Day-ahead Wind Power Predictions at Regional Scales

Post-processing Operational Weather Forecasts with a Hybrid Neural Network

Basu, Sukanta; Watson, Simon J.; Lacoa Arends, Eric; Cheneka, Bedassa

DOI
10.1109/EEM49802.2020.9221979

Publication date
2020

Document Version
Final published version

Published in
2020 17th International Conference on the European Energy Market, EEM 2020

Day-ahead Wind Power Predictions at Regional Scales: Post-processing Operational Weather Forecasts with a Hybrid Neural Network

Sukanta Basu*, Simon J. Watson†, Eric Lacoa Arends‡, and Bedassa Cheneka†
*Faculty of Civil Engineering and Geosciences
s.basu@tudelft.nl
†Faculty of Aerospace Engineering, Wind Energy Section
‡Faculty of Electrical Engineering, Mathematics & Computer Science
Delft University of Technology
Delft, The Netherlands

***Abstract*—A hybrid neural network model, comprising of a convolutional neural network and a multilayer perceptron network, has been developed for day-ahead forecasting of regional scale wind power production. This model requires operational weather forecasts as input and also has the capability to ingest data from ensemble forecasts. Even though the training of the model requires significant computational cost, the actual forecasting can be done within a few minutes on any recent personal computer. The proposed model has demonstrated noteworthy performance at a recent international forecasting competition.**

***Index Terms*—convolutional neural network, deep learning, multilayer perceptron, numerical weather forecasting, wind energy**

I. Introduction

Over the past few decades, numerous physics-based and data-driven wind power forecasting approaches have been proposed in the literature [1]–[3]. With the advent of deep learning and other statistical learning techniques in recent years, there has been a tremendous surge in the development of more sophisticated wind power forecasting frameworks (see [4], [5], and the references therein). In order to assess the relative performance of some of these new-generation approaches, the organizing committee of the 17th International Conference on the European Energy Market (EEM20) crafted a regional scale wind power forecasting competition (https://eem20.eu/forecasting-competition/).

The participants of the EEM20 competition were challenged with a total of six forecasting tasks. For each task, the participants were given gridded meteorological data from operational weather forecasts and were asked to forecast two-months of wind power production for four price regions over Sweden. During the competition, we formulated the estimation of wind power from concurrent meteorological data as a regression problem and developed a new hybrid neural network model. Our methodology is fundamentally different from the conventional time-series-based approaches. In the following sections we elaborate on the available input data, the proposed methodology along with selected results.

978-1-7281-6919-4/20/$31.00 ©2020 IEEE

II. Description of Data

The organising committee of the EEM20 forecasting competition supplied the participants with two years (2000–2001) of wind power production and meteorological datasets for the development of machine learning-based forecast models as well as their validations. In addition, they provided the geographical locations of the wind turbines over Sweden and their basic characteristics. In the following sub-sections, we briefly elaborate on these datasets.

A. Wind Power Production Dataset

The granularity of this dataset is hourly and it includes aggregated wind power production from four price regions (called SE1, SE2, SE3, and SE4). Prior to the start of the competition, the power production data for the year 2000 were released to the participants for the creation of the initial frameworks for forecast models. Throughout the competition, a few days after the completion of each forecast task, the associated power data were also released to the participants for model validation and/or further tuning of their models.

B. Wind Turbine Dataset

This particular dataset contains a few attributes of 4004 wind turbines which were operational during 2000–2001. In the top panel of Fig. 1, their locations have been marked out. The rated capacities of these turbines vary greatly – from 10 kW to 4.2 MW. Furthermore, a significant number of wind turbines came online during the years of 2000 and 2001; the resultant changes in the total installed capacities of the four price regions are documented in the bottom panel of Fig. 1.

Notably, some of the values of this dataset appeared to be rather spurious. For example, hub-heights (and rotor diameters) of 1 m do not seem to be realistic even for micro wind turbines. Since the competition participants were not

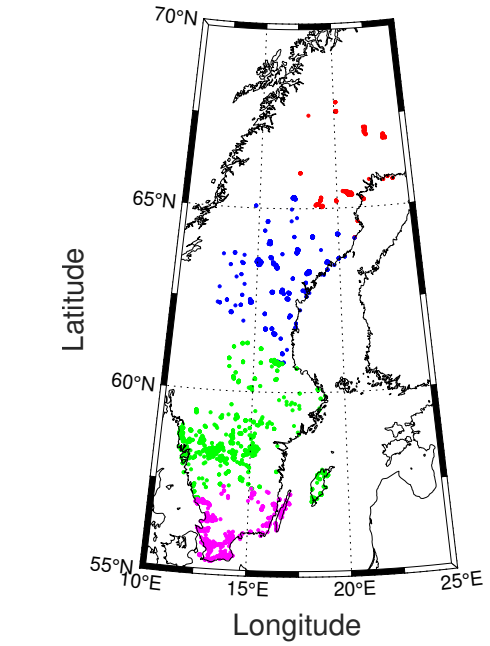

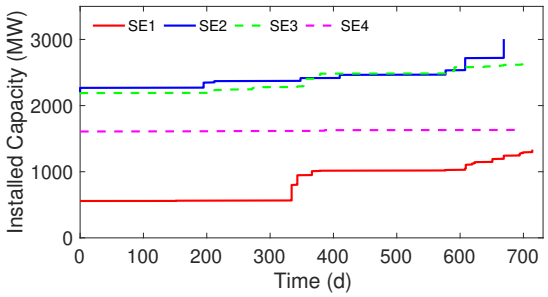

Fig. 1. The locations of the installed wind turbines are shown in the top panel. The red, blue, green, and purple points represent turbines from the price regions SE1, SE2, SE3, and SE4, respectively. In the bottom panel, the temporal changes of installed wind power production capacities in the four price regions from the start of year 2000 are documented.

allowed to use any external data sources, we did not make any attempt to correct these outliers. Furthermore, the competition organizers acknowledged other discrepancies with the turbine dataset (https://eem20.eu/forecasting-competition/):

> "This record is quality check [*sic*] as good as we could but there might be discrepancies compared to reality. For example, according to the Swedish Wind Power Association there were 4099 wind turbines in Sweden constituting 8984 MW of installed capacity. In the record provided there are only 4004 wind turbines constituting 8640 MW of installed capacity."

It is needless to say that such data quality problems may have impacted the performance of all the forecast models to some degree.

C. Numerical Weather Prediction (NWP) Dataset

Gridded fields of various meteorological variables (e.g., 10 m zonal and meridional velocity components, 2 m air temperature, wind gust, cloud cover) are included in this

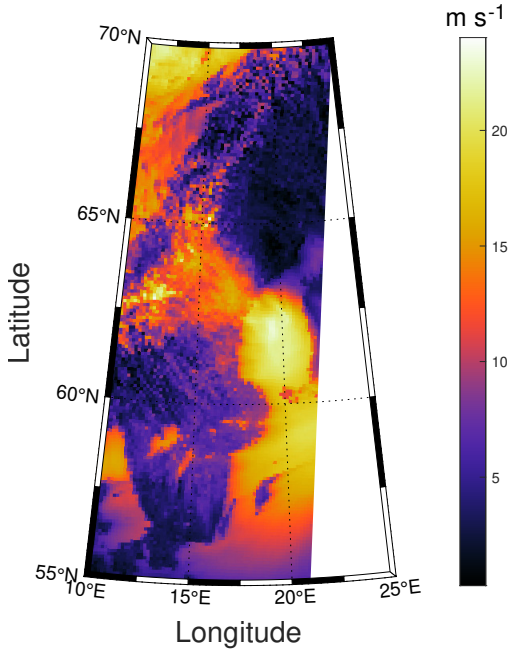

Fig. 2. An illustrative example of a numerical weather prediction model-generated spatial field of wind speeds at an elevation of 10 m above the land/sea surface. The data are extracted from ensemble member #1 and represent the wind field in our domain of interest at 18 UTC on January 7th, 2000.

weather model-generated dataset. As with the power production dataset, the organising committee released the NWP data for the year of 2000 prior to the commencement of the competition. During the competition, for each forecast task, the required NWP data were released a few days in advance. It should be noted that this specific dataset is solely used as primary input by the participants and not for model validation.

The NWP dataset was originally produced by the Norwegian Meteorological Institute (aka MET Norway) as part of their routine operational forecasting. They utilised the HARMONIE model [6] with a spatial resolution of 10 km over a computational domain covering Scandinavia and the Nordic Seas. The forecast data were output every hour. Due to the inherent chaotic nature of the atmosphere, the accuracy of the NWP models (including HARMONIE) tend to decrease with the increasing prediction horizon. To quantify the uncertainty and the predictability of weather forecasts, MET Norway created a 10-member 'ensemble' of weather forecasts. The first ensemble member (henceforth, E1) represents the control forecast which utilised the best available initial conditions (IC) and boundary conditions (BC) from a global model. The other nine forecasts were generated by perturbing the IC/BC with the aid of the Scaled Lagged Average Forecasting (SLAF) approach [7]. A detailed description of the current version of MET Norway's forecasting system can be found in [8].

An illustrative example of the gridded 10 m wind speed is shown in Fig. 2. Here, we plot the forecast values from E1 for a specific time instant (18 UTC) on January 7th, 2000. On that day, a cold front passed through Sweden and created strong localised wind ramps. The complexity of the near-surface

wind speed patterns across the four price regions are clearly noticeable in Fig. 2 and underscores the inherent challenge in wind power forecasting for all these regions using a single machine learning-based model.

III. Feature Engineering & Feature Selection

Broadly speaking, feature engineering involves extracting useful information from the raw input data in order to enhance the overall predictive capability of data-driven forecast models. There is a consensus in the machine learning community that feature engineering and feature selection are crucial steps in the model development process [9]. There exist numerous generic approaches for input data transformation and/or pruning of unwanted input. In this study, we do not opt for any sophisticated machine learning approach for feature engineering and feature selection; instead, we employ domain knowledge (i.e., meteorology and wind energy) and simple physics-based approaches for these tasks as elaborated below.

A. Vertical Extrapolation of Wind Speeds

The NWP dataset only provides zonal (U_{10}) and meridional (V_{10}) velocity components at 10 m height. For wind power forecasting, a more desirable input variable would be wind speeds at turbine hub-heights (H; unit = m). In lieu of such data, we first compute wind speeds at 10 m height (W_{10}) from U_{10} and V_{10}. Subsequently, we use the following power-law relationship:

$$W_H = W_{10}\left(\frac{H}{10}\right)^{\alpha}, \tag{1}$$

where, W_H are wind speeds at hub-height. The exponent α is called the shear exponent and it depends on surface roughness, local orography, and atmospheric stability [10].

Based on the limited NWP input data available to us, it was not possible for us to extract either a roughness map or any spatio-temporally varying stability metric (e.g., Obukhov length). Thus, we assumed α to be a constant, equal to 0.3, irrespective of location and time. Since wind farms in Sweden are often located in close proximity of forests, selection of such a high value for α is a physically meaningful choice.

B. Effects of Elevation & Temperature on Air Density

Wind power production (P) is linearly dependent on ambient density of air (ρ). Since ρ is inversely proportional to air temperature (T; unit: K), we used a non-dimensional correction factor ($273.15/T$) in our computations. Furthermore, air density more-or-less exponentially decreases with increasing elevation. To capture such an effect, we used the following simple formulation [11]:

$$\rho\,(z+H) = \rho_0 \exp\left(-\frac{z+H}{H_s}\right), \tag{2}$$

where, z is the height of the terrain, with respect to the mean sea level (MSL), at the location of a particular wind turbine. Since H represents the hub-height (above ground level), $(z+H)$ is the total height from MSL. Air density at MSL is denoted by ρ_0. The so-called scale height of density (H_s) is assumed to be equal to 8550 m [11].

C. Conversion to Wind Power

Given hub-height wind speeds (W_H), one can estimate wind power production values for a specific turbine if the manufacturer's power curve is available. Since we were not privy to such detailed information, we opted for a generic power curve depicted in Fig. 3. This normalised power curve (f_{PC}) roughly represents the IEC (International Electrotechnical Commission) class 1 turbines [12] and is used here to compute a first guess power production (P_{FG}) as follows:

$$P_{FG} = P_R f_{PC}(W_H)\left[\left(\frac{273.15}{T}\right)\exp\left(-\frac{z+H}{H_s}\right)\right], \tag{3}$$

where, P_R is the rated capacity of a wind turbine. The correction terms, in the square bracket, are described in the previous sub-section.

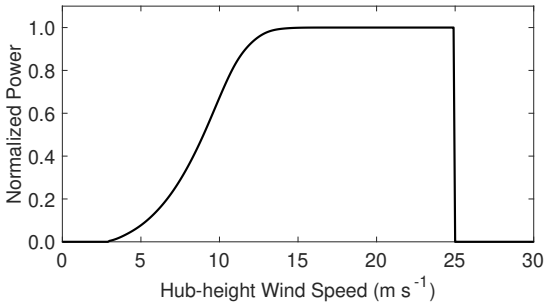

Fig. 3. Idealized power curve (f_{PC}) for IEC class 1 turbines [12]. The cut-in and cut-out wind speeds are 3 m s^{-1} and 25 m s^{-1}, respectively.

D. Gap Filling of Data

In the wind turbine dataset, 30 turbines did not have any information on hub-heights. We assumed that hub-heights and rater power capacity are proportional to each other, and in turn, used linear interpolation to estimate the missing H values based on P_R. Terrain height (z) for a few turbine locations (79 to be specific) were also missing in the dataset. We set $z = 0$ (i.e., mean sea level) as a crude approximation. Lastly, the NWP data were missing for three days: May 14$^{\text{th}}$ of 2000, September 26$^{\text{th}}$ of 2000, and July 30$^{\text{th}}$ of 2001. These dates were excluded from the forecasting competition as well as from this study.

E. Feature Selection

In our modelling approach, we only used the NWP-generated wind speeds and temperature fields. Based on wind power meteorology literature, we do not expect other available meteorological variables (e.g., wind gusts, cloud cover, relative humidity) to modulate wind power production in any significant manner. Thus, guided by our domain knowledge, we decided not to include them in our forecast model.

IV. A Hybrid Neural Network

In this study, we have developed a Hybrid model comprising of a Convolutional neural network (CNN) model and a Multilayer perceptron (MLP) model, which we dub the HCM neural network model, for wind power forecasting (refer to

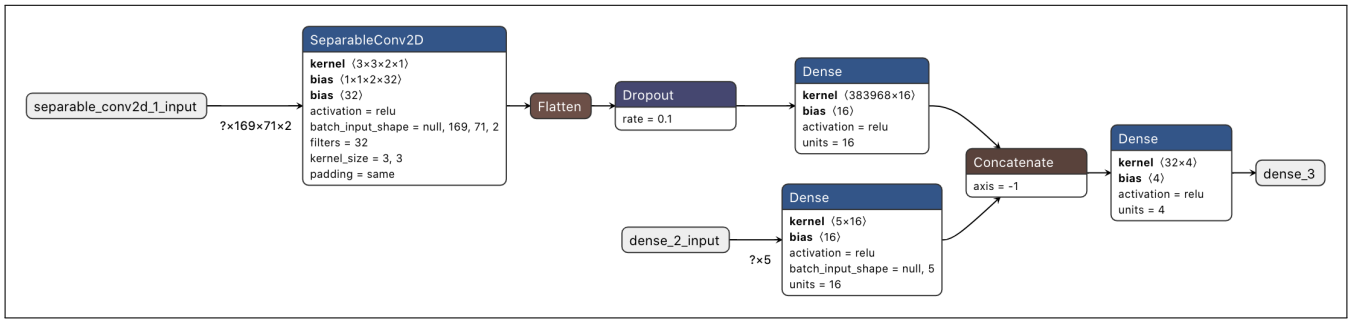

Fig. 4. A schematic of the proposed HCM neural network model. The upper branch represents the CNN model. Whereas, the lower branch corresponds to the MLP model. This schematic is generated using the Netron code [13].

Fig. 4). Our rationale behind such an unorthodox architecture is as follows: for the sake of simplicity, we wanted to have a single forecast model for the entire domain of interest instead of multiple models, each focusing on a different part of the domain. Above, in Fig. 2, it was highlighted that wind speeds over this domain exhibit significant spatial variability. From the contemporary deep learning literature involving meteorological applications [14]–[16], it became quite clear that the CNN model is the most promising candidate to capture such intricate spatial patterns. On the other hand, based on our past research in a different application arena [17], we anticipated the MLP model to better capture the temporal trends in the increasing installation capacities in the four price regions.

A. Convolutional Neural Network

As mentioned earlier, we have computed time-dependent, 2-D fields of 10 m wind speeds from the provided NWP dataset. We denote these fields as $W_{10}^{i}(x, y, t)$, where the superscript i represents the ensemble members ranging from 1 to 10. In this study, we have utilised wind speed data from the first ensemble member ($W_{10}^{1}(x, y, t)$). In addition, we have created a median wind speed field ($W_{10}^{p50}(x, y, t)$) based on the output from all the ensemble members. Henceforth, for brevity, we will simply denote these spatio-temporally varying wind speed fields as W_{10}^{1} and W_{10}^{p50} and drop their functional dependence on x, y, and t.

The spatial dimensions of both W_{10}^{1} and W_{10}^{p50} are: 169×71. These two fields are input as two separate channels to a 2-D convolutional layer (*SeparableConv2D*); the input layer is called *separable_conv2d_1_input* in Fig. 4. The *SeparableConv2D* layer consists of 32 feature maps each with filter size of 3×3. After flattening, the *SeparableConv2D* layer is connected to a dropout layer for regularization. Thereafter, the dropout layer is connected to a fully connected dense layer with 16 neurons. The rectifier linear unit (ReLU) is used as the activation function for all layers in this network.

B. Multilayer Perceptron Network

The input data generation for the MLP required a few steps:

1) for a given turbine, identify the nearest NWP grid-point to its location;
2) using the W_{10}^{1} value from that NWP grid-point, estimate W_{H}^{1} by invoking (1);
3) using the air temperature (T_{2}^{1}) value of the NWP grid-point in conjunction with the values of z, and H, calculate the density corrections;
4) compute the value of P_{FG} by using (3) along with W_{H}^{1} and the density corrections as input;
5) repeat the previous steps for all the 4004 turbines and for all time instances;
6) based on these time series, compute the aggregated P_{FG} time-series for each price region.

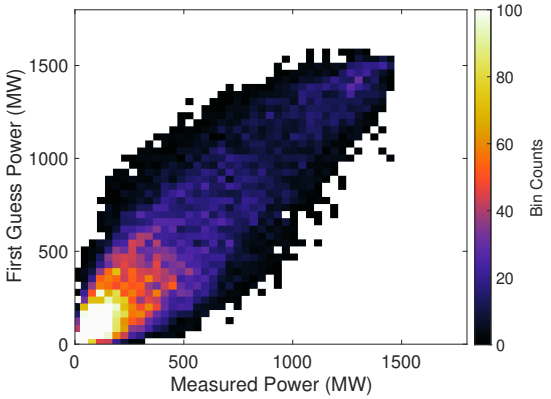

Fig. 5. Measured (P_O) versus first guess power (P_{FG}) production for the SE4 price region. These data correspond to the training period of January 1, 2000 – October 31, 2001.

The first guess power production (P_{FG}) time-series from the four price regions are used as four inputs of the MLP network. A comparison of P_{FG} against measured power production (P_O) is shown in Fig. 5. In order to capture the diurnal cycle of wind speeds, time (UTC) is used as the 5th input for the MLP network. These five input nodes, called *dense_2_input* in Fig. 4, are connected to a fully connected dense layer with 16 neurons. Similar to the CNN model, the ReLu activation function is used here as well.

C. Combination of CNN & MLP

The last layers of the CNN model and the MLP model are combined together before connecting to the output layer. By construction, both the CNN and MLP models have equal numbers of nodes (i.e., 16) in the final layer so that they contribute somewhat equally in the forecast process. The output layer contains four nodes, each containing the measured wind power production values from one of the price regions.

TABLE I
SUMMARY OF WIND POWER FORECASTS USING THE HCM NEURAL NETWORK MODEL

Runs	Type	CNN Channels	Training Period	Forecast Period	Pinball Score (MW)
C4	Competition	W_{10}^{1}, W_{10}^{p50}	January 1, 2000 – June 30, 2001	July 1 – August 31, 2001	41.96
R4-A	Reforecast	W_{10}^{1}	January 1, 2000 – June 30, 2001	July 1 – August 31, 2001	40.77
R4-B	Reforecast	W_{10}^{1}, W_{10}^{p50}	January 1, 2000 – December 31, 2000	July 1 – August 31, 2001	45.84
R4-C	Reforecast	W_{10}^{1}	January 1, 2000 – December 31, 2000	July 1 – August 31, 2001	45.71
C5	Competition	W_{10}^{1}, W_{10}^{p50}	January 1, 2000 – August 31, 2001	September 1 – October 31, 2001	51.77
R5-A	Reforecast	W_{10}^{1}	January 1, 2000 – August 31, 2001	September 1 – October 31, 2001	50.98
R5-B	Reforecast	W_{10}^{1}, W_{10}^{p50}	January 1, 2000 – December 31, 2000	September 1 – October 31, 2001	59.82
R5-C	Reforecast	W_{10}^{1}	January 1, 2000 – December 31, 2000	September 1 – October 31, 2001	58.87
C6	Competition	W_{10}^{1}, W_{10}^{p50}	January 1, 2000 – October 31, 2001	November 1 – December 31, 2001	66.28
R6-A	Reforecast	W_{10}^{1}	January 1, 2000 – October 31, 2001	November 1 – December 31, 2001	63.43
R6-B	Reforecast	W_{10}^{1}, W_{10}^{p50}	January 1, 2000 – December 31, 2000	November 1 – December 31, 2001	77.37
R6-C	Reforecast	W_{10}^{1}	January 1, 2000 – December 31, 2000	November 1 – December 31, 2001	71.74

The HCM network uses Keras [18] and Tensorflow [19] packages. Due to the space limitation, we are unable to provide more technical details of the CNN and MLP models. Interested readers are encouraged to check the textbook [20] for in-depth theoretical understanding. For more practical knowledge, please refer to [21] and [22].

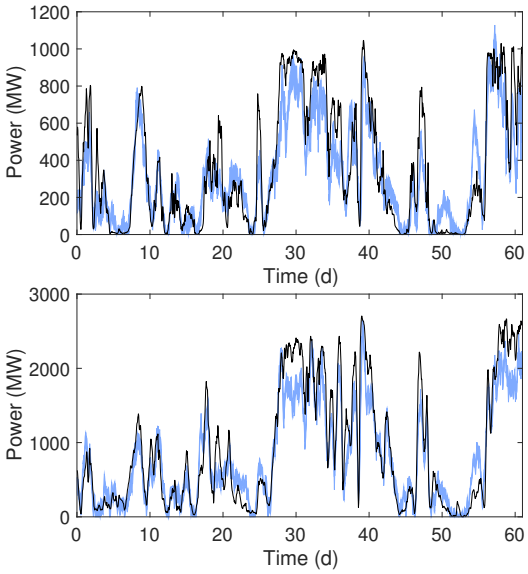

Fig. 6. Measured (black line) and the HCM model-based forecast (blue shaded area) time-series of wind power productions for SE1 (top panel), and SE2 (bottom panel) price regions. The blue shaded area includes 10th to 90th quantiles of the forecasts. These forecasts were made during the competition as part of Task #6 covering November 1st – December 31st, 2001.

D. Training, Validation, and Testing

During the training phase, the total training data are split into two parts. The first 90% of the data are directly used for optimising the network parameters. The HCM network contains a total of 6.14 million parameters. The popular ADAM algorithm is used as an optimizer in conjunction with the mean absolute error as a loss function. The last 10% of the data are used for the computation of validation errors. Via monitoring the evolution of this error, an early stopping strategy is used to avoid overfitting of the network. Typically, 40–50 epochs are needed for convergence.

The forecasting competition required various quantiles of wind power productions as output. For this purpose, we made use of a Monte Carlo-type approach and created one hundred separately tuned HCM networks for each forecast task. Note that each HCM network has identical architecture; however, the sequential ordering of the training data was randomised for all of them. As a result, each HCM network converged with different sets of optimized parameters. Using one hundred HCM networks, we forecast one hundred realizations of the wind power production values. Based on these realizations, we compute the required quantiles.

V. RESULTS: COMPETITION & REFORECASTS

During the EEM20 forecasting competition, the HCM model fared reasonably well and achieved an overall third rank. Two representative forecasts by the HCM model are shown in Fig. 6. After the competition, we are making several reforecasts and trying to further improve the model. Some of the results from these reforecasts, as well as from the competition, are tabulated in Table I.

Following the protocols of the EEM20 competition, the pinball loss function ([23], [24]) is utilized to compute the forecast errors. For each forecasting task, the errors are first computed for each price region separately. Then, they are averaged and reported on the last column of Table I. By inter-comparing the results from various runs in Table I, the positive impacts of input data from longer training periods become clear. It also becomes evident that the addition of W_{10}^{p50} as a CNN input channel deteriorates the model performance. This outcome is somewhat counter-intuitive and needs further investigation.

VI. CONCLUDING REMARKS

At the initial stages of the EEM20 forecasting competition, we were making use of the wind speed values from all the ten ensemble members. However, to reduce computational costs, we later decided to only use W_{10}^{1} and W_{10}^{p50} fields as inputs. Validation errors from a few trial runs were used for the identification of the 'best' ensemble member (i.e., E1). In our future work, we will investigate if wind data from all the ensemble members can improve the forecast scores.

Currently, we are also in the process of coupling the HCM network with the ERA5 reanalysis data [25]. In comparison to the NWP forecasts, the reanalysis data tend to be more accurate due to extensive assimilation of diverse observational data. By using such accurate meteorological data as input, we hope to disentangle the input error from the HCM modelling error.

Code and Data Availability

All the Jupyter notebooks and forecasts generated during the EEM20 competitions are publicly available at https://doi.org/10.5281/zenodo.3987741. We will also upload the HCM model-generated reforecast data on to this repository.

References

[1] M. Lange and U. Focken, *Physical Approach to Short-term Wind Power Prediction*. Springer, 2006, 208 pp.

[2] A. Costa, A. Crespo, J. Navarro, G. Lizcano, H. Madsen, and E. Feitosa, "A review on the young history of the wind power short-term prediction," *Renew. Sust. Energ. Rev.*, vol. 12, pp. 1725–1744, 2008.

[3] G. Giebel, R. Brownsword, G. Kariniotakis, M. Denhard, and C. Draxl, "State-of-the-art in short-term prediction of wind power: A literature reivew," ANEMOS. plus, Tech. Rep., 2011.

[4] J. Manero, J. Béjar, and U. Cortés, ""Dust in the wind...", deep learning application to wind energy time series forecasting," *Energies*, vol. 12, p. 2385, 2019.

[5] J. Zhang, J. Yan, D. Infield, Y. Liu, and F. Lien, "Short-term forecasting and uncertainty analysis of wind turbine power based on long short-term memory network and Gaussian mixture model," *Appl. Energy*, vol. 241, pp. 229–244, 2019.

[6] L. Bengtsson, U. Andrae, T. Aspelien, Y. Batrak, J. Calvo, W. de Rooy, E. Gleeson, B. Hansen-Sass, M. Homleid, M. Hortal *et al.*, "The HARMONIE–AROME model configuration in the ALADIN–HIRLAM NWP system," *Mon. Wea. Rev.*, vol. 145, pp. 1919–1935, 2017.

[7] E. Kalnay, "Historical perspective: earlier ensembles and forecasting forecast skill," *Q. J. Roy. Meteorol. Soc.*, vol. 145, pp. 25–34, 2019.

[8] I.-L. Frogner, A. T. Singleton, M. Ø. Køltzow, and U. Andrae, "Convection-permitting ensembles: challenges related to their design and use," *Q. J. Roy. Meteorol. Soc.*, vol. 145, pp. 90–106, 2019.

[9] A. Casari and A. Zheng, *Feature Engineering for Machine Learning*. O'Reilly Media, Inc., 2018, 200 pp.

[10] B. Storm and S. Basu, "The WRF model forecast-derived low-level wind shear climatology over the united states great plains," *Energies*, vol. 3, pp. 258–276, 2010.

[11] R. Stull, *Meteorology for Scientists and Engineers*. Brooks/Cole, 2000, 502 pp.

[12] J. King, A. Clifton, and B. Hodge, "Validation of power output for the WIND Toolkit," National Renewable Energy Lab, Golden, CO, United States, Tech. Rep. NREL/TP-5D00-61714, 2014.

[13] L. Roeder. (2020) Netron: visualizer for neural network, deep learning and machine learning models. [Online]. Available: https://github.com/lutzroeder/netron

[14] J. Zhang, P. Liu, F. Zhang, and Q. Song, "Cloudnet: Ground-based cloud classification with deep convolutional neural network," *Geophys Res Lett*, vol. 45, pp. 8665–8672, 2018.

[15] M. Sadeghi, A. A. Asanjan, M. Faridzad, P. Nguyen, K. Hsu, S. Sorooshian, and D. Braithwaite, "Persiann-cnn: Precipitation estimation from remotely sensed information using artificial neural networks–convolutional neural networks," *J. Hydrometeorol.*, vol. 20, pp. 2273–2289, 2019.

[16] W. E. Chapman, A. C. Subramanian, L. Delle Monache, S. P. Xie, and F. M. Ralph, "Improving atmospheric river forecasts with machine learning," *Geophys. Res. Lett.*, vol. 46, pp. 10 627–10 635, 2019.

[17] Y. Wang and S. Basu, "Using an artificial neural network approach to estimate surface-layer optical turbulence at Mauna Loa, Hawaii," *Opt. Lett.*, vol. 41, pp. 2334–2337, 2016.

[18] F. Chollet *et al.*, "Keras," https://keras.io, 2015.

[19] A. Martín *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: https://www.tensorflow.org/

[20] C. C. Aggarwal, *Neural Networks and Deep Learning: A Textbook*. Springer, 2018, 497 pp.

[21] F. Chollet, *Deep Learning with Python*. Manning Publications Co., 2018, 361 pp.

[22] J. Brownlee, *Deep Learning with Python*. Machine Learning Mastery, 2019, 257 pp.

[23] I. Steinwart and A. Christmann, "Estimating conditional quantiles with the help of the pinball loss," *Bernoulli*, vol. 17, pp. 211–225, 2011.

[24] T. Hong, P. Pinson, S. Fan, H. Zareipour, A. Troccoli, and R. J. Hyndman, "Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond," *Int. J. Forecasting*, vol. 32, pp. 896–913, 2016.

[25] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers *et al.*, "The ERA5 global reanalysis," *Q. J. Roy. Meteorol. Soc.*, vol. 146, pp. 1999–2049, 2020.