic Censored and Truncated Regression with crch. *The R Journal* 2016; 8(1): 173–181.
29. The H2O.ai team. *h2o: R Interface for H2O*. 2018. R package version 3.5.1.
30. Thorarindottir TL, Gneiting T. Probabilistic Forecasts of Wind Speed: Ensemble Model Output Statistics by Using Heteroscedastic Censored Regression. *Journal of the Royal Statistical Society A* 2010; 173(2): 371–388.
31. Friedman J. Greedy function approximation: a gradient boosting machine. *Annals of Statistics* 2001; 29(5): 1189–1232.
32. Natekin A, Knoll A. Gradient boosting machines, a tutorial. *Frontiers in Neuroinformatics* 2013; 7(DEC).
33. Hastie T, Tibshirani R, Friedman JH. *The elements of statistical learning: data mining, inference, and prediction*. Springer. 2nd ed. 2009.
34. Friedman JH. Stochastic gradient boosting. *Computational Statistics & Data Analysis* 2002; 38(4): 367–378.
35. Vincent C, Giebel G, Pinson P, Madsen H. Resolving Nonstationary Spectral Information in Wind Speed Time Series Using the Hilbert–Huang Transform. *Journal of Applied Meteorology and Climatology* 2010; 49(2): 253–267.

36. Gilbert C, Browell J, McMillan D. A Hierarchical Approach to Probabilistic Wind Power Forecasting. In: 2018 IEEE International Conference on Probabilistic Methods Applied to Power Systems (PMAPS); 2018: 1–6.
37. Wang Hz, Li Gq, Wang Gb, Peng Jc, Jiang H, Liu Yt. Deep learning based ensemble approach for probabilistic wind power forecasting. *Applied energy* 2017; 188: 56–70.
38. Catalão JPD, Pousinho HMI, Mendes VMF. Short-term wind power forecasting in Portugal by neural networks and wavelet transform. *Renewable energy* 2011; 36(4): 1245–1251.